

NBER WORKING PAPER SERIES

PRICES ARE STICKY AFTER ALL

Patrick J. Kehoe
Virgiliu Midrigan

Working Paper 16364
<http://www.nber.org/papers/w16364>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2010

We thank Kathy Rolfe and Joan Gieseke for excellent editorial assistance. Kehoe thanks the National Science Foundation for financial support. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Patrick J. Kehoe and Virgiliu Midrigan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Prices are Sticky After All
Patrick J. Kehoe and Virgiliu Midrigan
NBER Working Paper No. 16364
September 2010
JEL No. E24,E3,E32

ABSTRACT

Recent studies say prices change 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 if most price changes are regular, as they are in the standard New Keynesian model, then this interpretation is correct. But, if most price changes are temporary, as they are in the data, then it is incorrect. Temporary changes have two striking features: after a change, the nominal price returns exactly to its pre-existing level, and temporary changes are clustered in time. Our model, which replicates these features, implies that temporary changes cannot offset monetary shocks well, whereas regular changes can. Since regular prices are much stickier than temporary ones, our model, in which prices change as frequently as they do in the micro data, predicts that the aggregate price level is as sticky as in a standard model in which micro level prices change once every 12 months. In this sense, prices are sticky after all.

Patrick J. Kehoe
Research Department
Federal Reserve Bank of Minneapolis
90 Hennepin Avenue
Minneapolis, MN 55480-0291
and NBER
pkehoe@res.mpls.frb.fed.us

Virgiliu Midrigan
Department of Economics
New York University
19 W. 4th St.
New York, NY 10012
and NBER
virgiliu.midrigan@nyu.edu

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, 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 shock effects, whereas if good prices change often, then the aggregate price level is not sticky and can. In this model, in other words, the frequency of micro price changes has a one-to-one relationship with the degree of aggregate price stickiness.

Until recently, micro level prices have been assumed to be quite sticky—changing relatively infrequently, only about once a year and, 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 4.5 months (Bils and Klenow 2004). According to these studies, that is, prices are quite flexible at the micro level. Interpreted through the lens of the standard New Keynesian model, this evidence implies that aggregate prices are quite flexible too, so monetary policy cannot have large real effects.

We dispute this interpretation. Although it is true that this interpretation follows logically from a standard New Keynesian model, that model is grossly inconsistent with the pattern of price changes in the micro data. We build a simple extension of the standard model that is consistent with the micro data. We show that in this extension of the model, aggregate prices are, in fact, quite sticky and monetary policy has large real effects.

The major shortcoming of standard New Keynesian models is their inability to simultaneously account for the pattern of high-frequency price flexibility and low-frequency price stickiness that we document using two sets of micro data. The pattern shows up in a simple example displayed in Figure 1, where we plot a fairly typical micro price series, decomposed into two parts: a trend and temporary deviations from trend. In this series, prices often temporarily move away from a slow-moving trend line, which tends to change rarely. We call the trend line the *regular price* and the deviations the *temporary price*.¹ Notice that most

price changes in this series are temporary and have two distinctive features: after a temporary change, the nominal price returns exactly to its pre-existing level, and temporary changes are clustered in time—for example, a change in one direction is usually followed quickly by another change in the opposite direction. These two distinctive features imply that for individual price series, even though there is a great deal of *high-frequency price flexibility* (in that actual prices change frequently), there is also a great deal of *low-frequency price stickiness* (in that the trend price changes infrequently).

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

We extend the standard model to allow firms to temporarily deviate from a sticky pre-existing price. We quantify our model and show that it reproduces the empirical micro pattern of regular and temporary price changes. We then study our model’s implication for the degree of aggregate price stickiness in response to monetary policy shocks. It is quite different from that of a New Keynesian model. Unlike that model, our model, which is consistent with the micro data pattern, still implies large real effects from monetary shocks.

The key prediction of our model is that temporary micro price changes cannot offset monetary shocks well, whereas regular changes can—but they do so only infrequently because regular prices are much stickier than temporary prices. To measure how sticky our model says aggregate prices are, we translate its results into those of a standard model. Our model, in which micro prices change as frequently as they do in the data, predicts that the aggregate price level is as sticky as it is in a standard model in which micro prices change once every 12 months. Aggregate prices, the model says, are sticky after all.

Our study is based on the salient features of the micro price data. To document these features, we study two sets of data: monthly price data from the U.S. Bureau of Labor Statistics (BLS) and weekly data from a chain of grocery stores, Dominick’s Finer Food retail chain. Our primary focus is on the statistics from the larger, more comprehensive BLS data

set, but we use the Dominick's data to demonstrate the robustness of that analysis. We find that the Dominick's data set is also informative, despite its limited coverage, because these data are collected weekly and thus can better measure high-frequency price movements than the monthly data of the BLS. Moreover, the Dominick's data include quantities sold as well as prices, so that we can examine the extent to which temporary changes account for a disproportionate amount of goods sold.

We document in both data sets the two types of micro price changes and quantify their distinctive features. Recall that one of two key features of temporary price changes is that after such a change, the nominal price typically returns exactly, to the penny, to its level before the change. In the Dominick's data, this event occurs 80% of the time; in the BLS data, 50% of the time. In contrast, in both data sets, regular price changes almost never return the nominal price to its pre-existing level. The other key feature of temporary price changes is that they are clustered over time. In both data sets, for example, the probability that a temporary price spell ends in any particular period is about 50%. This means that a temporary change in one direction is often quickly followed by a change in the opposite direction.

Consider now the intuition for our main result that the aggregate price level is sticky even though micro prices change frequently. The model produces price series for individual goods similar to that in Figure 1. In the model, therefore, 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. Moreover, since temporary price changes are highly clustered in time, they are less able to offset persistent changes in monetary policy. For example, a firm that changes its prices four times in January and not at all the rest of the year is less able to respond to persistent money supply changes than a firm that also changes its prices four times a year, but spreads those changes out over the year to, say, once a quarter. For these two reasons, even though micro prices change frequently, the aggregate price level is sticky. Our key insight 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 in the data there is substantial low-frequency price stickiness, the aggregate price level is sticky as well.

Our model of temporary and regular price changes is motivated by evidence on the pricing practices of actual firms. In particular, Zbaracki et al. (2004, 2007) provide evidence that pricing is done at two levels: upper-level managers (at headquarters) set list prices, while lower-level managers (at the store level) choose the actual transaction (posted) prices. These researchers find that the managerial costs of changing list prices are much greater than the physical costs of changing posted prices. Moreover, lower-level managers must effectively pay a time cost to be allowed to depart from the regular price set by the upper-level managers. This interaction between lower- and upper-level managers is illustrated, for example, in the following quotation from an interview with a sales manager (Zbaracki et al. 2004, p. 524):

I was a territory manager so I had no pricing authority. The only authority I had was to go to my boss and I would say, "OK, here is the problem I've got." He would say "Fill out a request and we will lower the price for that account." So that is how the pricing negotiations went. At that time I went up the chain to make any kind of adjustments I had to make My five guys have a certain level [of discount] they can go to without calling me. When they get to the certain point they have to get my approval.

We model this two-level decision-making process in a simple, reduced-form way. We assume that retailers set two prices: a list price and a posted price. The posted price is the price at which goods are sold. In our model, the list price matters because charging a posted price other than the list price entails a fixed cost. We think of this cost as standing in for the cost of the lower-level manager obtaining approval from the upper-level manager for temporarily charging a price that deviates from the list price. Changing list prices themselves entails another fixed cost, the managerial cost of upper-level decision making. We assume that the cost of changing posted prices is zero. (We also solved a version of the model with a fixed cost of changing posted prices and got results nearly identical to those presented here. For simplicity, we focus on the model without such costs.)

Our model is purposely chosen to be a parsimonious extension of the standard menu

cost model of, say, Golosov and Lucas (2007) and Midrigan (2007). Our extension has a different technology for changing prices, as we have described, and the addition of temporary shocks, which give a motive for temporary price changes. Specifically, in our model, firms are subject to two types of idiosyncratic disturbances: persistent productivity shocks and transitory shocks to either the cost or the elasticity of demand for the firm's product. The latter shocks are meant to capture in a simple way an idea popular in the industrial organization literature: that firms face demand for their products with time-varying elasticity. Our simple extension allows the model to produce patterns of both temporary and regular price changes that are similar to those in the data.

A sizable literature in industrial organization has suggested explanations for the temporary price discounts (or *sales*) which account for the majority of temporary changes in the data.² Unfortunately, all of these explanations are about real prices and, hence, cannot explain a striking feature of the data: that the nominal price, after a temporary price discount, often returns exactly to the nominal pre-existing price. Indeed, this feature is the subtle low-frequency price stickiness which is at the heart of our results. Namely, even though there is a large amount of high-frequency variation in prices associated with temporary price changes, there is much less low-frequency variation, which is ultimately what matters for how aggregate prices respond to low-frequency variation in monetary policy.

On the empirical side, our work here is most closely related to that of Bils and Klenow (2004) and Nakamura and Steinsson (2008). As we have noted, Bils and Klenow show that the frequency of all price changes is fairly high, about once every 4.5 months in the BLS data. Nakamura and Steinsson study the same data and show that once temporary price cuts are removed, prices change infrequently, about every 8–11 months. The implicit rationalization 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.

Some may interpret our results as providing a theoretical rationale for removing temporary price cuts from the data, as Nakamura and Steinsson (2008) and Golosov and Lucas (2008) have done. Thus, here we briefly use our theory to evaluate a similar procedure—discarding all temporary changes in favor of just regular prices. We find that compared to our model, this regular price procedure only slightly overstates the degree of aggregate price

stickiness.

