Federal Reserve Bank of Minneapolis
Research Department Staff Report 413

June 2012

Prices Are Sticky After All*

Patrick J. Kehoe

Federal Reserve Bank of Minneapolis, University of Minnesota
and Princeton University

Virgiliu Midrigan

Federal Reserve Bank of Minneapolis and New York University

ABSTRACT

Recent studies say prices change about every four months. Economists have interpreted this high frequency as evidence against the importance of sticky prices for the monetary transmission mechanism. Theory implies that this interpretation is correct if most price changes are regular, but not if most are temporary, as in the data. Temporary changes have a striking feature: after such a change, the nominal price tends to return exactly to its pre-existing level. We study versions of the Calvo model and the menu cost model which replicate this feature. Both models imply that temporary changes cannot offset monetary shocks well, whereas regular changes can. Since regular prices are much stickier than temporary ones, our models, in which prices change as frequently as in the micro data, predict that the aggregate price level is as sticky as in a standard Calvo model or a standard menu cost model in which micro level prices change about once a year. In this sense, prices are sticky after all.

*This paper is a greatly revised version of earlier drafts titled "Sales and the Real Effects of Monetary Policy" and "Temporary Price Changes and the Real Effects of Monetary Policy." We thank Kathy Rolfe and Joan Gieseke for excellent editorial assistance. Kehoe thanks the National Science Foundation for financial support.

A widely held view in macroeconomics is that monetary policy can be effective primarily because aggregate prices are sticky; when monetary policy changes, the aggregate price level cannot respond quickly enough to offset the intended real effects. This price stickiness is clearly at the heart of the widely used New Keynesian analysis. In standard New Keynesian models of both the Calvo and the menu cost varieties, the degree of aggregate price stickiness is determined by the frequency of price changes at the micro level: if individual good prices change rarely, then the aggregate price level is highly sticky and cannot offset monetary shocks, whereas if good prices change often, then the aggregate price level is not sticky and can.

Until recently, micro level prices have been assumed to be quite sticky—changing relatively infrequently, only about once a year; hence, aggregate prices have been assumed to be highly sticky. Recently, however, researchers have examined large micro price data series and determined that individual good prices change much more frequently than previously thought, about once every 4.3 months (Bils and Klenow 2004). Thus, according to these studies prices are quite flexible at the micro level. Interpreted through the lens of the standard New Keynesian models, this evidence implies that aggregate prices are quite flexible too.

We dispute this interpretation. Although it follows logically from standard New Keynesian models, those models are grossly inconsistent with the pattern of price changes in the micro data. We build simple extensions of both the Calvo model and the standard menu cost model that are consistent with the micro data and show that in these models aggregate prices are as sticky as in a standard menu cost model in which micro level prices change about once a year.

The major inconsistency between the standard New Keynesian models and the data is the models' inability to simultaneously account for the high- and low-frequency patterns of price variation that we document using monthly price data from the U.S. Bureau of Labor Statistics (BLS). At high frequencies, prices often temporarily move away from a slow-moving trend line called the *regular price*, but after such a *temporary price* change the nominal price often returns exactly to its pre-existing level. These distinctive features imply that even though an individual price series has a great deal of *high-frequency price flexibility* (the actual price changes frequently), the series also has a great deal of *low-frequency price stickiness* (the regular price changes infrequently).[1]

Standard New Keynesian models of both the Calvo and menu cost varieties have only one type of price change and thus have no hope of generating this feature of the data. In particular, these models generate either highly flexible prices at both high and low frequencies or highly sticky prices at both frequencies. What they cannot generate is what we see in the micro data: very flexible prices at high frequencies and very sticky prices at low frequencies.

To remedy this deficiency of standard New Keynesian models, we need a theory of why firms temporarily change their prices. Such theories can be found in the models of sales from industrial organization. Unfortunately, these theories are about real prices and, hence, cannot explain the striking feature of temporary price changes: after a temporary price change, the nominal price often returns exactly to the nominal pre-existing price. Moreover, these theories cannot explain the temporary price increases often seen in the data.[2]

The models we study, while simple, overcome the shortcomings of both the standard New Keynesian models as well as the models of sales from industrial organization. We extend the Calvo model and the standard menu cost model by allowing firms to temporarily deviate from a sticky pre-existing price. We quantify these models and show that they reproduce the empirical micro pattern of regular and temporary price changes.

We then show that these models imply that the aggregate price level responds slowly to monetary shocks. This result is driven by the distinctive features of temporary micro price changes. In the models prices change frequently, but most of those changes reflect temporary deviations from a much stickier regular price. When a firm changes its price temporarily in a given period because of an idiosyncratic shock, it is also able to react to changes in monetary policy. These responses are, however, short-lived. And whenever the price returns to the old price, it no longer reflects the change in monetary policy. For this reason, even though micro prices change frequently, the aggregate price level is sticky. Our key insight

[1] In terms of documenting this basic pattern in the data, an important reference is Nakamura and Steinsson (2008), who focus on temporary price decreases (or sales) and show that sales price changes account for the bulk of all price changes in the data. They also show that sales price changes are more transient than regular price changes and tend to return to the original level following a sale. For a survey of this literature, see Klenow and Malin (2010).

[2] See, for example, models based on demand uncertainty (Lazear 1986), thick-market externalities (Warner and Barsky 1995), loss-leader models of advertising (Chevalier, Kashyap, and Rossi 2003), intertemporal price discrimination (Sobel 1984).

is that what matters for how the aggregate price level responds to low-frequency changes in monetary policy is the degree of low-frequency micro price stickiness. Since the micro data have substantial low-frequency price stickiness, the aggregate price level is sticky as well.

Our result has implications for the debate between Bils and Klenow (2004) and Nakamura and Steinsson (2008) on the stickiness of prices. Bils and Klenow (2004) find that when they leave sales in their data prices change often, once every 4.3 months and argue that prices are fairly flexible. Nakamura and Steinsson study the same data and show that once temporary price cuts are removed, prices change infrequently, about every 7–11 months and argue that prices are fairly sticky.

The rationalization suggested by Bils and Klenow (2004, p. 955) for leaving sales in the data is that "temporary sales represent a true form of price flexibility that should not be filtered out, say because the duration and the magnitude of temporary sales respond to shocks." The argument for removing temporary price cuts is that they are somehow special and, to a rough approximation, can be ignored when determining the amount of price stickiness in the data. For example, Nakamura and Steinsson (2008, p. 1417) suggest that "some types of sales may be orthogonal to macroeconomic conditions."

We use economic theory to settle this debate. We begin with a simple extension of the Calvo model to make our point because it is simple and is viewed as the workhorse New Keynesian model. We go on to show that our result is robust to explicitly introducing menu costs to changing prices, rather than the more reduced-form Calvo approach.

In both models, the assumptions we make on the technologies for changing prices are purposefully engineered to allow the model to reproduce the observed pattern of micro price changes. In particular, we assume that firms set two prices: a list price and an actual transactions (posted) price, and face frictions for changing the list price and for having the posted price differ from the list price. In the Calvo model, these frictions are that the list price can only be changed at certain random dates and that the posted price can only differ from the list price at other random dates. In the menu cost model these frictions are menu costs of changing the list price and for charging a posted price other than the list price.

The resulting models, though simple, are broadly consistent with some aspects of the pricing practices of actual firms. In particular, Zbaracki et al. (2004) and Zbaracki, Bergen and Levy (2007) provide evidence that pricing is done at two levels: upper-level managers (at headquarters) set list prices, while lower-level managers (at stores) choose the actual transaction (posted) prices. These researchers find that lower-level managers face constraints in their ability to post a price that departs from the list price set by the upper-level managers. We think of our models as capturing this two-level decision-making process in a simple, reduced-form way.

Consider first our extension of the Calvo model. In the standard Calvo model a fraction of firms are allowed to permanently reset their list price in any given period and cannot deviate from this price. We extend this model by also allowing a fraction of firms to temporarily deviate from their list price in any given period. We show that this simple one-parameter extension of the standard Calvo model can account for the pattern of high- and low-frequency price stickiness in the data. We show that even though prices change frequently at the micro-level, the model predicts substantial amounts of aggregate price stickiness. This extension is so simple that it can be trivially embedded in the vast array of applied New Keynesian models that are frequently used for policy analysis.

In an important paper Eichenbaum, Jaimovich, and Rebelo (2011), henceforth EJR, take issue with Calvo model. They argue that the Calvo model is inconsistent with key features of the micro data. In particular, EJR carefully document that micro data on prices and costs show sharp evidence of the type of state-dependence in prices that only menu cost models deliver. Briefly, EJR show that prices typically change only when costs change and that prices are much more likely to change the farther away the actual price is from the desired price.