Our work is also related to the empirical work of Hosken and Reiffen (2004) and Eichenbaum, Jaimovich, and Rebelo (forthcoming), who document that there is significant low-frequency price stickiness in the micro data despite the large high-frequency price variation.

On the theory side, Guimarães and Sheedy (forthcoming) offer an alternative explanation for temporary price discounts (sales) arising from firms pursuing mixed-price strategies. Finally, Rotemberg (forthcoming) 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 our two U.S. data sets: the BLS monthly data and the Dominick’s weekly data. Here we describe several regularities, or *facts*, that we see in these data. These facts help clarify the distinction between temporary and regular price changes and illustrate their properties. We use these facts to motivate our model.

A. The Data Sets

The BLS data set is the CPI Research Database 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. For our use in this work, Emi Nakamura and Jon Steinsson kindly computed statistics for us by applying an algorithm we devised (which we describe in the appendix) to the CPI Research Database. These statistics are available in their forthcoming work.

The Dominick’s data set includes nine years (1989–97) of weekly store-level reports from 86 stores in the Chicago area on the prices of more than 4,500 individual products, organized into 29 product categories. The products available in this data base range from nonperishable foodstuffs (for example, frozen and canned food, cookies, crackers, juices, sodas, and beer) to various household supplies (for example, detergents, fabric softeners, and

bathroom tissue) as well as pharmaceutical and hygienic products. (For a detailed description of the data and Dominick’s pricing practices, see the work of Hoch, Drèze, and Purk (1994), Peltzman (2000), and Chevalier, Kashyap, and Rossi (2003).)

B. Categories of Price Changes

To identify a pattern of price changes in the data, we wrote a simple algorithm which 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. The algorithm is somewhat involved, so we relegate a formal description to the appendix.

C. The Facts

In Figure 2, we illustrate the results of applying our algorithm to several particular price series from the Dominick’s data. On each of the four graphs, for each of the four products, the dashed lines are the raw data (the original posted prices), and the solid lines are the regular price series constructed with our algorithm. Just a glance at these graphs makes several facts about price changes clear: across the board, price changes are frequent and large, but most of them are temporary, and most temporary prices return to the pre-existing regular price.

Table 1 reports statistics summarizing the facts about price changes that result from applying the algorithm. The first column of data is statistics from the BLS data set, in which

all statistics are computed at a monthly frequency. We report revenue-weighted averages of the corresponding statistics at the level of product categories. (The product level statistics are available from Nakamura and Steinsson (forthcoming).) The second column of data in Table 1 is statistics from the Dominick's data set, in which all statistics are computed at a weekly frequency and each good weighted by its revenue share. The third column is statistics constructed by us from the Dominick's data by sampling the weekly data at a monthly frequency; here, as with the BLS data, we do not weight individual products by their revenue share.

Among the monthly BLS data statistics, we highlight some of the key features that motivate our model. First, most price changes in the data are temporary: 72% of price changes are temporary, with an average duration of about two months ($1/.53$). Second, about 50% of the time, the nominal price after a temporary price change returns to the exact nominal level it had before the change (the old regular price). Because of this feature of the data, although overall the frequency of price changes is large (22% of all prices change every month, so the average duration is 4.5 months), the frequency of regular price changes is much smaller, 7% per month, with an average duration of about 14.5 months.³

The feature of the data that will become especially significant to this analysis is that micro prices have a subtle type of low-frequency price stickiness. For example, the table shows that 75% of the time during a year, firms charge exactly the same price for a good, namely, its annual mode. (Note that this statistic does not depend on our algorithm for identifying regular price changes.) When we combine that fact with the fact that firms change prices once every 4.5 months, we see that prices tend to come back often to the same nominal level. We will show, using our model, that it is this feature of the data that allows monetary shocks to have sizable effects despite the high frequency of price changes.

Consider now the second data column in Table 1, where we report similar statistics for the Dominick's weekly data. The basic patterns evident in the BLS data are even more evident here. Nearly all price changes are temporary (94%), and after such changes, 80% of these prices come back to the pre-existing price. As a result, even though prices change relatively frequently, once every 3 weeks, the regular price changes occur much less often, once every 8 months. Unlike the BLS data, the Dominick's data also contain information

on quantities. Using those data, we find that a disproportionate fraction of goods are sold during periods of temporary prices: even though temporary price changes occur only about a quarter of the time, almost 40% of goods are sold during these periods.

When comparing the Dominick's data to the BLS data, note that some of the difference is coming from the frequency of sampling: Dominick's data are sampled weekly; the BLS data, monthly. To illustrate the role of sampling, we also sampled the Dominick's data at the monthly frequency and report the results in the third data column of Table 1. Doing so dramatically lengthens the implied duration of price spells, from 3 weeks ($1/.33$ weeks) in the weekly data to nearly 3 months ($1/.36$ months) in the monthly data. This difference illustrates our contention that, at least with Dominick's data, the monthly sampled prices miss many of the high-frequency movements in prices that are reversed within a month.

2. A Menu Cost Model with Temporary Price Changes

Now we build a menu cost model with temporary price changes and use it to evaluate the relationship between the frequency of micro price changes and the degree of aggregate price 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 model does.

Our model is a simple extension of the standard menu cost model of Golosov and Lucas (2007). To account for the pattern of temporary and regular price changes in the data, we make two additional assumptions.

First, motivated in part by the work of Zbaracki et al. (2004) on the pricing practices of firms, we assume that firms choose two prices for each good: a *list* price P_{Lt} as well as a *posted* price P_t that the consumer faces. Intuitively, 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 takes a costly action to make it differ. Formally, the list price is relevant because every time the firm posts a price that differs from the list price, the firm must incur a fixed cost ϕ . As a result, the posted price will deviate from the list price infrequently, only when the benefit from doing so exceeds the fixed cost. We assume that changing list prices is costly and entails a fixed cost

κ . (For simplicity, we set the cost of changing the posted price to zero and in an appendix available on request show that allowing for such a cost has virtually no effect on our results.)

Second, we allow for both transitory and permanent idiosyncratic productivity shocks. The transitory and permanent shocks will help the model deliver the temporary and regular price changes in the data. As Golosov and Lucas (2007) do, we think of these shocks as a stand-in for all the idiosyncratic forces that make changing prices optimal for firms.

A. Setup

Formally, we study a monetary economy populated by a large number of infinitely lived consumers and firms and a government. In each time period t , this economy experiences one of finitely many events s_t . We denote by $s^t = (s_0, \dots, s_t)$ 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 and idiosyncratic shocks to a firm's productivity. In terms of the money supply shocks, we assume that the (log of) money growth follows an autoregressive process of the form

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

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 . We describe the idiosyncratic shocks below.

Consumers and Technology

In each period t , the commodities in this economy are labor, money, a continuum of intermediate goods indexed by $i \in [0, 1]$, and a final good. The final good is used for consumption and investment as well as for materials used in production by intermediate good firms.

In this economy, consumers consume, trade bonds, work, and hold real money balances. They also own the capital stock and rent it to intermediate good producers. The consumer problem is to choose consumption $c(s^t)$, nominal labor $l(s^t)$, investment $x(s^t)$, nominal money

balances $M(s^t)$, and a vector of bonds $\{B(s^t, s_{t+1})\}_{s_{t+1}}$ to maximize utility

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

subject to the budget constraint

$$\begin{aligned} & P(s^t) \left[c(s^t) + x(s^t) + \frac{\xi}{2} \left(\frac{x(s^t)}{k(s^{t-1})} - \delta \right)^2 k(s^{t-1}) \right] + M(s^t) + \sum_{s^{t+1}} Q(s^{t+1}|s^t) B(s^{t+1}) \\ & \leq W(s^t) l(s^t) + \Pi(s^t) + M(s^{t-1}) + B(s^t) + R(s^t) k(s^t), \end{aligned}$$

where $P(s^t)$ is the price of the final good, $x(s^t) = k(s^t) - (1 - \delta)k(s^{t-1})$ is investment, $W(s^t)$ is the nominal wage, $\Pi(s^t)$ are nominal profits, and $R(s^t)$ is the rental rate of capital. Note that capital is subject to adjustment costs, the size of which is determined by ξ . Here we have complete, contingent, one-period nominal bonds. We let $B(s^{t+1})$ denote 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 $t + 1$. We let $Q(s^{t+1}|s^t)$ denote the price of this bond in period t and state s^t .

Consider, next, the technology for the intermediate good producers. The producer of intermediate good i uses the technology

$$(3) \quad y_i(s^t) = a_i(s^t) z_i(s^t) [k_i(s^t)^\alpha l_i(s^t)^{1-\alpha}]^\nu m_i(s^t)^{1-\nu},$$

where $y_i(s^t)$ is the output of good i , $k_i(s^t)$ is the amount of capital it rents, $l_i(s^t)$ is the labor input, and $m_i(s^t)$ is the amount of materials used in the production process. The good-specific productivity has two components: a transitory one, $z_i(s^t)$, and a permanent one, $a_i(s^t)$. The transitory component follows a Markov process. The permanent component evolves according to

$$(4) \quad \log a_i(s^t) = \log a_i(s^{t-1}) + \varepsilon_i(s^t),$$

where $\varepsilon_i(s^t)$ is an i.i.d. shock (or *innovation*) to the permanent part of productivity with a distribution that we describe below. To keep the model stationary, we assume that intermediate good firms exit with probability ρ_e and are replaced by firms with $a_i(s^t) = 1$. Below we describe the problem that the intermediate good firms solve.