EJR go on to raise critical issues for standard menu cost models as well. They do so by arguing that these models have their own failings with respect to the data. First, in the data prices are more volatile than costs and nearly all prices are associated with cost changes. Standard models cannot generate both of these features simultaneously. Second, standard menu cost models cannot generate the type of high- and low-frequency price variation observed in the data. They argue that an important challenge for macroeconomists is to build menu cost models consistent with these facts.

Our extension of the standard menu cost model addresses the EJR challenge. In particular, we show that our extension can account for all of the features of the data that they document. We extend the standard menu cost model, in which changing a list price entails a fixed cost, by adding the option of paying a separate fixed cost and temporarily charging a posted price other than the list price.

In addition to responding to the EJR challenge, our menu cost model can also address the arguments of those who claim that allowing for temporary price changes can greatly diminish the real effects of money shocks. The first part of their argument is that if the timing of temporary price changes can respond to money shocks, such price changes will, perhaps greatly, increase the flexibility of aggregate prices. The second part is that since in the data a disproportionate amount of goods is sold during periods with temporary price changes, these periods are disproportionately important in allowing for aggregate price flexibility. Our menu cost model incorporates the two mechanisms present in these arguments. Nevertheless, we show that even though prices change frequently at the micro-level, the model predicts substantial amounts of aggregate price stickiness.

On the empirical side, our work here is most closely related to that of Bils and Klenow (2004) and Nakamura and Steinsson (2008). One distinction between our work and these authors' work is that we document the patterns of all temporary price changes, both increases and decreases instead of restricting attention to only the price decreases. We note that once we filter out such changes, our regular price series has a duration of 14.5 months, which is significantly longer than the 7-11 month duration found by previous researchers. The reason for this difference is that those researchers identify temporary price increases as regular price changes while we do not.

On the theory side, Guimarães and Sheedy (2011) and Head, Liu, Menzio, and Wright (2011), offer an alternative explanation for the pattern of price changes in the data arising from firms pursuing mixed-price strategies, along the lines of Varian (1980) and Burdett and Judd (1983). While elegant, these models do not attempt to address the EJR challenge to sticky price models. Finally, Rotemberg (2011) offers another explanation for why temporary prices return to their previous level. His work shows how costs to the firm of changing list prices—costs that act similarly to menu costs—can arise from the preferences of consumers.

1. The Pattern of Price Changes in the U.S. Data

We begin by documenting how prices change in the BLS monthly data set which represents about 70% of U.S. consumer expenditure. Here we describe several facts that we see in these data. These facts help clarify the distinction between temporary and regular price changes and illustrate their properties. We will later use these facts to motivate our model.

A. The Data Set

The data set we study is the CPI Research Database constructed by the BLS and used by Nakamura and Steinsson (2008). This data set contains prices for thousands of goods and services collected monthly by the BLS for the purpose of constructing the consumer price index (CPI) and covers about 70% of U.S. consumer expenditures.

B. Categories of Price Changes

To identify a pattern of price changes in the data, we wrote a simple algorithm that categorizes each change as either *temporary* or *regular*. We define for each product an artificial series called a *regular* price series. This price is essentially a running mode of the original series. Given this series, every price change that is a deviation from the regular price series is defined as *temporary*, whereas every price change that coincides with a change in the regular price is defined as *regular*.

An intuitive way to think about our analysis is to imagine that at any point in time, every product has an existing regular price that may experience two types of changes: *temporary* changes, in which the price briefly moves away from and then back to the regular price, and much more persistent *regular* changes, which are changes in the regular price itself. Our algorithm is based on the idea that a price is *regular* if the store charges it frequently in a window of time adjacent to that observation. The *regular* price is thus equal to the modal price in any given window surrounding a particular period, provided the modal price is used sufficiently often in that window. We set the window to 5 months. The algorithm is somewhat involved, so we relegate a formal description to the appendix.

Our algorithm differs from the one employed Bils and Klenow (2004), Klenow and Kryvtsov (2008), and Nakamura and Steinsson (2008) in that we treat temporary price increases and temporary price decreases symmetrically. All of these researchers construct their regular price series after removing sales from the data where sales are marked as such by the BLS[3]. Hence, by construction, these researchers only filter out temporary price decreases and, hence, treat temporary price increases as regular price changes. (For a notable exception to this work that also treats temporary increases and decreases symmetrically, see the work of EJR who study price and cost data for one firm.)

C. The Facts

Table 1 reports statistics summarizing the facts about price changes that result from applying our algorithm. These statistics are revenue-weighted averages of the corresponding statistics at the level of product categories.

We highlight several features of the data that motivate our model. First, prices change often, 22% of all prices change every month, so the average duration is 4.5 months, and most price changes are temporary (72%). Second, regular prices, in contrast, change rather infrequently: only 6.9% change per month so they have an average duration of about 14.5 months.[4] Third, the fraction of periods at which a price equals the temporary price is 10%.

We interpret these facts as implying that most price changes are temporary deviations from a slow-moving trend given by the regular price. Thus, the data shows a great deal of *high frequency price flexibility* in the presence of substantial *low frequency price stickiness*. Of course, our notion of slow-moving trend depends on the algorithm we use to define a regular price. We find it comforting that if we use a simple alternative measure of trend, namely the annual mode, we see a similar pattern: about 75% of all prices are at their annual mode.

2. A Calvo Model with Temporary Price Changes

We now build a Calvo model with temporary price changes and use it to study the relationship between the frequency of micro price changes and the degree of aggregate price

[3] Nakamura and Steinsson (2008) explain that in practice, the BLS denotes a price as a sales price when there is a "sale" sign next to the price when it is collected. In a robustness section, Nakamura and Steinsson (2008) also discuss an algorithm that defines sales prices as V-shaped declines in prices.

[4] Note that this duration of 14.5 months is higher than the corresponding 7–11 month number of Nakamura and Steinsson (2008), primarily because our algorithm takes out temporary price increases as well as temporary price decreases, or sales, that Nakamura and Steinsson focus on. Hence, our regular price series has fewer changes than the one computed by Nakamura and Steinsson (2008).

stickiness. Here, we describe the model, quantify it, and demonstrate that it does a much better job of reproducing the pattern of changes in the data than the standard Calvo model does.

Our *benchmark* model is a simple extension of the standard Calvo model. Recall that in the Calvo model there are two possibilities for any given period t: with probability α a firm can change its list price and with probability $1-\alpha$ the firm must charge the pre-existing list price P_{Lt-1}. Either way the firm always sells at its list price.

To account for the pattern of high- and low-frequency price stickiness in the data, we make a simple one-parameter modification to the technology for price adjustment. We now assume three possibilities for a given period t: with probability α_L a firm can change its list price, with probability α_T a firm can charge any price P_{Tt} that it wants (for that period only), and with probability $1-\alpha_L-\alpha_T$ the firm must charge the pre-existing list price P_{Lt-1}. Note that this simple modification allows firms to have temporary deviations from its current list price and nests the standard Calvo model as a special case (with $\alpha_T = 0$).

These assumptions are motivated in part by the work of Zbaracki et al. (2004) on the pricing practices of firms. We think of the list price as the price set by the upper-level manager and the posted price as the price actually charged to the consumer. The posted price will equal the list price unless the lower-level manager is (randomly) allowed to make a temporary deviation.

A. Setup

Formally, we study a monetary economy populated by a large number of infinitely lived consumers, firms, and a government. In each time period t, this economy experiences one of finitely many events $s_{t.}$ We denote by $s^t = (s_0, \ldots, s_t)$ the history (or *state*) of events up through and including period t. The probability, as of period 0, of any particular history s^t is $\pi(s^t)$. The initial realization s_0 is given.

In the model, we have aggregate shocks to the economy's money supply. We assume that the (log of) money growth follows an autoregressive process of the form

(1) $\mu(s^t) = \rho_\mu \mu(s^{t-1}) + \varepsilon_\mu(s^t),$

where μ is money growth, ρ_μ is the persistence of μ, and $\varepsilon_\mu(s^t)$ is the *monetary shock*, a normally distributed i.i.d. random variable with mean 0 and standard deviation σ_μ.

Consumers and Technology

In each period t, the commodities in this economy are labor, capital, money, a continuum of intermediate goods indexed by $i \in [0,1]$, and a final good.