Consider, last, the final good producers. These firms, who are perfectly competitive, purchase a continuum of intermediate goods and sell a final good to consumers and intermediate good firms. The problem of a final good firm is to choose the amount of each intermediate good $y_i(s^t)$ to purchase in order to maximize

$$P(s^t) y(s^t) - \int_0^1 P_i(s^t) y_i(s^t) di$$

subject to the final good production function,

$$y(s^t) = \left[\int y_i(s^t)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}.$$

We define $P(s^t)$ as the price of the final good, $P_i(s^t)$ as the price of good i purchased from an intermediate good firm, and θ as the elasticity of substitution among intermediate inputs. The solution to this problem is, then,

$$(5) \quad y_i(s^t) = \left(\frac{P_i(s^t)}{P(s^t)} \right)^{-\theta} y(s^t),$$

which will be the demand function faced by the producer of intermediate good i . Here,

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

is the minimum cost of producing one unit of the final good and, because of perfect competition, the final good's price.

The Intermediate Good Firm Problem

Consider, now, the problem of an intermediate good firm in this economy. The firm has a fixed cost, measured in units of labor, of changing its list prices. We refer to this cost

as a *menu cost*. Let $P_{L,i}(s^{t-1})$ denote the firm's list price from the previous period. This list price is a state variable for the firm at the subsequent state s^t . The firm has two sets of pricing decisions. It can leave the list price unchanged at no cost, or it can pay a fixed cost κ and change the list price to $P_{L,i}(s^t)$. A firm can then either pay nothing and charge the list price $P_i(s^t) = P_{L,i}(s^t)$ or pay ϕ and charge any price other than the list price. (We think of κ as an upper-level managerial cost of changing the list price and ϕ as the cost of a lower-level manager deviating from that list price, say, by offering a temporary discount.)

In our model, firms face a mixture of idiosyncratic shocks—permanent and transitory. Here firms typically use a list price change to respond to the more permanent shocks and temporarily deviate from this list price in order to respond to transitory shocks.

To write the firm's problem formally, first note that the firm's period nominal profits, excluding fixed costs at price $P_i(s^t)$, are

$$D(P_i(s^t); s^t) = (P_i(s^t) - V_i(s^t))y_i(s^t),$$

where

$$(6) \quad V_i(s^t) = \frac{\zeta \left(R(s^t)^\alpha W(s^t)^{1-\alpha} \right)^\nu P(s^t)^{1-\nu}}{a_i(s^t) z_i(s^t)}$$

is the nominal cost of producing one unit of intermediate good i and $y_i(s^t)$ is given by (5), where ζ is a constant that depends on the parameters of the production function. The present discounted value of profits of the firm, expressed in units of period 0 money, is given by

$$(7) \quad \sum_t \sum_{s^t} Q(s^t) (1 - \rho_e)^t [D(P_i(s^t); s^t) - W(s^t) (\kappa \delta_{L,i}(s^t) + \phi \delta_{T,i}(s^t))],$$

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 (7), 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.

Equilibrium

Consider, now, this economy's market-clearing conditions and the definition of *equilibrium*. The market-clearing condition on labor,

$$l(s^t) = \int_i [l_i(s^t) + \kappa \delta_{L,i}(s^t) + \phi \delta_{T,i}(s^t)] di,$$

requires that the sum of the labor used in production and the costs of making both list and temporary price changes equals total labor. The market-clearing condition for the final good is

$$c(s^t) + x(s^t) + \frac{\xi}{2} \left(\frac{x(s^t)}{k(s^{t-1})} - \delta \right)^2 k(s^{t-1}) + \int_0^1 m_i(s^t) di = y(s^t).$$

Note, for later use, that real output (value-added) at base period prices can be written as

$$P(s^0) \left(y(s^t) - \int_0^1 m_i(s^t) di \right).$$

The market-clearing condition on bonds is $B(s^t) = 0$.

An *equilibrium* for this economy is a collection of allocations for consumers $\{c_i(s^t)\}_i$, $M(s^t)$, $B(s^{t+1})$, $k(s^t)$, $x(s^t)$, and $l(s^t)$; prices and allocations for firms $\{P_i(s^t), y_i(s^t), l_i(s^t), k_i(s^t), m_i(s^t)\}_i$; and aggregate prices $W(s^t)$, $P(s^t)$, and $Q(s^{t+1}|s^t)$, all of which satisfy the following conditions: (i) the consumer allocations solve the consumers' problem; (ii) the prices and allocations of firms solve their maximization problem; (iii) the market-clearing conditions hold; and (iv) the money supply processes satisfy the specifications above.

For convenience, we write the equilibrium problem recursively. At the beginning of s^t , after the realization of the current monetary and productivity shocks, the state of an individual firm i is characterized by its list price in the preceding period, $P_{L,i}(s^{t-1})$; its permanent productivity component, $a_i(s^t)$; and the transitory productivity component, $z_i(s^t)$. We normalize all of the nominal prices and wages by the current money supply and let $p_{L,-1,i}(s^t) = P_{L,i}(s^{t-1})/M(s^t)$ and $v_i(s^t) = V_i(s^t)/M(s^t)$ and use similar notation for other prices. With this normalization, we can write the state of an individual firm i in s^t as $(p_{L,-1,i}(s^t), a_i(s^t), z_i(s^t))$.

Let $\lambda(s^t)$ denote the measure over all firms of these state variables. The only aggregate uncertainty is money growth, and the process for money growth is autoregressive; therefore, the aggregate state variables are $[\mu(s^t), \lambda(s^t)]$. Dropping explicit dependence of s^t and i , we write the state variables of a firm as $(p_{L,-1}, a, z)$ and the aggregate state variables as $S = (\mu, \lambda)$. Let the static gross profit function, normalized by the current money supply M , be denoted by

$$(8) \quad d(p_i, a_i, z_i, S) = (p_i - v(a_i, z_i, S))y(p_i, S),$$

where $v(a_i, z_i, S)$ is the unit cost of producing good i and $y(p_i, S)$ is the quantity demanded of good i . Let $\lambda' = \Lambda(\lambda, S)$ denote the *transition law* on the measure over the firms' state variables.

Because of discounting, firms never change a list price in a given period without selling at that price during that particular period. In any period, therefore, a firm has three relevant options: leave its old list price unchanged and pay 0 to charge a price equal to its list price (N), leave its old list price unchanged and pay ϕ to charge a price different than its list price (T), or pay κ to change the list price and sell at that new list price (L).

In any period, the value of a firm that does nothing (N)—does not change its price and sells at its existing list price—is

$$V^N(p_{L,-1}, a, z; S) = d(p_{L,-1}, a, z, S) + (1 - \rho_e)E \left[\sum_{S'} Q(S', S) V(p_{L,-1}, a', z'; S') | a, z \right].$$

(Here the expectations are taken only with respect to the idiosyncratic shocks a and z . Since these shocks are idiosyncratic, the risk about their realization is priced in an actuarially fair way. Of course, our formalization is equivalent to having an intertemporal price defined over idiosyncratic and aggregate shocks and then simply summing over both of those.)

The value of a firm that charges a temporary (T) price $p_T \neq p_{L,-1}$ is

$$V^T(p_{L,-1}, a, z; S) = \max_{p_T} [d(p_T, a, z, S) - \phi w(S)] + (1 - \rho_e)E \left[\sum_{S'} Q(S', S) V(p_{L,-1}, a', z'; S') | a, z \right],$$

and that of a firm that changes its list (L) price is

$$V^L(a, z; S) = \max_{p_L} [d(p_L, a, z, S) - \kappa w(S)] + (1 - \rho_e) E \left[\sum_{S'} Q(S', S) V(p_L, a', z'; S') | a, z \right],$$

where $V(p_{L,-1}, a, z, S) = \max(V^N, V^T, V^L)$.

Inspection of the value function V^T makes clear that the optimal temporary price is static and is chosen so that the marginal gross profit $d_p(p, a, z, S) = 0$, so that the optimal temporary price is simply

$$(9) \quad p_{T,i} = \frac{\theta}{\theta - 1} v_i(S).$$

Note that this temporary price is a simple markup over the nominal unit cost of production and is the *frictionless* price, that is, the price which a flexible price firm would charge when faced with such a unit cost. In contrast, if the list price is changed, then the optimal pricing decision for the new list price, p_L , is dynamic. In particular, p_L will not typically equal p_T .

As (9) makes clear, if a price is changed temporarily, then the inherited list price $p_{L,-1}$ does not affect the temporary price, so we can write $p_T(a, z, S)$. Similarly, as inspection of the value function V^L makes clear, conditional on having a list price change, the inherited list price $p_{L,-1}$ is also irrelevant, so we can write $p_L(a, z, S)$.

The Workings of the Model

Our model works differently from existing menu cost models because of the ability of a firm in our model to use temporary price changes to respond to shocks. To provide intuition for our model's predictions, we now describe the firm's optimal decision rules—in particular, when the firm chooses to make a temporary price change and when it chooses to make a list price change. Briefly, we show that firms use changes in the temporary price primarily to respond to temporary shocks and use changes in the list price to respond to permanent shocks.

Consider the firm's optimal decision rules in a quantitative version of the menu cost model (the details of which will be described later). These rules are a function of the individual states, namely, the normalized list price $p_{L,-1} = P_{L,-1}/M$, the permanent productivity level

a , and the transitory productivity level z , as well as the aggregate state variable—the money supply growth rate—and the distribution of firms λ .

We illustrate the firm’s optimal decision rules in Figure 3. This figure shows two decision rules in the $(z, p_{L,-1})$ space for a given level of permanent productivity: the list price $p_L(z)$, conditional on the firm’s choice to change the list price, and the temporary price $p_T(z)$, conditional on the firm’s choice to set a temporary price. Note that since the permanent productivity shock has a unit root, the list price and the temporary price are nearly identical. Recall also that the temporary price is identical to the frictionless price, the price that would prevail if prices were not sticky. This frictionless price decreases one-for-one with productivity shock increases.

Figure 3 also shows the regions of the state space in which the firm optimally chooses to make a list price change (L), to make a temporary price change (T), or to not change its price (N). The figure shows that if the current list price $p_{L,-1}$ is close enough to the frictionless price (that is, if the price lies in the region labeled N), then the optimal decision for the firm is to forgo paying any costs and just charge the old list price.

This figure also shows that a firm chooses to charge a temporary price when two conditions are met: the temporary shock is either sufficiently high or sufficiently low, and the old list price is close to what would be optimal when the temporary shock is at its mean. This pattern of temporary price-setting arises from two features of our quantitative model: the cost of changing the list price is not that much higher than the cost of deviating from the list price, and temporary productivity shocks have high mean reversion. Briefly, if the list price is far from the level that is optimal when the temporary shock is at its mean, then the firm realizes it will likely have to change its list price soon anyway, so it simply changes its list price in order to respond to the transitory shock.

Figure 4 shows a simulation of shocks and a firm’s decision rules in our quantitative model for 7 years (84 months). Panel A shows the three shocks. For convenience, we report the inverse of the log of the permanent and transitory productivity shocks (so that these shocks can be interpreted as percentage changes in costs) and the log of the money supply. Panel B shows the posted and list prices. Comparing these two panels, we see that, for the most part, firms seem to change list prices in order to offset permanent productivity shocks

and money shocks and temporarily deviate from the list prices in order to offset transitory productivity shocks.

Panel C shows how posted prices differ from frictionless prices in the model. Clearly, even though posted prices change frequently, they do not offset all the money shocks. For example, from month 50 to month 80, posted prices change 7 times. Even so, the posted price at the end of this period is considerably above the frictionless price, and this difference is mostly due to a persistent decline in the money supply.

Panel D makes clear that the regular prices constructed by our algorithm often agree with the list prices from the theoretical model, although not always.

B. Quantification and Prediction

We want to use the facts about price changes that we have isolated in the two U.S. data sets as the basis for our model and its evaluation. To do that, we must quantify the model with values from U.S. data. Here we describe how we choose the model's functional forms and parameter values. We then investigate whether our parsimonious model can be made to 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

$$U\left(c, \frac{M}{P}, l\right) = \frac{\eta}{\eta - 1} \log\left(\omega c^{\frac{\eta-1}{\eta}} + (1 - \omega) \left(\frac{M}{P}\right)^{\frac{\eta-1}{\eta}}\right) - \psi l.$$

We choose $\eta = .39$ and $\omega = .94$. As Chari, Kehoe, and McGrattan (2000) show, such a choice of parameters is consistent with estimates from the money demand literature. We choose the value of ψ , the disutility of labor parameter, to ensure that without aggregate shocks, consumers supply one-third of their time to the labor market. We set the depreciation rate of capital δ to be consistent with 10% annual depreciation. We choose the adjustment cost on capital ξ so that the standard deviation of investment relative to that of output is equal to 3, a number similar to that in the U.S. data.

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 the elasticity of substitution in the literature. (See, for example, Nevo 1997 and Chevalier, Kashyap, and Rossi 2003.) For the intermediate good production function, we set α to be $1/3$ and $1 - v = .7$. These values imply that the capital share in value-added is $1/3$ and the material share in gross output is 47%. (Recall that to calculate the material share in gross output, we need to divide $1 - v$ by the markup $\theta/(\theta - 1)$.) These are consistent with the work of Basu (1995) and Nakamura and Steinsson (forthcoming). We choose the exit rate of intermediate good firms $\rho_e = 1.8\%$. Gertler and Leahy (2008) argue that this probability of exit is consistent with the rate at which products are replaced in the Nakamura and Steinsson (2008) data.

We would like to isolate the real effects of exogenous monetary policy shocks as a simple way of measuring the degree of nominal rigidity in the model. A popular way to do so is as in the work 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.⁴ 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 policy shocks.⁵ 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 $\sigma_m = .0018$. (We also redid the analysis in this paper using a Taylor rule and obtained similar results that are available on request. We focused on the money growth rule because the degree of aggregate price stickiness is easier to interpret under a money growth rule than under an interest rate rule.)

The rest of the parameters are chosen so that the model can closely reproduce the salient features of the micro price data we have described: κ , the (menu) cost the firm incurs when changing its list price; ϕ , the cost of deviating from the list price; and the specifications of the productivity shocks.

Consider, first, the specification of the permanent productivity shocks. The distribution of the innovations $\varepsilon_i(s^t)$ for these requires special attention. Midrigan (2007) shows that when $\varepsilon_i(s^t)$ is normally distributed, a model like ours generates counterfactually low disper-

sion 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 Poisson probability λ_a and are, conditional on arrival, uniformly distributed on the interval $[-\bar{\nu}, \bar{\nu}]$. We follow this approach and assume that

$$\varepsilon_i(s^t) = \begin{cases} \nu_i(s^t) & \text{with probability } \lambda_a \\ 0 & \text{with probability } 1 - \lambda_a, \end{cases}$$

where $\nu_i(s^t)$ is distributed uniformly on the interval $[\underline{\nu}, \bar{\nu}]$. The productivity process thus has two parameters: the arrival rate of shocks λ_a and the support of these shocks $\bar{\nu}$.

Paying special attention to the distribution of these 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 productivity shocks are normally distributed. But as Midrigan (2007) shows, with a fat-tailed distribution of productivity 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.

Consider, next, the process for the transitory productivity shocks. To keep the model simple, we assume that these shocks follow a Markov chain, with $z_t \in \{-\bar{z}, 0, \bar{z}\}$ referred to as the *low*, *medium*, and *high* productivity values, with transition probabilities

$$\begin{bmatrix} \rho_s & 1 - \rho_s & 0 \\ \rho_l & 1 - \rho_l - \rho_h & \rho_h \\ 0 & 1 - \rho_s & \rho_s \end{bmatrix}.$$

Here, the subscripts l and h indicate the low and high productivity values. Hence, ρ_l is the probability of experiencing a decrease in productivity from 0 to $-\bar{z}$, and ρ_h is the probability of experiencing an increase in productivity from 0 to \bar{z} . Finally, ρ_s is the probability of staying in a non-medium state. Our parameterization of these shocks thus has four parame-

ters $\{\bar{z}, \rho_s, \rho_l, \rho_h\}$. Here, ρ_l and ρ_h govern the probability of temporary price increases and decreases, while ρ_s determines the duration of temporary price changes.

We choose these parameters to minimize the distance between 13 moments in the data and the model listed in panel A of Table 2. These 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. These moments include the frequency of monthly price changes (posted and regular), the fraction of price changes that are temporary, the fraction of periods with temporary price discounts, the proportion of returns to the old regular price, the probability of a temporary price spell ending, the fraction of prices at and below the annual mode, as well as the size and dispersion of price changes (as measured by the interquartile range, or IQR), both including and excluding temporary price changes.

Panel B of Table 2 lists the parameter values that have allowed the model to best match the moments in the data. The menu cost κ of changing a list price is 0.81% of a firm's steady-state profits, while the cost of deviating from this list price ϕ is 0.68% of a firm's steady-state profits. Permanent productivity shocks arrive with probability λ_a of 7.8% and have an upper bound of $\bar{\nu} = .184$. For the transitory productivity shock, we choose $\bar{z} = .143$. The parameters governing the Markov transition matrix are $\rho_s = .52$, $\rho_h = .038$, and $\rho_l = .026$. Thus, the medium productivity state is most persistent, whereas firms that are in the low- or high-productivity states expect to return to the medium with high probability $1 - \rho_s = .48$.

The Micro Moments in the Data and the Model

Returning to panel A of Table 2, we see that our parsimonious extension of a standard menu cost model can account well for the micro moments of the BLS data. Recall that in the data we computed statistics about regular prices using our algorithm. We use the same algorithm 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.)

The frequency of posted price changes is high: 22% in both the data and the model; the frequency of regular price changes is much lower: 6.9% in both the data and the model.

Most price changes are temporary: 72% in the data and 76% in the model. Temporary prices often return to the regular price that existed before the temporary change: 50% in the data and 70% in the model. We also see that temporary price changes are transitory: the probability that a temporary price spell ends is equal to 53% in the data and 54% in the data. Periods with temporary prices account for 10% in the data and 11% in the model, and most of these periods are ones with temporary price declines (6% in both the data and the model).

So far we have used a specific algorithm to distinguish between temporary and regular price changes. We think of these price changes as characterizing high- and low-frequency price variation. We would like to argue that our broad conclusions carry through for other ways of distinguishing high- and low-frequency price variation. Another way to see that there is considerable low-frequency price stickiness is to consider the fraction of times a good's price is equal to the annual mode. (This alternative measure of low-frequency price stickiness has been suggested by Hosken and Reiffen (2004) and was also used recently by Eichenbaum, Jaimovich, and Rebelo (forthcoming).) We apply this measure to both the data and the model, and in both it is quite high: 75% in the data and 73% in the model. When prices are not at their annual mode, they tend to be below the annual mode more often than above it in both the data and the model (13% of the time in the data and about 17% in the model). In sum, we see that our conclusions that the data exhibit considerable low-frequency price stickiness are robust to alternative ways of examining the data.