In this economy, consumers consume, invest, work, hold real money balances, and trade one-period state-contingent nominal bonds. The consumer problem is to choose consumption $c(s^t)$, investment $x\left(s^t\right)$, labor $l(s^t)$, nominal money balances $M(s^t)$, and a vector of bonds $\{B(s^t, s_{t+1})\}_{s_{t+1}}$ to maximize utility

$$(2) \quad \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi\left(s^t\right) U\left(c\left(s^t\right), l\left(s^t\right), \frac{M(s^t)}{P(s^t)}\right)$$

subject to a budget constraint

$$(3) \quad P\left(s^t\right)\left[c(s^t) + x\left(s^t\right) + \frac{\xi}{2}\left(\frac{x\left(s^t\right)}{k(s^{t-1})} - \delta\right)^2 k(s^{t-1})\right] + M\left(s^t\right) + \sum_{s^{t+1}} Q\left(s^{t+1}|s^t\right) B\left(s^{t+1}\right)$$

$$\leqslant W\left(s^t\right) l\left(s^t\right) + \Pi\left(s^t\right) + M\left(s^{t-1}\right) + B\left(s^t\right) + R\left(s^t\right) k\left(s^t\right)$$

where $P(s^t)$ is the price of the final good, $W(s^t)$ is the nominal wage, $\Pi(s^t)$ is nominal profits, and $R(s^t)$ is the rental rate on capital. Here, $B(s^{t+1})$ denotes the consumers' holdings of such a bond purchased in period t and state s^t with payoffs contingent on some particular state s^{t+1} in $t+1$ and $Q(s^{t+1}|s^t)$ denotes the price of this bond.

Consider, next, the technology for the intermediate good producers. The producer of intermediate good i produces output $y_i(s^t)$ using capital $k_i\left(s^t\right)$, labor $l_i(s^t)$, and materials $n_i\left(s^t\right)$ according to

$$(4) \quad y_i(s^t) = \left[k_i\left(s^t\right)^\alpha l_i(s^t)^{1-\alpha}\right]^\nu n_i\left(s^t\right)^{1-\nu}.$$

This technology implies that an intermediate good firm faces a nominal unit cost of production

$$(5)\quad V\left(s^t\right)=\psi\left(R\left(s^t\right)^{\alpha}W\left(s^t\right)^{1-\alpha}\right)^{\nu}P\left(s^t\right)^{1-\nu}$$

where ψ is a constant. These firms are monopolistically competitive and we describe their problem below.

Next, a competitive final good sector combines varieties of the intermediate goods into a final good which is used for consumption, investment, and materials according to

$$(6)\quad y\left(s^t\right)=\left(\int_0^1 y_i\left(s^t\right)^{\frac{\theta-1}{\theta}}di\right)^{\frac{\theta}{\theta-1}},$$

where θ is the elasticity of substitution across intermediate inputs. The resource constraint for final goods is

$$(7)\quad c(s^t)+x(s^t)+\frac{\xi}{2}\left(\frac{x\left(s^t\right)}{k(s^{t-1})}-\delta\right)^2 k(s^{t-1})+\int n_i\left(s^t\right)di\leq y\left(s^t\right).$$

A final good firm chooses the intermediates inputs $\{y_i(s^t)\}$ to maximize

$$P(s^t)y(s^t)-\int P_i(s^t)y_i(s^t)di$$

subject to (6). The solution to this problem gives the demand for intermediate good i

$$(8)\quad y_i(s^t)=\left(\frac{P_i(s^t)}{P(s^t)}\right)^{-\theta}y(s^t)$$

and the zero profits condition implies that

$$P(s^t)=\left(\int_0^1 P_i\left(s^t\right)^{1-\theta}di\right)^{\frac{1}{1-\theta}}$$

The Intermediate Good Firm Problem

The period profits of an intermediate goods firm that charges a price $P_i(s^t)$ is given by $(P_i(s^t)-V(s^t))y_i(s^t)$ where $y_i(s^t)$ is given by (8).

Consider the problem of an intermediate goods firm that in period t has a pre-existing list price $P_{Li}(s^{t-1})$. There are three possibilities. With probability $1-\alpha_L-\alpha_T$ this firm has to charge its pre-existing list prices, $P_i(s^t) = P_{Li}(s^{t-1})$. With probability α_T the firm can charge any price in that particular period and hence charges the static optimal price, namely the solution to

$$\max_P (P - V(s^t)) y_i(s^t)$$

subject to (8). We denote the solution to this problem, referred to as the *temporary* price, as

$$P_{Ti}(s^t) = \frac{\theta}{\theta-1} V(s^t).$$

Finally, with probability α_L the firm can change its list price and hence chooses its list price P to maximize the present value of its future stream of profits at all dates and states at which that price is still in effect, namely

$$(P - V(s^t)) y_i(s^t) + \sum_{r=t+1}^{\infty} \sum_{s^r} Q\left(s^r | s^t\right) (1-\alpha_L)^{r-(t+1)} (1-\alpha_T-\alpha_L) \left[P - V\left(s^r\right)\right] y_i\left(s^r\right)$$

where for all $r \geq t$, $y_i(s^r) = \left(P/P(s^r)\right)^{-\theta} y(s^r)$. Taking the first order conditions, normalizing all nominal variables by the money supply, log-linearizing, and quasi-differencing gives that the reset list price is

$$(9) \quad p_{L,t}^R = (1-\alpha_L)\,\beta E_t p_{L,t+1} + \frac{1-(1-\alpha_L)\,\beta}{1-\alpha_T\beta}\left[v_t - \alpha_T \beta E_t v_{t+1}\right] + \frac{1-\alpha_T-\alpha_L}{1-\alpha_T\beta} \beta E_t g_{t,t+1}$$

where $g_{t+1} = \ln(M_{t+1}/M_t)$ is the growth rate of the money supply and lower-case variables denote log-deviations of normalized variables from the steady state. Note that when $\alpha_T = 0$, this formula reduces to the standard Calvo expression for the reset price.

The aggregate price level in log-linearized, normalized form is

$$(10) \quad p_t = \alpha_L p_{L,t}^R + \alpha_T p_{T,t} + (1-\alpha_L-\alpha_T)(\bar{p}_{L,t-1} - g_t)$$

where the average list price, $\bar{p}_{L,t}$, evolves according to

(11) $\bar{p}_{L,t} = \alpha_L p^R_{L,t} + (1 - \alpha_L)(\bar{p}_{L,t-1} - g_t)$

since a fraction α_L of firms reset their list prices and the rest do not.

B. Quantification and Prediction

We want to use the facts about price changes that we have isolated in the BLS data as the basis for our model and its evaluation. To do that, we must quantify the model. Here we describe how we choose the model's functional forms and parameter values. We then investigate whether our parsimonious model can account for the facts about prices that we have documented. We find that it can.

Functional Forms and Parameters

We set the length of the period in our model as one month and, therefore, choose a discount factor of $\beta = .96^{1/12}$. We assume that preferences are given by

$$u(c, m, l) = \frac{\eta}{\eta - 1} \log\left(\chi c^{\frac{\eta-1}{\eta}} + (1 - \chi) m^{\frac{\eta-1}{\eta}}\right) - \zeta l.$$

We follow Chari, Kehoe and McGrattan (2002) and set the weight on consumption, χ, equal to 0.94, and the elasticity of substitution between consumption and real balances, η, equal to 0.39. The parameter governing the disutility from work, ζ, simply sets the units in which we measure leisure and we choose it so that consumers supply one-third of their time to the labor market.

For the final good production function, we set θ, the elasticity of substitution across intermediate good inputs, to be 3. This number is in the middle of estimates of this elasticity in the literature. (See, for example, Nevo (1997) and Chevalier, Kashyap, and Rossi (2003).) We set the elasticity of capital, α, in the intermediate good firm production function equal to 1/3, and the elasticity of materials, ν, equal to 0.70. Given the 50% markup implied by our choice of θ, this implies a share of materials of slightly below 50%, consistent with U.S. evidence. Finally, we assume a capital depreciation rate of 1% per month, and set the size of the capital adjustment costs, ξ, equal to 29.65, so that the model reproduces the relative standard deviation of investment to consumption of 4 in the U.S. data.

We want to isolate the real effects of exogenous monetary shocks as a simple way of measuring the degree of nominal rigidity in the model. A popular way to do so is the approach of Christiano, Eichenbaum, and Evans (2005) and Gertler and Leahy (2008), who study the response of the economy to shocks in the money growth rate. We adopt the interpretation of Christiano, Eichenbaum, and Evans (2005), who extract the process for the exogenous component of money growth that is consistent with the monetary authority following an interest rate rule.[5] In that spirit, we set the coefficients in the money growth rule by first projecting the growth rate of (monthly) M1 on current and 24 lagged measures of monetary shocks.[6] We then fit an AR(1) process for the fitted values in this regression and obtain an autoregressive coefficient equal to .61 and a standard deviation of residuals of $\sigma_m = .0018$.

The parameters governing the frequency of price changes, α_L and α_T, are chosen so that the model can closely reproduce the salient features of the micro price data we have described. Specifically, we choose these two parameters jointly so that the model can simultaneously reproduce the frequency of price changes of 22% per month, as well as the frequency of regular price changes of 6.9% per month. (Here we define these price changes by applying the same statistical algorithm to the data generated from the model that we used on the BLS data. Note that regular prices produced by our algorithm mostly, but not always, coincide with the list price in the model.)