Following Golosov and Lucas (2007) and Midrigan (2007), 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% for all of them). So is the dispersion of these changes as measured by the IQR: 9% for all price changes in both the model and the data and 8% for regular price changes in both.

C. A Comparison with the Standard Model

We next compare the patterns of low- and high-frequency price stickiness in the standard model which has one type of price change with the same patterns in our model. We show that, unlike our model, the standard model cannot simultaneously reproduce the high-

frequency price flexibility (that is, the low level of high-frequency price stickiness) of micro price changes and the low-frequency price stickiness of micro price changes observed in the data.

To demonstrate that, we consider a sequence of standard menu cost models. In each of these models, we vary the frequency of micro price changes and keep the size and interquartile range of price changes equal to those in the data. We convert this frequency into months and consider it a measure of the degree of high-frequency price stickiness. 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 5. 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 model—is represented in panel A by a large dot. In the BLS data and in our 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 5, are quite 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 menu cost 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 effectiveness of monetary policy, in terms of aggregate price stickiness, relative to what the standard model says. Our model’s main message is that despite the high frequency of micro price changes observed in

the data—and despite what the standard model says—aggregate prices are quite sticky. We attribute this result to the key features of temporary price changes that we have identified. And we evaluate one way in which researchers have tried to continue to use the standard model despite its difficulty with temporary price change data.

A. A Measure of Aggregate Price Stickiness

For this analysis, we must somehow measure the degree of aggregate price stickiness in our model. We want a measure that captures how slowly the aggregate price level reacts to a change in the money supply. The slower is this reaction, the greater will be the difference between the impulse response for the price level and the impulse response for the money supply. Based on that logic, we choose as our measure of aggregate price stickiness the difference (actually, the integral of the difference) between the impulse response of money and prices to an innovation to the money supply relative to the average impulse response of money (over the first two years following the innovation). Notice that our measure has a one-to-one relationship with the real effects of money, as measured by the integral of the impulse response of output to an innovation in the money supply. To express our measure in convenient units, we express it as the frequency of micro price changes in a standard model (without temporary price changes) that produces this same difference between the impulse response of money and prices. We find this measure convenient because the units are in terms of the key parameter of the standard New Keynesian model and is, thus, well understood.

B. Impulse Responses in Our Model

We begin by examining the impulse responses to a monetary policy shock in our model. In Figure 6, we show the impulse responses to a shock to the money growth rate in period 1 that, in the limit, produces a 5% increase in the level of the money supply. (All variables are expressed as percentage deviations from the original steady state.)

Panel A of Figure 6 displays the model's responses of nominal variables: the money supply, the price level, and the nominal marginal cost. Notice that because in this model each intermediate firm uses the final good as an input, the nominal marginal cost responds slowly to the monetary shock. Hence, producers that set their prices based on the present value of these marginal costs also respond slowly.

In panel B of Figure 6, we see that on the shock's impact, the model implies that output increases by about 6.8% and then declines, reaching half of its peak value in about 8 months. Consumption follows a similar pattern, although its response is only half that of output. All other variables (investment, labor, use of intermediate inputs) have a similar qualitative pattern, so are not reported here.

The column in panel A of Table 3 labeled *our model* summarizes some of the key properties of the model. The first entry reminds us that, at the micro level, prices are flexible, in that they change once every 4.5 months. The second entry quantifies the degree of aggregate price stickiness to be about 40%, measured, again, as the average difference between the impulse response of the money supply and the price level over the first 24 months after the impulse divided by the average impulse response of the money supply during this period. To get some sense of this number, note that if aggregate prices never moved after a monetary shock, our aggregate price stickiness measure would be 100%, whereas if prices instantly adjusted to the flexible price level, our measure would be 0%. The table's last two entries quantify how the aggregate price stickiness in our model manifests itself in output: it leads to a sizable output response of about 3% and a half-life of that response of about 8 months.

C. A Comparison of the Degree of Price Stickiness in our Model with that in the Standard Model

For some perspective on our model's aggregate price implications, we now compare them to those of the standard menu cost model used in the literature (as in, for example, the work of Midrigan (2007) and Gertler and Leahy (2008)). The implications are very different.

The standard model we use here is a special case of our model in which we eliminate the option of temporary changes by setting $\phi = \infty$. We also eliminate the transitory productivity shocks to make our analysis parallel to the existing approaches, which use only one type of idiosyncratic shock. We choose parameters governing the permanent productivity process and the size of the menu cost so that the standard model matches the frequency, average size (11%), and interquartile range of price changes (9%) in the data.

In Figure 7 we report how the aggregate degree of price stickiness in the standard model varies with the degree of micro price stickiness. The figure shows that the standard model implies a one-to-one relationship between micro price stickiness and aggregate price

stickiness. In particular, when micro prices change frequently, the degree of aggregate price stickiness is low. Conversely, when micro prices change infrequently, the degree of aggregate price stickiness is high.

This is quite a contrast to the predictions of our model. In the data, prices change every 4.5 months. When the standard model reproduces this high frequency of price changes, as it does at point *A* in the figure, this model implies small amounts of aggregate price stickiness. In contrast, when our model produces this high frequency of price changes, as it does at point *B*, it predicts much larger amounts of aggregate price stickiness (40% vs. 17%).

We can also use Figure 7 to translate our measure of stickiness into more commonly used terms. We ask, To reproduce the degree of aggregate price stickiness in our model, what degree of micro price stickiness would a standard model require? Point *C* in the figure provides the answer: micro prices in the standard model would need to change about once every 12 months.

The impulse responses of the standard model that exactly match the degree of price stickiness in our model are displayed in Figure 8. Clearly, the impulse responses of output and prices in our model are nearly identical to those in a standard model with 12 month micro price stickiness.

In sum, even though our model is consistent with frequent micro price changes, it still predicts a quite sticky aggregate price level.

D. Evaluating the Existing Approach in the Standard Model

As we have emphasized, the micro data exhibit two types of price changes: regular changes, which happen at a moderately low frequency, and temporary deviations from regular prices, which happen at a high frequency. We have dealt with this phenomenon by building a model that accounts for both types of price changes. Other researchers have chosen not to explicitly model both but instead use a standard model with only one type. To deal with temporary changes, these researchers discard all temporary price changes from the data and set the frequency of price changes in the model to match the frequency of regular price changes in the data. This approach, which we refer to as the *regular price* approach, can, at best, only roughly approximate the data. Since the vast majority of quantitative exercises

use this approach, we find it informative to evaluate just how rough an approximation it delivers.

Returning to Figure 7, we see that a standard model using the regular approach predicts a degree of price stickiness of about 46%, or about 1.15 times greater than in our model. We conclude that the regular price approach only slightly overstates the degree of price stickiness and the real effects of monetary policy shocks. We thus find that the results of Golosov and Lucas (2007), Midrigan (2007), and Nakamura and Steinsson (2008), who have used the regular price approach in studying menu cost models, do not greatly overstate the real effects of monetary shocks.

In work available upon request, we find a similar result for the regular price approach in a Calvo model with a constant hazard of price changes: the regular price approach does not greatly overstate the real effects of monetary shocks.

E. Explanation

The key prediction of our model is that temporary price changes, although very frequent, do not allow the aggregate price level to react to monetary policy shocks. We now provide some intuition for this result and argue that this phenomenon is explained primarily by the distinctive features of temporary price changes that we have identified. We also show that strategic interactions emphasized by Guimarães and Sheedy (forthcoming) play no role in our result.

Intuition

We claim that the nature of temporary price changes does not allow the aggregate price level to react to monetary policy shocks. To gain some intuition for that explanation, let us walk through the model's pricing process again.

First, consider the response of the aggregate price level to a one-time change in the log of the money supply equal to Δm in period 1, assuming that each firm's desired posted and list prices increase by this amount. For simplicity, assume that the probability that any given firm changes its list price in any given period is constant and equal to α_L . Similarly, assume that the probability that any given firm has a temporary deviation from its old list price is constant and equal to α_T .

The change in the aggregate price level can be written, to a first approximation, as

$$\Delta p_t = \int \Delta p_{it} di,$$

where Δp_t is the change in the log of the aggregate price level and Δp_{it} , the change in the log of an individual firm's price. In any period after the shock, there are three types of firms: those that have already reset their list prices, those that currently have a temporary change, and those that are still charging the list price in effect prior to the change in the money supply.

The change in the aggregate price level is, therefore, equal to

$$\Delta p_t = \lambda_{L,t} \times \Delta m + \lambda_{T,t} \times \Delta m + (1 - \lambda_{L,t} - \lambda_{T,t}) \times 0 = (\lambda_{L,t} + \lambda_{T,t}) \Delta m,$$

where $\lambda_{L,t}$ is the measure of firms that have reset their list price as of period t and $\lambda_{T,t}$ is the measure of firms that have a temporary change in period t . Clearly, $\lambda_{T,t} = \alpha_T$, since the measure of firms that have a temporary price change is constant. To compute $\lambda_{L,t}$, notice that α_L firms have reset their list prices one period after the shock, $\alpha_L + (1 - \alpha_L)\alpha_L$ in the second period (the firms that have done so in the first period plus those that do so in the second period), $\alpha_L + (1 - \alpha_L)\alpha_L + (1 - \alpha_L)^2\alpha_L$ in the third period, and so on. So we have

$$\lambda_{L,t} = \alpha_L \sum_{i=0}^{t-1} (1 - \alpha_L)^i = 1 - (1 - \alpha_L)^t.$$

Consider now two extreme scenarios. Suppose first that all price changes are list price changes, so that α_L is high and $\alpha_T = 0$. If α_L is high, then the measure of firms that have had a chance to react to the monetary shock quickly increases with t . For example, if $\alpha_L = 0.2$, the fraction of firms that have reset their list price and permanently responded to the money shock is equal to 89% after 10 periods. Hence, if all price changes are list price changes, the aggregate price level is quite flexible.