Notice from Table 2, Panel B, that $\alpha_L = 7.47\%$ and $\alpha_T = 7.90\%$. To get a feel of what these numbers imply note that when a firm receives an opportunity to temporarily change its price it typically undertakes two price changes, one to the temporary price in that period and then one back to the list price in the subsequent period. When a firm, however, receives an opportunity to change its list price it undertakes only one price change: it changes the list price and leaves it there. Thus, even if $\alpha_L = \alpha_T$ the model would imply that two-thirds

[5] Specifically, Christiano, Eichenbaum, and Evans (2005) specify an interest rate rule in their empirical work as $R_t = f(\Omega_t) + \varepsilon_t$, where R_t is the short-term nominal rate, Ω_t is an information set, and ε_t is the monetary shock. They interpret the monetary authority as adjusting the growth rate of money so as to implement this rule. They then identify the process for money growth in their vector autoregression consistent with this interest rate rule. That process is well approximated by an AR(1) similar to the one we use.

[6] The results we report here use a new measure of shocks due to Romer and Romer (2004), which is available for 1969–96. We have also used the measure of Christiano, Eichenbaum, and Evans (2005) and get very similar results.

of the price changes are temporary and one-third are regular so that the frequency of regular price changes is one third that of all price changes.

Notice, in Table 2, that the model, in addition to reproducing these 2 statistics exactly, can also account for the other measures of low- and high-frequency price stickiness in the data. The fraction of price changes that are temporary is equal to 75% (72% in the data), and firms charge a temporary price 9% of the time (10% in the data). The model also accounts well for our alternative measure of low-frequency price stickiness: 74% of prices are at their annual mode (75% in the data).

C. A Comparison with the Standard Calvo Model

We next compare the patterns of low- and high-frequency price stickiness in the standard model that has one type of price change with the same patterns in our model, referred to as the *benchmark* (Calvo) model. We show that, unlike our model, the standard (Calvo) model cannot simultaneously reproduce the micro data's high-frequency price flexibility and low-frequency price stickiness.

To demonstrate that, we consider a sequence of parameterizations of a standard Calvo model in which firms change prices with probability α. Recall that the standard model is a special case of our model with $\alpha_L = \alpha$ and $\alpha_T = 0$. We vary the frequency of micro price changes, α, in the standard model, convert it into months, and consider it a measure of the degree of high-frequency price stickiness. We keep all other parameters equal to those in our model with temporary price changes. Then for each model, we simulate a long price series and apply our algorithm to construct the regular price series. We compute the frequency of these regular price changes, convert it into months, and consider it a measure of the degree of low-frequency price stickiness.

The results are displayed in Figure 1. The curve in panel A shows that if micro prices are highly sticky in the standard model, then regular prices are too; the degrees of high- and low-frequency stickiness match. This is not the pattern we have seen in the data. That pattern—and the pattern produced by our benchmark model—is represented in panel A by a large dot. In the BLS data and in our benchmark model, prices have a low degree of high-frequency stickiness, about 4.5 months, but they also have a high degree of low-frequency (regular price) stickiness, about 14.5 months.

We also do an analogous experiment with the standard model for our alternative measure of low-frequency price stickiness, the fraction of prices at the annual mode. The results of that experiment, displayed in panel B of Figure 1, are consistent with the results of the regular price experiment. This consistency strongly suggests that our conclusions are not dependent on the exact way in which we measure low-frequency price stickiness or the details of our algorithm that defines regular prices.

3. The Degree of Aggregate Price Stickiness

We have shown that our benchmark model with temporary price changes can reproduce the main features of the BLS micro price data—and much better than a standard model can. We now turn to analyzing what our model has to say about the real effects of monetary policy, in terms of aggregate price stickiness, relative to what the standard model says. We find that the benchmark model predicts that aggregate prices are quite sticky despite the high frequency of micro price changes.

A. A Measure of Aggregate Price Stickiness

For this analysis, we must first define a measure of the degree of aggregate price stickiness in our model. We want a measure that captures how slowly the aggregate price level, P_t, reacts to a change in the money supply, M_t. We define *aggregate price stickiness* as the average difference in the first two years after the shock between the impulse responses of money and prices to a monetary shock scaled by the average impulse response of money. Note that up to a scalar of normalization, our measure is the difference between the cumulative impulse response (CIR) of money and prices.

To interpret this measure note that when it is large, then $\log M_t - \log P_t$ is large along the impulse response, which means that when the money supply increases, prices lag behind—so that prices are sticky. If prices fully react to changes in the money supply, so that the impulse response of prices is equal to that of money, then our measure of aggregate price stickiness is equal to zero. In contrast, if prices do not react at all to changes in the money supply, then our measure of aggregate price stickiness is equal to 1.

B. The Model's Implications

According to this measure and our benchmark model, aggregate prices are quite sticky despite how flexible prices are at the micro level.

To see that, we first document how the key variables respond to a particular monetary shock. We shock the money growth rate in period 1 by 19.5 basis points so that the level of the money supply increases 50 ($= 19.5/(1-\delta)$) basis points in the long run. This is approximately the size of a one standard deviation shock, which is 18 basis points.

In Figure 2, we display what the model predicts. The aggregate price level (in panel A) responds slowly to the shock. The gross domestic product (GDP), defined as final goods production net of spending on materials, reaches a peak of about 52 basis points in the first month after the shock and then gradually declines.

We quantify these responses with summary statistics. In Table 3 we display the result of calculating our measure of aggregate price stickiness given this particular shock. Recall that our measure is the average difference, over the first 24 months after the shock, between the impulse responses of the money supply and the price level divided by the average money supply impulse response over that period. The model's degree of aggregate price stickiness is about 58.6%.

The table also summarizes how the aggregate stickiness manifests itself in GDP. The money shock leads to an average GDP response of about 34.1 basis points in the first 24 months after the shock.

We have reported on one measure of the real effects of money, namely, the impulse response to a monetary shock. Another common measure of these effects is the volatility and persistence of output induced by such shocks at business cycle frequencies. In Table 3 we report that the standard deviation of HP-filtered output is equal to 0.81% while its serial correlation is 0.82.

C. The Benchmark Calvo Model vs. the Standard Calvo Model

For some perspective on the benchmark Calvo model's aggregate price implications, we now compare them with the standard Calvo model discussed above (which has $\alpha_L = \alpha$ and $\alpha_T = 0$ and the rest of the parameters as in our benchmark model). We find that for the standard model to reproduce the degree of aggregate price stickiness in our model, the frequency of price changes needs to be about 12 months. We find this comparison useful because it allows us to translate our measure of aggregate price stickiness into units that are familiar to those working in the New Keynesian literature.

In Figure 3, we report how, in this standard model, the degree of aggregate price stickiness varies with the degree of micro price stickiness ($1/\alpha$). Clearly, the model implies that frequent micro price changes correspond to low aggregate price stickiness, and infrequent micro changes, to high aggregate stickiness.

Recall that micro prices change every 4.5 months in the data. When a standard model reproduces this high frequency of micro price changes, as it does at point A in Figure 3, it predicts a low degree of aggregate price stickiness (24%). This is quite a contrast with our benchmark model, which predicts (at point B) a much higher degree (about 58.6%), as we have seen.

Translating our results into more commonly used units might be helpful here. Let us ask: What frequency of micro price changes does a standard model need in order to reproduce the degree of aggregate price stickiness in our model? In Figure 3, point C shows the answer: the standard model needs micro prices to change about once every 12 months, a very low frequency compared with that in the data.

The impulse responses of the standard model that match the degree of aggregate price stickiness in our benchmark model are displayed in Figure 4. Recall that, by construction the area between the impulse responses of money and prices is equal in the two models. Interestingly, we see that once we match this area, the shapes of the impulse responses of output and prices in the two models are nearly identical as well.

In sum, the impulse response to a money in our model with frequent price changes is well-approximated by a standard model in which prices change about once a year. In this sense, prices are sticky after all.

D. Intuition from a simplified version

Thus, even though our benchmark model is consistent with the frequent micro price changes in the data, it still predicts a quite sticky aggregate price level. How can this be?

How can temporary price changes, although very frequent, not allow the aggregate price level to react to monetary policy shocks? We argue that the answer is in the distinctive features of temporary price changes seen in the U.S. data.