Now suppose, however, that all price changes are temporary. Since after temporary changes, the nominal price reverts to the pre-existing list price, these changes only allow firms to temporarily react to a monetary policy shock. If $\alpha_L = 0$ and $\alpha_T = 0.2$, the measure of

firms that have responded to the monetary shock is the same, 20%, forever. Thus, if all price changes are temporary price changes, the aggregate price level is quite sticky. Notice that our example makes clear that a key difference between temporary and list price changes is that temporary changes are special: since they return the nominal price back to its pre-existing list price, they allow firms to only temporarily respond to a change in monetary policy.

Finally, the second distinctive feature of temporary changes is their clustering. To see the role of this feature, note that every episode of temporary deviation from a list price involves two price changes: one that initiates the temporary change and another that reverses it. Hence, if $\alpha_T = 0.2$ and $\alpha_L = 0$, the frequency of price changes is twice as high and equal to $2\alpha_T$. The second round of price changes does not allow firms to react to monetary policy shocks and thus does not lower the stickiness of the aggregate price level (even though it does lower the degree of micro price stickiness).

On the Unimportance of Strategic Interactions

Based on recent work by Guimarães and Sheedy (forthcoming), some may conjecture that the intuition for our results comes from strategic interactions in price-setting, interactions that arise from our use of materials as a factor of production and from our use of capital and interest-elastic money demand. In particular, some may think that the reason our model predicts a large amount of aggregate price stickiness is that a given firm that undertakes a temporary price change in our model has little incentive to respond to a money shock. The logic is that since this given firm purchases inputs from other firms with sticky prices, this firm's input costs are sticky. We show that this conjecture is false: our results are not driven by strategic complementarity in price-setting.

To see this, consider an economy in which there is no such complementarity. In this *no strategic complementarities* economy, we exclude materials and capital from production and make money demand interest-inelastic by introducing money using a cash-in-advance constraint. To see that these changes eliminate any strategic complementarities, note that, given our log-linear preferences, the nominal unit cost of production is given by the nominal

wage that satisfies

$$(10) \quad W(s^t) = \psi P(s^t) c(s^t) = \psi M(s^t),$$

where the second equality follows from the cash-in-advance constraint. Equation (10) makes clear the precise sense in which the model has no strategic interactions in price-setting: the desired price of any given firm is a function of only an exogenously given process, the money supply, and not of the actions of other firms.

We report the results from this economy in panel B of Table 3. As we did with our original economies, we here consider two versions of the no strategic complementarities economy: one *with* temporary price changes and one *without* them. Again we ask, What degree of micro price stickiness in the economy without temporary changes reproduces the degree of aggregate price stickiness in the model with temporary changes? Table 3 shows that the answer is again about once every 12 months.

We therefore conclude that the special nature of temporary price changes in our original model, rather than general equilibrium considerations, accounts for our results. This stands in sharp contrast to the results of Guimarães and Sheedy (forthcoming), who emphasize the role of strategic substitutability of sales. In the variation of our model considered here, a given firm's incentive to have a temporary price change is independent of the actions of other firms, and we nevertheless find that temporary price changes do not greatly contribute to the flexibility of the aggregate price level.

4. Robustness Exercises Using Dominick's Data

We now consider two extensions of our model that use statistics from the Dominick's data set rather than the BLS data set. Both extensions provide a check on the robustness of our results. In one extension, we simply reparameterize our model to be consistent with Dominick's data. In the other, we also replace the transitory productivity shocks in our model with transitory shocks to the elasticity of demand for the firm's product, leaving all else unchanged. We find that our results based on the BLS data are robust to both of these extensions.

A. Weekly Data with Quantities

Our use of the Dominick's data is motivated by two concerns about the BLS data. One concern is that because the BLS data are monthly, by construction, they miss much of the high-frequency movements in prices that are reversed within a month (such as temporary sales). Because the BLS data set misses these movements, when we use it to set parameters in our model, the resulting model may vastly overstate the degree of price stickiness. Our second concern is that since the BLS data set has no information on quantities, it leaves open the possibility that almost all purchases are made when goods are on sale, so that the prices when goods are not on sale, the regular prices, are essentially irrelevant. This possibility might also imply that our exercise with BLS data, which is silent on the quantity facts, vastly overstates the degree of price stickiness.

The Dominick's data set allows us to address both of these concerns. Since these data are weekly (and since Dominick's resets its prices only once a week), this data set does not miss high-frequency movements in prices. Also, the quantity information in the Dominick's data set allows us to investigate whether the vast majority of purchases are made when a good is on sale.

B. Model with Demand Shocks

We have described the Dominick's data set in detail earlier, so we begin this analysis by describing our other extension: a modification to our model to include a different kind of shock, a demand shock.

This extension is motivated by theory and data. The motivation from theory is that a common explanation in the industrial organization literature for temporary price changes is intertemporal price discrimination in response to time-varying price elasticities of demand. In particular, the idea is that firms willingly lower markups in periods during which a large number of buyers of products happen to have high elasticities.

The motivation from data to include a demand shock is based on two observations. First, temporary price cuts are associated with reductions in price-cost margins. (See, for example, the work of Chevalier, Kashyap, and Rossi (2003).) Second, as we have documented, periods with temporary price cuts account for a disproportionately large fraction of goods

sold. Taken together, these features suggest that, in the data, the demand elasticity that firms face is time-varying, and this feature leads firms to have time-varying markups.

Motivated by both theory and data, then, we introduce time-varying elasticities by having consumers with differing demand elasticities and by including good-specific shocks to preferences. We assume the same preferences here as in the earlier model, so that utility is still given by (2).

Now our consumers belong to a two-member *household*. Specifically, each household supplies labor as a single entity to the market, but now values consumption according to

$$c(s^t) = \chi^\chi (1 - \chi)^{1-\chi} c^A(s^t)^\chi c^B(s^t)^{1-\chi},$$

where $c^A(s^t)$ and $c^B(s^t)$ are composite goods given by

$$c^A(s^t) = \left(\int_0^1 c_i^A(s^t)^{\frac{\theta-1}{\theta}} di \right)^{\frac{\theta}{\theta-1}}$$

$$c^B(s^t) = \left(\int_0^1 z_i(s^t)^{\frac{1}{\gamma}} c_i^A(s^t)^{\frac{\gamma-1}{\gamma}} di \right)^{\frac{\gamma}{\gamma-1}}.$$

Our interpretation of these preferences is that they represent a household with two members, consumer A and consumer B , who jointly supply labor but have separate consumption bundles. We suppose that consumer B has a higher elasticity of substitution across goods, $\gamma > \theta$. Moreover, consumer B 's weight on the different goods, $z_i(s^t)$, is a random variable. Note that $z_i(s^t)$ is specific to any particular good but common to all households. In this sense, $z_i(s^t)$ is a good-specific demand shock, i.i.d. across goods. We assume here, as earlier, that $z_i(s^t)$ is a Markov process (which we describe in detail below). This structure generates time-varying elasticities of demand in a representative household environment. The general idea of allowing for a type of heterogeneity while avoiding the complication of recording individual state variables is similar to that of Lucas (1992).

The household's problem here is to choose how much consumption to allocate to each consumer, as well as how much to invest, work, and purchase state-contingent securities subject to the following budget constraint:

$$\begin{aligned}
& \int_0^1 P_i(s^t) [c_i^A(s^t) + c_i^B(s^t)] di + P^F(s^t) \left[x(s^t) + \frac{\xi}{2} \left(\frac{x(s^t)}{k(s^{t-1})} - \delta \right)^2 k(s^{t-1}) \right] \\
& + M(s^t) + \sum_{s^{t+1}} Q(s^{t+1}|s^t) B(s^{t+1}) \\
& \leq W(s^t) l(s^t) + \Pi(s^t) + M(s^{t-1}) + B(s^t) + R(s^t) k(s^t),
\end{aligned}$$

where $P^F(s^t)$ is the price of the final good used for investment, and all other variables are as defined in the earlier model.

This problem yields the following consumer demand functions for the different varieties of goods:

$$(11) \quad c_i^A(s^t) = \chi \left(\frac{P_i(s^t)}{P^A(s^t)} \right)^{-\theta} \left(\frac{P^A(s^t)}{P(s)} \right)^{-1} c(s^t)$$

$$(12) \quad c_i^B(s^t) = (1 - \chi) z_i(s^t) \left(\frac{P_i(s^t)}{P^B(s^t)} \right)^{-\gamma} \left(\frac{P^B(s^t)}{P(s)} \right)^{-1} c(s^t),$$

where the aggregate price index $P(s) = P^A(s)^{\chi} P^B(s)^{1-\chi}$ and

The producer of intermediate good i in this economy uses basically the same technology as before:

$$(13) \quad y_i(s^t) = a_i(s^t) [k_i(s^t)^\alpha l_i(s^t)^{1-\alpha}]^\nu m_i(s^t)^{1-\nu}.$$

One big difference is that we eliminate the transitory productivity process and include only the permanent component of the firm's productivity $a_i(s^t)$ that, as earlier, evolves according to (4).