We develop intuition for these answers by considering a stripped-down version of our benchmark model without capital, materials, and interest-elastic money demand, so that $M(s^t) = P(s^t)y(s^t)$. Moreover, we assume that utility is logarithmic in consumption and linear in leisure. These assumptions imply that nominal marginal cost (the wage rate) is proportional to the money supply. Finally, we assume that the log of the money supply, $m(s^t)$ is a random walk so that $\rho_\mu = 0$ in (1) and

$$m(s^{t+1}) = m(s^t) + \varepsilon_\mu(s^{t+1})$$

Consider first a Calvo version of the model in which a firm is allowed to reset its price with probability α. Under these assumptions, a firm that is given an opportunity to reset its list price after a one-time money shock chooses to respond one-for-one to the shock. Dropping the s^t notation it is easy to show that output is given by

$$(12) \quad y_t = (1-\alpha)y_{t-1} + (1-\alpha)(m_t - m_{t-1})$$

so starting from a steady state with $y_{-1} = m_{-1} = 0$, the cumulative impulse response to a money shock of size $m_0 = 1$ is

$$(13) \quad (1-\alpha)\left[1 + (1-\alpha) + (1-\alpha)^2 + \ldots\right] = \frac{1-\alpha}{\alpha}$$

Now, consider a *temporary price* version of this model by supposing that, in addition to being able to change its list price with some given probability, say, α_L, it can also temporarily deviate from its list price with probability α_T. Here the temporary price will also respond one-for-one with the money shock. It is easy to show that under the above assumptions output is given by

$$(14) \quad y_t = (1-\alpha_L)y_{t-1} + (1-\alpha_L-\alpha_T)(m_t - m_{t-1})$$

so the cumulative impulse response to the same money shock is

$$(15)\quad (1-\alpha_L-\alpha_T)\left[1+(1-\alpha_L)+(1-\alpha_L)^2+\ldots\right]=\frac{1-\alpha_L-\alpha_T}{\alpha_L}$$

Let us compare the impulse responses from these models. The impact effect in the Calvo model is $1-\alpha$ while it is $1-\alpha_L-\alpha_T$ in the temporary price version. After the impact period, output decays at rate $1-\alpha$ and $1-\alpha_L$ in the two models. Since the cumulative impulse response is primarily determined by the rate of decay, these response will be similar as long as α is close to α_L and α_T is not too large.

We show this result more precisely by asking a similar question to the one we posed in our quantitative model. In that model we asked: How often must prices change in the standard Calvo version to give the same degree of aggregate price stickiness as the benchmark Calvo model with some given α_L and α_T? We noted that our measure of price stickiness is proportional to difference between the cumulative impulse response of money and prices. Since $m_t-p_t=y_t$, this measure is the same as the cumulative impulse response of output. Focusing on the infinite (rather than the two year) cumulative response for simplicity, we equate (13) and (15) to get

$$(16)\quad \alpha=\frac{\alpha_L}{1-\alpha_T}.$$

Thus, if $\alpha_L=.075$ and $\alpha_T=.079$, as we found in our quantitative exercise, then $\alpha=.081$. Thus, a standard Calvo model needs prices to change once every 12.3 (1/.081) months to give the same aggregate price stickiness as the temporary price version. This is true, even though in the temporary price version prices change once every 4.5 months.

The key to our result is that the rate at which output decays in the temporary price version is solely a function of the frequency of list price changes α_L (and not of all price changes). To understand why this is the case consider the impulse response of the price level. Note that in any period after the shock, there are three types of firms: those that have already reset their list prices since the money shock occurred, those that have not reset their list price but currently have a temporary change and those that have not reset their list price but do not currently have temporary change (and hence are still charging the original list price).

To calculate the cumulative change in the aggregate price level, we simply add up the firms in the different categories and use the fact that any firm that has either a list price change or a temporary price change reacts one-for-one to the money shock. Therefore, the response of prices in period t is

$$(17)\quad p_t = \lambda_{L,t} + (1 - \lambda_{L,T})\,\alpha_T$$

where $\lambda_{L,t} = \alpha_L \sum_{i=0}^{t-1} (1 - \alpha_L)^i = 1 - (1 - \alpha_L)^t$ is the cumulative sum of the firms that have reset their list price by period t. To understand the expression for $\lambda_{L,t}$, note that one period after the shock α_L firms have reset their list prices, two periods after the shock $\alpha_L + (1 - \alpha_L)\,\alpha_L$ have reset them, and so on.

Notice that the rate at which the price level increases with t is solely a function of the frequency of list price changes, α_L. The reason for this result is that list price changes are permanent: once a firm changes its list price, it permanently responds to the money supply shock, and that holds regardless of when it made this change. Temporary price changes, in contrast, last only one period: after one period these prices simply return to their previous list price. Hence, firms that have had temporary price deviations in the past have returned their price to their pre-existing level and do not affect the cumulative price level.

In sum, temporary price changes are special. Because they return the nominal price back to its pre-existing level, they allow firms to only temporarily respond to a change in monetary policy. Hence, following a monetary shock, they affect neither the rate at which price level increases nor the rate at which output decays.

4. A Menu Cost Model with Temporary Price Changes

So far we have focused on the widely-used Calvo sticky price framework. Researchers typically interpret the Calvo model as a simple reduced form for a menu cost model. The natural question then arises: Do our results extend to a menu cost framework? That is, does a simple extension of the menu cost model that is consistent with the patterns of price changes seen in the micro data predict that aggregate prices are sticky? Here we extend the menu cost model and show that our results are robust.

Importantly, we also take up the challenge for menu costs models posed by EJR.

These authors argue that standard menu cost models are inconsistent with four features of the data they study. First, in their data prices are more volatile than marginal costs. Second, prices tend to return to a slow-moving trend. Third, there is substantial high frequency price flexibility together with substantial low frequency price stickiness. Finally, prices rarely change without changes in costs. We show that our menu cost model can reproduce all of these features of the data.

A. Additional Facts and Overview of the Model

Earlier we presented five facts from the BLS data about the regular and temporary movements in prices. Here we discuss two additional sets of facts that we report in Table 4. The first set has to do with the size and dispersion of price changes in the BLS data, that are the focus of menu cost literature that builds on Golosov and Lucas (2007). The second set has to do with the relation between prices and costs emphasized by EJR for their proprietary data set.

From the table we see the first set of additional facts: price changes are both large and dispersed. In particular, the mean size of price changes are 11% for all prices and for regular prices. Price changes are dispersed in that the interquartile range (IQR) of all price changes is 9% and the IQR of regular price changes is 8%.

We also see the second set of additional facts: prices are more volatile than costs and prices and costs tend to move together. From the table we see that the standard deviation of prices relative to that of costs is 1.33, so that prices are a third more volatile than costs. We also see that most price changes are associated with cost changes: in only 7% of periods in which there are price changes there are no cost changes.

Clearly, a model that can generate all of these facts needs to be much richer than our simple benchmark model. Moreover, part of the challenge of EJR is to generate these facts in a model in which firms choose the timing of price changes optimally (rather than their timing being exogenously given as in the benchmark model). To that end we extend the standard menu cost model of Golosov and Lucas (2007) in several ways.

To account for the pattern of temporary and regular price changes in the data, we make three additional assumptions. First, we allow for both transitory and permanent idiosyncratic productivity shocks. These shocks help the model deliver the large temporary and regular price changes in the data. Second, we introduce time-varying demand elasticities by having good-specific demand shocks. Time-varying demand elasticities are a popular explanation for temporary price changes in the IO literature (see, for example, Sobel 1984 and Pesendorfer 2002) and allow our model to match the fact that prices are more volatile than costs.

Third, we now assume that in addition to paying a fixed cost κ to change the list price, the firm also has the option to pay a fixed cost ϕ to charge a price other than the list price for one period. (In our robustness section below we discuss three variants of this price setting technology and show that all three lead to similar results.)

As in the Calvo model, we think of the list price as the price set by the upper-level manager and the posted price as the price actually charged to the consumer. In contrast to the Calvo model, however, here the decision to deviate from the list price is no longer exogenous and random but rather endogenous. Thus, here the timing of temporary price deviations can respond to all shocks, including monetary shocks.

Overall, we think of our model as a parsimonious extension of an otherwise standard menu cost model that allows it to generate both temporary and regular price changes of the type documented by EJR.

B. Setup

The consumer's problem is identical to that in our Calvo model. What differs is the technologies for producing intermediate and final goods.

Intermediate good i is produced according to

$$(18)\quad y_i(s^t) = a_i(s^t)z_i\left(s^t\right)\left[k_i\left(s^t\right)^{\alpha} l_i(s^t)^{1-\alpha}\right]^{\nu} n_i\left(s^t\right)^{1-\nu},$$

where $a_i\left(s^t\right)$ is a permanent productivity component and $z_i\left(s^t\right)$ is a transitory productivity component. The permanent component follows a random walk process, while the transitory component follows an autoregressive process. We describe both below. To ensure stationary, we assume a fraction ρ_e of firms exit every period and are replaced by new entrants that draw a value of $a_i\left(s^t\right) = z_i\left(s^t\right) = 1$.

As earlier, we assume that there is a continuum of final good firms, that combine varieties of the intermediate goods into a final good. We modify the technology for producing final goods to

$$(19) \quad y\left(s^t\right) = y^A\left(s^t\right)^{1-\omega} y^B\left(s^t\right)^{\omega}$$

with

$$(20) \quad y^A\left(s^t\right) = \left(\int_0^1 y_i^A\left(s^t\right)^{\frac{\theta-1}{\theta}} di\right)^{\frac{\theta}{\theta-1}} \text{ and } y^B\left(s^t\right) = \left(\int_0^1 v_i\left(s^t\right)^{\frac{1}{\gamma}} y_i^B\left(s^t\right)^{\frac{\gamma-1}{\gamma}} di\right)^{\frac{\gamma}{\gamma-1}}.$$

where $v_i\left(s^t\right)$ is a good-specific shock and $\gamma > \theta$. As we show below, this two-tier specification of technology, in conjuction with the good-specific shocks, implies that demand for intermediate goods is characterized by time-varying elasticity.

The resource constraint for final goods is, as earlier, (7). The final good firm chooses the intermediates inputs $\{y_i(s^t)\}$ to maximize

$$P(s^t)y(s^t) - \int P_i(s^t)\left[y_i^A(s^t) + y_i^B\left(s^t\right)\right] di$$

subject to (19)-(20). The solution to this problem gives the demand for intermediate good i, which we can write as $y_i(s^t)/y(s^t) =$

$$(21) \quad (1-\omega)\left(\frac{P_i(s^t)}{P^A(s^t)}\right)^{-\theta}\left(\frac{P^A\left(s^t\right)}{P\left(s^t\right)}\right)^{-1} + v_i\left(s^t\right)\left(\frac{P_i(s^t)}{P^B(s^t)}\right)^{-\gamma}\left(\frac{P^B\left(s^t\right)}{P\left(s^t\right)}\right)^{-1}.$$

The zero profits condition implies that

$$P(s^t) = (1-\omega)^{-(1-\omega)}\omega^{-\omega}(P^A(s^t))^{1-\omega}(P^B(s^t))^{\omega},$$

where $P^A(s^t) = \left(\int_0^1 P_i\left(s^t\right)^{1-\theta} di\right)^{\frac{1}{1-\theta}}$ and $P^B\left(s^t\right) = \left(\int_0^1 v_i\left(s^t\right) P_i\left(s^t\right)^{1-\gamma} di\right)^{\frac{1}{1-\gamma}}$.

A useful feature of the resulting demand function is that it has time-varying elasticity. Clearly, as the demand shock $v_i\left(s^t\right)$ increases, so does the total demand elasticity for good i, since $\gamma > \theta$. Such a shock would therefore lead the intermediate goods firm to optimally lower its markup, and therefore change its price even in the absence of cost changes. We assume that $v_i(s^t)$ follows a first-order autoregressive process that we describe below.

The nominal unit cost of producing good i is

$$(22)\quad V_i\left(s^t\right) = \frac{V\left(s^t\right)}{a_i\left(s^t\right) z_i\left(s^t\right)},$$

where $V(s^t)$ is given by (5). The firm's period profits, gross of fixed costs, are therefore $\Pi_i(s^t) = (P_i(s^t) - V_i(s^t))\, y_i(s^t)$. The nominal present discounted value of profits of the firm is given by

$$(23)\quad \sum_t \sum_{s^t} Q(s^t)(1-\rho_e)^t \big[\Pi_i\left(s^t\right) - W(s^t)\big(\kappa\delta_{L,i}(s^t) + \phi\delta_{T,i}(s^t)\big)\big],$$

where $\delta_{L,i}(s^t)$ is an indicator variable that equals one when the firm changes its list price $(P_{L,i}(s^t) \neq P_{L,i}(s^{t-1}))$ and zero otherwise, and $\delta_{T,i}(s^t)$ is an indicator variable that equals one when the firm temporarily deviates from the list price $(P_i(s^t) \neq P_{L,i}(s^t))$ and zero otherwise. In expression (23), the term $W(s^t)\kappa\delta_{L,i}(s^t)$ is the labor cost of changing list prices, which we think of as the *menu cost*, and $W(s^t)\phi\delta_{T,i}(s^t)$ is the cost of deviating from the list price.

C. Quantification

We assign the model the same parameters describing preferences and technology as in the Calvo model. We set the higher demand elasticity, γ, equal to 6, at the upper range of estimates in existing work and the lower elasticity, θ, equal to 2.15. With these elasticities the model implies that firms sell about twice as much output in periods with temporary markdowns than they do otherwise, a number consistent with evidence from grocery stores.[7]

The rest of the parameters are chosen so that the model can closely reproduce the salient features of the micro price data from the BLS, as well as the statistics reported by Eichenbaum, Jaimovich and Rebelo (2011). These parameters include κ, the (menu) cost the firm incurs when changing its list price; ϕ, the cost of deviating from the list price; the specifications of the productivity and demand shocks, as well as the parameters describing the technology with which final goods firms aggregate intermediate inputs.

[7] For example, in the Dominick's data the ratio of quantity sold in periods with markdowns to that in other periods is 2.15. In our model this ratio is 2.11.

Consider, first, the specification of the permanent productivity shocks. Midrigan (2011) shows that when productivity shocks are normally distributed, a model like ours generates counterfactually low dispersion in the size of price changes. Midrigan argues that a fat-tailed distribution is necessary in order for the model to account for the distribution of the size of price changes in the data. We find that a parsimonious and flexible approach to increasing the distribution's degree of kurtosis is to assume, as Gertler and Leahy (2008) do, that productivity shocks arrive with a Poisson probability and are, conditional on arrival, uniformly distributed. We follow this approach and assume that the permanent productivity component, $a_i\left(s^t\right)$, evolves according to

$$\ln a_i\left(s^t\right) = \ln a_i\left(s^{t-1}\right) + \varepsilon_{a,i}\left(s^t\right),$$

where $\varepsilon_{a,i}\left(s^t\right) \sim U\left[-\bar{a}, \bar{a}\right]$ with probability λ_a and 0 with probability $1-\lambda_a$. The transitory component, $z_i\left(s^t\right)$, evolves according to

$$z_i\left(s^t\right) = \rho_z z_i\left(s^{t-1}\right) + \varepsilon_{z,i}\left(s^t\right),$$

where $\varepsilon_{z,i}\left(s^t\right) \sim U\left[z^L, z^H\right]$ with probability λ_z and 0 with probability $1-\lambda_z$.

The optimal markup of a firm in this economy, absent adjustment costs, is a function of $\mathbf{v}_i\left(s^t\right) = v_i\left(s^t\right)\left(a_i\left(s^t\right) z_i\left(s^t\right)\right)^{\gamma-\theta}$. To reduce computational complexity, we specify the demand shock $v_i\left(s^t\right)$ so that the composite term $\mathbf{v}_i\left(s^t\right)$ is independent of the productivity shocks $a_i\left(s^t\right)$ and $z_i\left(s^t\right)$. In particular, we assume

$$\mathbf{v}_i\left(s^t\right) = \rho_v \mathbf{v}_i\left(s^{t-1}\right) + \varepsilon_{v,i}\left(s^t\right),$$

where $\varepsilon_{v,i}\left(s^t\right) \sim U\left[0,1\right]$ with probability λ_v and 0 with probability $1-\lambda_v$. The bounds for the shock are simply normalized to lie on the unit interval since they are not separately identified from ω, the relative weight on the type B aggregator in the final good production function.

Paying special attention to the distribution of idiosyncratic shocks is necessary because this distribution plays an important role in determining the real effects of changes in the money supply. Golosov and Lucas (2007) show, for example, that the effects of monetary shocks are approximately neutral when idiosyncratic shocks are normally distributed. But as Midrigan (2011) shows, with a fat-tailed distribution of idiosyncratic shocks, shocks to the money supply have much larger real effects because changes in the identity of adjusting firms are muted as the kurtosis of the distribution of productivity shocks increases.

We choose all these parameters to minimize the squared deviation between they salient moments in the data and the moments in the data and the model listed in panel A of Table 4. The moments include the facts about temporary and regular price changes, as well as other measures of the degree of low- and high-frequency price variation in the BLS data we have discussed, the size and dispersion of price changes, as well as the statistics on the relative variability of prices and costs from EJR. We also include information from the Dominick's data on the relative amount of quantities sold in periods with temporary price changes to pin down the lower demand elasticity θ. Panel B of Table 2 lists the parameter values that allow the model to best match the moments in the data.

The Micro Moments in the Data and the Model

Our parsimonious extension of a standard menu cost model accounts well for the micro moments we have documented. Recall that in the data we computed statistics about regular prices using our algorithm. We use the same algorithm now to construct statistics about regular prices in the model. (Recall that the regular prices produced by our algorithm mostly, but not always, coincide with the list price in the theory.)