Consider, next, the final good firms. These firms are perfectly competitive and purchase a continuum of intermediate goods from intermediate good firms and sell a final good to consumers (who use it as investment) and intermediate good firms (that use it as materials). The problem of a final good firm is to choose the amount of each intermediate good $q_i(s^t)$ to

purchase in order to maximize

$$P^F(s^t) q(s^t) - \int_0^1 P_i(s^t) q_i(s^t) di$$

subject to

$$(14) \quad q(s^t) = \left(\int_0^1 q_i(s^t)^{\frac{\theta-1}{\theta}} di \right)^{\frac{\theta}{\theta-1}},$$

which yields the demand function

$$(15) \quad q_i(s^t) = \left(\frac{P_i(s^t)}{P^F(s^t)} \right)^{-\theta} q(s^t).$$

Notice that, for simplicity, we have assumed that the elasticity of substitution for producing the final good, θ , is the same as the elasticity of substitution for type A consumers.

Now consider the intermediate good firm's problem. This type of firm sells to type A consumers, type B consumers, and final good firms. The demand for the intermediate good firm's goods is therefore given by

$$(16) \quad c_i^A(s^t) + c_i^B(s^t) + q_i(s^t),$$

where these functions are given by (11), (12), and (15).

A useful feature of the resulting demand function is that it has time-varying elasticity. Clearly, as the demand shock $z_i(s^t)$ increases, so too does the relative demand for good i coming from the high-elasticity (type B) consumers. Indeed, the total demand elasticity for good i increases with $z_i(s^t)$. A bit of algebra shows that this total demand elasticity increases productivity, specifically with $a_i(s^t)^{\gamma-\theta}$ and $\gamma - \phi > 0$. Intuitively, then, we know that a positive productivity shock leads the firm to lower the price for its good and increases the relative demand of the high-elasticity consumers.

The assumptions about the technology of price adjustment are the same here as in the earlier model. The problem of the intermediate good firm is, too, except that its production function is now given by (13) and the demand function is now given by (16).

The equilibrium definition for this economy is similar to that in our earlier model, except that now the market-clearing condition for the final good is

$$(17) \quad x(s^t) + \frac{\xi}{2} \left(\frac{x(s^t)}{k(s^{t-1})} - \delta \right)^2 k(s^{t-1}) + \int_0^1 m_i(s^t) di = y(s^t)$$

and the market-clearing condition for each intermediate good is

$$(18) \quad c_i^A(s^t) + c_i^B(s^t) + q_i(s^t) = y_i(s^t).$$

Note, for later use, that real output (value-added) at base period prices can be written as

$$P^A(s^0)c^A(s^t) + P^B(s^0)c^B(s^t) + P^F(s^0) \left(y(s^t) - \int_0^1 m_i(s^t) di \right).$$

C. Quantification and Prediction

Now we have two variations of our model with temporary price changes, which we can use to study the Dominick's data. One variation, the *productivity shocks* model, is the model we studied earlier with the BLS data; the other is the economy with transitory demand shocks instead of transitory productivity shocks just described, the *demand shocks* model. We must now parameterize both of these models in order to reproduce features of the Dominick's data before we can have them simulate responses to monetary shocks. In both models, we assume that the length of a period is one week, the frequency with which the Dominick's data are sampled.

In both models, we assign the same functional forms and parameters as earlier for the production function, the rate at which capital depreciates, the rate at which firms exit, and so on. For example, we set $\delta = 1 - (1 - 0.01)^{\frac{1}{4}}$, so that the monthly rate at which capital depreciates is 1%. We set $\rho_e = 1 - (1 - 0.018)^{\frac{1}{4}}$, so that 1.8% of firms exit in any given month (4 weeks). We adjust all other parameters, including the persistence and standard deviation of money growth shocks, similarly. We calibrate the size of the capital adjustment costs to again ensure that the standard deviation of investment relative to that of output is equal to 3, a number similar to that in the U.S. data. All of these parameters are reported in Table 4.

The rest of the parameters are chosen so that the two models can closely reproduce

the salient features of the Dominick’s data. We describe these separately for each model.

Productivity Shocks Model

For the model with productivity shocks only, the key parameters are κ , the (menu) cost the firm incurs when changing its list price; ϕ , the cost of having a temporary deviation from the list price; as well as the specifications of the transitory and permanent productivity shocks. The moments we choose to pin down these parameters are the same moments we used earlier with the BLS data. Here, we also ask the model to account for two quantity facts: the fraction of goods that sold in periods of temporary price changes and the fraction sold in periods with temporary price discounts. We compute these facts from the Dominick’s data by applying an algorithm to that data set similar to the algorithm used with the BLS data, in order to identify regular price changes.

These values, as measured in the data and in the productivity shocks model, are nearly identical. (See Table 5.) Recall that price changes are much more frequent in the Dominick’s data than in the BLS data. The Dominick’s frequency of price changes is 33% per week (in both the model and the data) compared to 22% per month in the BLS data. This higher frequency of price changes reflects the much larger role that temporary price changes play in grocery stores and at the weekly frequency: 94% of price changes are temporary in the Dominick’s data (95% in the model). Notice also that regular price changes are much less frequent here (2.9% per week in both the model and the data), an outcome of the fact that most temporary price changes (80% in the data and 88% in the model) return to the old regular price.

Our productivity shocks model also reproduces well the other moments that describe the pattern of low- and high-frequency price variation in the Dominick’s data. The fraction of periods with temporary prices is high both in the model (24%) and in the data (25%), and most of these are periods with temporary price discounts (20% in both the data and the model). Although prices change very frequently, a substantial proportion of prices are equal to the annual mode (58% in the data, 55% in the model). Deviations from the annual mode are mostly downward (30% of prices are below the annual mode in both the model and the data). Moreover, the model accounts well for the mean and dispersion of the size of

price changes in the data. Notice that price changes are on average somewhat larger (17% in absolute value in the model and in the data) than regular price changes (11% in both the model and the data) and more dispersed.

Notice, finally, in the last two rows of Table 5, that this version of the productivity shocks model does a good job of reproducing the quantity facts in the Dominick's data. In both the data and the model, periods with temporary price changes account for a disproportionate amount of goods sold. Even though prices are temporary 24% of the weeks in the data and 25% of the time in the model, these periods account for 39% of the goods sold in the data and 36%, in the model. Periods in which prices are temporarily below the regular price account for the bulk of these sales: 35% of goods are sold during such episodes in the model and 33% in the data. The reason the model does so well at reproducing these quantity facts is that its demand elasticity ($\theta = 3$) is consistent with the price elasticities of demand in the data.

Again, Table 4 (panel A) reports the parameters that allow the productivity shocks model to achieve this fit. Notice that now the cost of a regular price change (1.89% of the firm's steady-state profits) is roughly twice as high as the cost of temporarily deviating from the list price (1.01% of the firm's steady-state profits) and that both types of shocks are now more frequent relative to those in the BLS experiments.

Demand Shocks Model

For the model with demand shocks, we must choose the new parameters, those that describe the new shocks.

In this demand shock economy, the optimal markup of an intermediate good firm is a function $za^{\gamma-\theta}$. To reduce the dimensionality of the state-space, we assume that $\tilde{z} = za^{\gamma-\theta}$ is a random variable that takes two values, $\tilde{z}_t \in \{0, \bar{z}\}$. (Essentially, this makes the z shock correlated with a in a way that ensures that given \tilde{z}_t , the optimal markup is not separately a function of a .) Moreover, we normalize $\bar{z} = 1$, since this parameter is not separately identified; it plays the same role as χ , the parameter determining the relative weight of

consumers. Finally, we assume that the law of motion for \tilde{z}_t is given by

$$\begin{bmatrix} 1 - \rho_h & \rho_h \\ 1 - \rho_s & \rho_s \end{bmatrix}.$$

Here, ρ_h is the probability that a firm with $\tilde{z}_t = 0$ (a firm that sells only to type *A* consumers) will have $\tilde{z}_{t+1} = 1$ next period (and thus sell to both types of consumers). Moreover, ρ_s is a parameter that governs the persistence of the high demand state.

In addition to ρ_h, ρ_s , we must also choose χ , the weight of type *A* (low-elasticity) consumers, as well as θ and γ , the demand elasticities of the two types of consumers. We set $\gamma = 6$, an upper bound of estimates of demand elasticity in the industrial organization literature, and jointly choose $\chi, \rho_h, \rho_s, \theta$ as well as the parameters determining the size of the price adjustment costs and the evolution of permanent productivity shocks, in order to match the size, frequency, and persistence of temporary price changes and the quantity facts in the Dominick's data set. The set of moments we use to parameterize the model with demand shocks is thus the same as for the productivity shocks model.

Returning to Table 5, we see that the model with demand shocks fits the Dominick's data almost as well as does the productivity shocks model. Once again, this model reproduces well the frequency of regular and all price changes and the size and dispersion of the two types of price changes, as well as the fact that periods with temporary price changes account for a disproportionate fraction of goods sold.