For the details, see panel A of Table 4. The frequency of posted price changes is high: 22% in the data and 23% in the model; the frequency of regular price changes is much lower: 6.9% in both the data and the model.[8] Most price changes are temporary: 72% in the data and 78% in the model. Temporary price changes are transitory: the probability that a temporary price spell ends is equal to 53% in the data and 66% in the model. Periods with temporary prices account for 10% of all periods in the data and 11% in the model, and about 60% of these periods are ones with temporary price declines in both the data and the model.

The model also accounts well for our alternative measure of low-frequency price stickiness: 75% of prices are at their annual mode in that data while 73% are in the model.

[8]We compute the frequency of price changes only for those products that are not replaced, both in the model and in the data.

Following Golosov and Lucas (2007) and Midrigan (2011), we also examine the size and dispersion of price changes. The mean size of all price changes and regular price changes is high in both the data and the model (11% in the data and 12% in the model). So is the dispersion of these changes as measured by the interquartile range (IQR): 9% for all price changes and 8% for regular price changes in the data versus 8% and 8% in the model, respectively.

Importantly, the model successfully accounts for the key statistics from EJR. In both the data and the model, prices are about a third more volatile than costs: their relative standard deviation is 1.33 in the data and 1.32 in the model. Also, in both the data and the model, most price changes are associated with cost changes: in both of them in only 7% of periods in which there are price changes are there no cost changes.

D. The Degree of Aggregate Price Stickiness

Here we discuss the degree of aggregate price stickiness in our menu cost model with temporary price changes. We then compare this model's implications to those of a standard menu cost model (without temporary price changes).

In Table 5 we see that the degree of aggregate price stickiness for our menu cost model is 52.5% while the average output response to a 50 basis point monetary shock is 29.6 basis points in the first two years after a shock.

The standard model used in our comparison retains the permanent productivity shocks of our benchmark model, but follows Golosov and Lucas (2007) in abstracting from other shocks. We adjust the parameters governing the permanent productivity process and the size of the menu cost so that the standard model matches the average size (11%), and IQR of price changes (9%) in the data as well as the degree of aggregate price stickiness in our menu cost model with temporary changes. When we do so we see that micro prices must change once every 10.1 months for the standard menu cost model to reproduce the degree of aggregate price stickiness in our menu cost model with temporary price changes.

Notice that we find that the degree of aggregate price stickiness in the menu cost model (10 months) is lower than it is in the Calvo model (12 months). To understand why recall that the menu cost model has two additional mechanisms that tend to lower the degree of aggregate price stickiness. First, the timing of temporary price changes can potentially respond to money shocks. Second, a disproportionate amount of goods is sold during periods with temporary price changes. These two mechanisms, though present in the menu cost model, are, however, weak quantitatively and do not overturn our earlier results.

In sum, the menu cost model shows our earlier result is robust: Even though prices change frequently at the micro-level, the impulse response of the model is well approximated by a standard menu cost model in which prices change infrequently, roughly once a year. In this sense, our result that prices are sticky after all is robust to two very different ways of modeling sticky prices.

E. Robustness Checks

Here we discuss the large number of robustness checks we conducted on our menu cost model. Some of these checks are on the details of the price-setting technologies while others are on the nature of idiosyncratic and aggregate shocks. We find that our result is robust to all of these checks. We report on all these checks in our online appendix.

We begin by exploring alternative price-setting technologies. One reason is that while we think of the work Zbaracki et al. (2004) as suggestive about the existence of costs of deviating from the regular price, this work is clearly not precise enough to pin down the exact details of what the lower level manager can do after contacting the upper level manager. We consider three alternative specifications about what paying the fixed cost entitles the firm to do. In the *sticky temporary price* version, this extra cost gives the manager the right to continuously charge a given temporary price as long as that manager sees fit. In the *flexible temporary price* version, this extra cost gives the manager the right to vary the temporary price it charges freely for a given period of time, say three months. Finally, in the *free switching to a temporary price* version, this extra cost gives the manager the right to choose one temporary price and to freely switch between that one price and the regular price for a fixed amount of time, say three months. We find that our results are robust to these alternative pricing technologies.

We also consider a large number of variations of our benchmark model. We explore the role of capital, interest-elastic money demand, and real rigidities; an alternative data set (Dominick's); random menu costs that allow the model to generate small price changes; alternative specifications of the productivity shocks (Gaussian and allowing for correlation between them); added shocks to the elasticity of demand; and alternative specifications of monetary policy. We find that the quantitative implications of our main result—that the aggregate price level is as sticky as it is in a standard model in which micro prices change only very infrequently—is robust to all these features.

5. Concluding Remarks

Micro price data shows a great deal of high-frequency price flexibility but low-frequency price stickiness. We have shown that two classes of sticky price models that are consistent with these features of the data imply a large degree of aggregate price stickiness. In this sense, even though prices change frequently at the micro level, aggregate prices are sticky.

References

Bils, Mark, and Peter J. Klenow. 2004. Some evidence on the importance of sticky prices. *Journal of Political Economy* 112: 947–85.

Burdett, Kenneth and Kenneth L. Judd. 1983. Equilibrium Price Dispersion. *Econometrica*, 51 (4): 955-969.

Chevalier, Judith A., and Anil K. Kashyap. 2011. Best prices. NBER Working Paper 16680.

Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi. 2003. Why don't prices rise during periods of peak demand? Evidence from scanner data. *American Economic Review* 93: 15–37.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy* 113: 1–45.

Eichenbaum, Martin S., Nir Jaimovich, and Sergio Rebelo. 2011. Reference prices, costs, and nominal rigidities. *American Economic Review* 101: 234–62.

Gertler, Mark, and John Leahy. 2008. A Phillips curve with an *Ss* foundation. *Journal of Political Economy* 116: 533–72.

Golosov, Mikhail, and Robert E. Lucas, Jr. 2007. Menu costs and Phillips curves. *Journal of Political Economy* 115: 171–99.

Guimarães, Bernardo, and Kevin D. Sheedy. 2011. Sales and monetary policy. *American Economic Review* 101: 844-76.

Head, Allen, Lucy Qian Liu, Guido Menzio, and Randall Wright. 2011. Sticky Prices: A New Monetarist Approach. NBER Working Paper 17520.

Hosken, Daniel, and Daniel Reiffen. 2004. Patterns of retail price variation. *Rand Journal of Economics* 35: 128–46.

Kehoe, Patrick, and Virgiliu Midrigan. 2010. Prices are sticky after all. NBER WP 16364.

Klenow, Peter J., and Benjamin A. Malin. 2010. Microeconomic evidence on price-setting. In *Handbook of Monetary Economics*, vol. 3A: 231-84, B. Friedman and M. Woodford, eds.

Lazear, Edward P. 1986. Retail pricing and clearance sales. *American Economic Review* 76: 14–32.

Midrigan, Virgiliu. 2011. Menu costs, multi-product firms, and aggregate fluctuations. *Econometrica* 79(4): 1139-80.

Nakamura, Emi, and Jón Steinsson. 2008. Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics* 123: 1415–64.

Nakamura, Emi, and Jón Steinsson. 2010. More facts about prices. Supplement to “Five facts about prices: A reevaluation of menu cost models.” Manuscript. Columbia University.

Nevo, Aviv. 1997. Measuring market power using discrete choice models of demand: An application to the ready-to-eat cereal industry. Working Paper. Univ. of Mass.

Romer, Christina D., and David H. Romer. 2004. A new measure of monetary shocks: Derivation and implications. *American Economic Review* 94: 1055–84.

Rotemberg, Julio J. Forthcoming. Fair pricing. *Journal of the European Economic Association*.

Rotemberg, J. J. 1987. The New Keynesian Microfoundations. *Macroeconomics Annual 2*.

Sobel, Joel. 1984. The timing of sales. *Review of Economic Studies* 51: 353–68.

Varian, Hal R. 1980. A model of sales. *American Economic Review* 70: 651–59.

Warner, Elizabeth J., and Robert Barsky. 1995. The timing and magnitude of retail store markdowns: Evidence from weekends and holidays. *Quarterly Journal of Economics* 110: 321–52.

Zbaracki, Mark J., Mark Bergen, and Daniel Levy. 2007. The anatomy of a price cut. Working Paper. University of Western Ontario.

Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *Review of Economics and Statistics* 86: 514–33.

Table 1: Facts about price changes in BLS data

Frequency of all price changes	22.0%
Frequency of regular price changes	6.9%
Percentage of price changes that are temporary	72%
Fraction of periods with temporary prices	10%
Fraction of prices at annual mode	75%

Table 2: Parameterization. The Calvo Model

A. Moments

	BLS Data	Model
Frequency of all price changes	0.22	0.22
Frequency of regular price changes	0.069	0.069
Fraction of price changes that are temporary	0.72	0.75
Fraction of periods with temp. prices	0.10	0.09
Fraction of prices at annual mode	0.75	0.74

B. Parameter Values

Calibrated	
Probability of changing list price, α_L, %	7.47
Probability of deviating from list price, α_T, %	7.90
Assigned	
Period length	1 month
Annual discount factor	0.96
AR(1) growth rate of M	0.61
S.D. of shocks to growth rate of M, %	0.18
Capital elasticity, α	0.33
Materials elasticity, ν	0.70
Weight on C in utility, χ	0.94
Money demand elasticity, η	0.39
Capital depreciation, δ	0.01
Capital adjustment cost, ξ	21.95

Table 3: Aggregate Implications. The Calvo Models

Statistic	Benchmark Model (with temporary changes)	Standard Model (without temporary changes)
Micro-price stickiness, months	4.5	12.2
	Impulse Response to a 50 b.p. monetary shock	
Aggregate price stickiness, %	58.6	58.6
Average output response, b.p.	34.1	34.1
Maximum output response, b.p.	52.2	54.0
	Business Cycle Statistics	
Std. dev output, %	0.81	0.84
Autocorrelation output	0.82	0.82

Notes:

Aggregate price stickiness is measured as the average difference between M and P responses, relative to the M response.
Responses are computed for the first 2 years after the shock.
Business cycle statistics reported for HP(14400) filtered data

Table 4: Parameterization. The Menu Cost Model

A. Moments

	BLS Data	Model
Frequency of all price changes	0.22	0.23
Frequency of regular price changes	0.069	0.069
Fraction of price changes that are temporary	0.72	0.78
Fraction of periods with temp. prices	0.10	0.11
Fraction of prices at annual mode	0.75	0.73
Probability that temporary price spell ends	0.53	0.66
Fraction of periods with price temp. down	0.06	0.06
Mean size of price changes	0.11	0.12
Mean size of regular price changes	0.11	0.11
IQR of all price changes	0.09	0.08
IQR of regular price changes	0.08	0.08
Std. dev. changes in prices vs. costs	1.33	1.32
Fraction of price changes w/o cost changes	0.07	0.07

B. Parameter Values

Calibrated	
Menu cost of regular price change, κ, % SS profits	0.38
Cost of temp. price deviation, ϕ, % SS profits	0.12
Arrival rate of permanent shock, λ_a	0.083
Upper bound of permanent productivity shock, *a_bar*	0.191
Arrival rate of transitory shock, λ_z	0.081
Bound on transitory productivity shock, $[z^L, z^H]$	[-0.17, 0.19]
Persistence of transitory productivity, ρ_z	0.40
Arrival rate of demand shock, λ_v	0.007
Persistence of demand shock, ρ_v	0.10
Weight on type A aggregator, $1-\omega$	0.973
Elasticity of type A aggregator	2.15
Assigned	
Period length	1 month
Annual discount factor	0.96
AR(1) growth rate of M	0.61
S.D. of shocks to growth rate of M, %	0.18
Capital elasticity, α	0.33
Materials elasticity, ν	0.70
Weight on C in utility, χ	0.94
Money demand elasticity, η	0.39
Capital depreciation, δ	0.01
Capital adjustment cost, ξ	21.95
Elasticity of type B aggregator	6

Table 5: Aggregate Implications. The Menu Cost Models

Statistic	Benchmark Model (with temporary changes)	Standard Model (without temporary changes)
Micro-price stickiness, months	4.5	10.1
	Impulse Response to a 50 b.p. monetary shock	
Aggregate price stickiness, %	52.5	52.5
Average output response, b.p.	29.6	29.6
Maximum output response, b.p.	40.7	44.7
	Business Cycle Statistics	
Std. dev output, %	0.67	0.72
Autocorrelation output	0.86	0.86

Notes:

Aggregate price stickiness is measured as the average difference between M and P responses, relative to the M response.
Responses are computed for the first 2 years after the shock.
Business cycle statistics reported for HP(14400) filtered data

Figure 1: Relationship between high- and low-frequency stickiness. Standard Calvo model

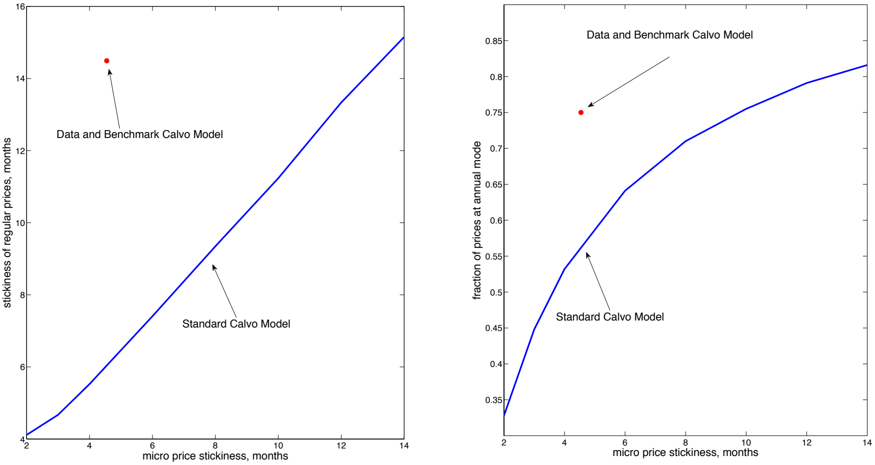

Figure 2: Impulse responses to 50 b.p. monetary shock.
Benchmark Calvo model

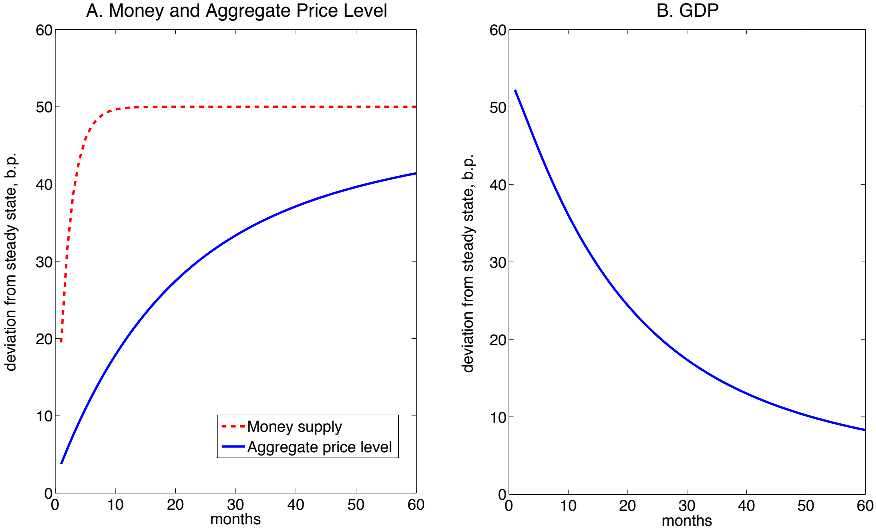

Figure 3: Aggregate price stickiness vs. Micro price stickiness. Standard Calvo model

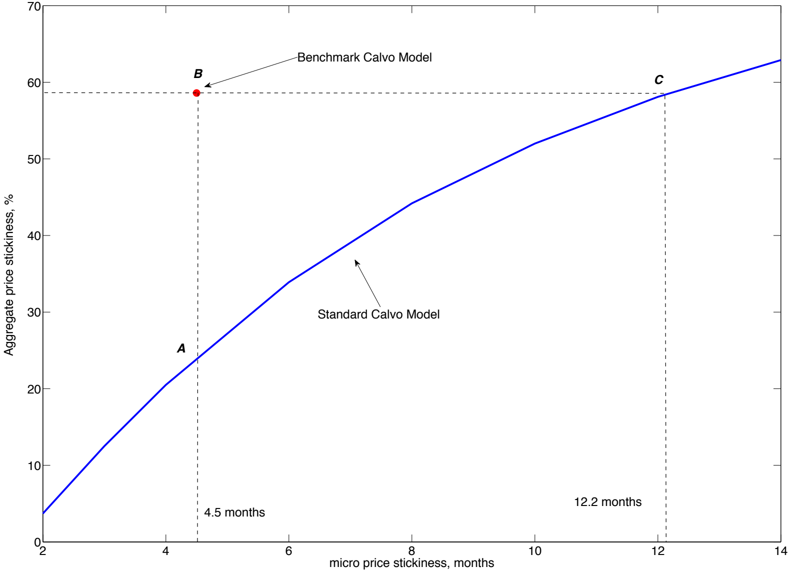

Figure 4: Impulse responses to 50 b.p. monetary shock. Benchmark Calvo and Standard Calvo model with 12.2-month stickiness