Panel B of Table 4 lists the parameter values that produce this fit. Notice that the weight on high-elasticity (type *B*) consumers, $1 - \chi$, is equal to 0.11, the value required to account for the fact that as much as 39% of the output is sold in periods with temporary price changes (which occur in the model mostly in states with $\tilde{z}_t = 1$, since this state is more transitory). Moreover, a demand elasticity of $\theta = 2$ for low-elasticity (type *A*) consumers is necessary for the model to account for the average size of all price changes in the data of 17%.

D. The Degree of Aggregate Price Stickiness

Finally, we study the degree of aggregate price stickiness in our Dominick’s models to see whether our earlier results using BLS data are supported. We find that they are.

We begin with the productivity shocks model, which simply uses our model and the standard model with parameters chosen to mimic the Dominick data. We first shock money growth in the models in the same way as in the models based on the BLS data: an innovation that leads to an eventual increase of 5% in the money supply. We then calculate the degree of aggregate price stickiness in our (reparameterized) model and find that a standard model with a frequency of price changes of 6.2 months produces the same degree of aggregate price stickiness as in our model. In panels A and B of Figure 9 and in Table 6, we display the responses of prices and output in the two models.

We get a similar result from the demand shocks model. Here, a standard model with a frequency of price changes of 7.3 months produces the same degree of aggregate price stickiness as does the demand shocks model. In panels A and B of Figure 10 and in Table 6, we see that the responses of prices and output to a monetary shock are quite similar in the two models.

We thus conclude that the exact source of price changes (transitory demand or productivity shocks) does not matter much for our result: in both models, the degree of aggregate price stickiness is fairly high—aggregate prices change only every 6 or 7 months—even though micro prices in the Dominick’s data change much more frequently (every 3 weeks).

These results are consistent with our earlier results using the BLS data. Since the frequency of price changes in the Dominick’s data is much higher than in the BLS data (3 weeks vs. 4.5 months), we view this 6- to 7-month number as a conservative lower bound on the degree of price stickiness in the economy as a whole. In this sense, as with the BLS data, we conclude that prices are sticky after all.

5. Conclusion

Standard New Keynesian models imply that if prices change frequently at the micro level, then aggregate prices are not sticky. In the micro data, prices do change frequently, but our models, which are consistent with the micro data, still imply that aggregate prices

are sticky. How is this possible? The answer is that the one-to-one relationship between the frequency of micro price changes and the degree of aggregate price stickiness breaks down when there are two types of price changes, regular and temporary, as there are in the data. Briefly, temporary price changes have two striking features that distinguish them from regular changes: after a temporary change, prices tend to return to their pre-existing levels, and temporary price changes are clustered in time. These features imply that even though temporary price changes happen with high frequency, micro prices also have a type of low-frequency price stickiness that generates stickiness in aggregate prices.

6. Appendix: The Algorithm to Construct the Regular Price

Here we describe, first intuitively and then precisely, our algorithm for constructing a regular price series for each product in the data. We have applied this algorithm ourselves to the Dominick’s data set and have asked Emi Nakamura and Jon Steinsson to apply it to the BLS data set.

Our algorithm is based on the idea that a price is a *regular price* if the store charges it frequently in a window of time adjacent to that observation. We start by computing for each period the mode of prices p_t^M that occur in a window which includes prices in the previous five periods, the current period, and the next five periods.⁶ Then, based on the modal price in this window, we construct the regular price recursively as follows. For the initial period, set the regular price equal to the modal price.⁷ For each subsequent period, if the store charges the modal price in that period, and at least one-third of prices in the window are equal to the modal price, then set the regular price equal to the modal price. Otherwise, set the regular price equal to the preceding period’s regular price.

We want to eliminate regular price changes that occur when the store’s posted price does not change, but only if the posted and regular prices coincide in the period before or after the regular price change. To do that, if the initial algorithm generates a path for regular prices in which a change in the regular price occurs without a corresponding change in the actual price, then we replace the last period’s regular price with the current period’s actual price for each period in which the regular and actual prices coincide. Similarly, we replace the current period’s regular price with the last period’s actual price if the two have coincided in the previous period.

Now we provide the precise algorithm we use to compute the regular price and describe how we apply it.

1. Choose parameters: $l = 2$ (= *lag*, or size of the window: the number of months before or after the current period used to compute the modal price. For the Dominick’s data, we set $l = 5$ weeks), $c = 1/3$ (= *cutoff* used to determine whether a price is temporary), $a = .5$ (= the number of periods in the window with the *available* price required in order to compute a modal price).

We apply the algorithm below for each good separately:

Let p_t be the price in period t ; T , the length of the price series.

2. For each time period $t \in (1 + l, T - l)$,

- If the number of periods with available data in $(t - l, \dots, t + l)$ is $\geq 2al$, then
 - Let $p_t^M = \text{mode}(p_{t-l}, \dots, p_{t+l})$.
 - Let f_t = the fraction of periods (with available data) in this window subject to $p_t = p_t^M$.
- Else, set $f_t, p_t^M = 0$ (missing data).

3. Define the regular price in period t , p_t^R , using the following recursive algorithm:

- If $p_{1+l}^M \neq 0$, then set $p_{1+l}^R = p_{1+l}^M$ (initial value).
- Else, set $p_{1+l}^R = p_{1+l}$ for $t = 2 + l, \dots, T$.
 - If $(p_t^M \neq 0 \ \& \ f_t > c \ \& \ p_t = p_t^M)$, then set $p_t^R = p_t^M$.
 - Else, set $p_t^R = p_{t-1}^R$.

4. Repeat the following algorithm five times:

- Let $\mathcal{R} = \{t : p_t^R \neq p_{t-1}^R \ \& \ p_{t-1}^R \neq 0 \ \& \ p_t^R \neq 0\}$ be the set of periods with regular price changes.
- Let $\mathcal{C} = \{t : p_t^R = p_t \ \& \ p_t^R \neq 0 \ \& \ p_t \neq 0\}$ be the set of periods in which a store charges the regular price.
- Let $\mathcal{P} = \{t : p_{t-1}^R = p_{t-1} \ \& \ p_{t-1}^R \neq 0 \ \& \ p_{t-1} \neq 0\}$ be the set of periods in which a store's last period price was the regular price.
- Set $p_{\{\mathcal{R} \cap \mathcal{C}\}-1}^R = p_{\{\mathcal{R} \cap \mathcal{C}\}}$. Set $p_{\{\mathcal{R} \cap \mathcal{P}\}-1}^R = p_{\{\mathcal{R} \cap \mathcal{P}\}}$.

Notes

¹We define the regular price series precisely using an algorithm described in detail in the appendix. Loosely speaking, the regular price series is a type of running mode of the original prices, and the temporary prices are simply deviations from that running mode.

²These explanations include, for example, search frictions (Butters 1977, Varian 1980, Burdett and Judd 1983), demand uncertainty (Lazear 1986), thick-market externalities (Warner and Barsky 1995), loss-leader models of advertising (Chevalier, Kashyap, and Rossi 2003), and intertemporal price discrimination (Sobel 1984).

³Note that this duration of 14.5 months is higher than the corresponding 8–11 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. Even though temporary price cuts account for most of the temporary price changes, the number of price increases is large enough relative to the number of regular price changes to influence the resulting frequency.

⁴Specifically, 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 policy 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 (VAR) which is consistent with this interest rate rule. That process is well-approximated by an AR(1) similar to the one we use.

⁵The 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.

⁶We do this calculation only if at least one-half of the prices in this window are available.

⁷If in the window around this price more than half of the data is missing, then we set the initial reference price equal to the actual price.

References

- Basu, Susanto. 1995. Intermediate goods and business cycles: Implications for productivity and welfare. *American Economic Review* 85 (3, June): 512–31.
- Bils, Mark, and Peter J. Klenow. 2004. Some evidence on the importance of sticky prices. *Journal of Political Economy* 112 (5, October): 947–85.
- Burdett, Kenneth, and Kenneth L. Judd. 1983. Equilibrium price dispersion. *Econometrica* 51 (4, July): 955–69.
- Butters, Gerard R. 1977. Equilibrium distributions of sales and advertising prices. *Review of Economic Studies* 44 (3, October): 465–91.
- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2000. Sticky price models of the business cycle: Can the contract multiplier solve the persistence problem? *Econometrica* 68 (5, September): 1151–79.
- 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 (1, March): 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, February): 1–45.
- Eichenbaum, Martin S., Nir Jaimovich, and Sergio Rebelo. Forthcoming. Reference prices, costs, and nominal rigidities. *American Economic Review*.
- Gertler, Mark, and John Leahy. 2008. A Phillips curve with an Ss foundation. *Journal of Political Economy* 116 (3, June): 533–72.
- Golosov, Mikhail, and Robert E. Lucas, Jr. 2007. Menu costs and Phillips curves. *Journal of Political Economy* 115 (2, April): 171–99.
- Guimarães, Bernardo, and Kevin D. Sheedy. Forthcoming. Sales and monetary policy. *American Economic Review*.
- Hoch, Stephen J., Xavier Drèze, and Mary E. Purk. 1994. EDLP, hi-lo, and margin arithmetic. *Journal of Marketing* 58 (October): 16–27.
- Hosken, Daniel, and Daniel Reiffen. 2004. Patterns of retail price variation. *Rand Journal of Economics* 35 (1, Spring): 128–46.

Lazear, Edward P. 1986. Retail pricing and clearance sales. *American Economic Review* 76 (1, March): 14–32.

Lucas, Robert E., Jr. 1992. On efficiency and distribution. *Economic Journal* 102 (411, March): 233–47.

Midrigan, Virgiliu. 2007. Menu costs, multi-product firms, and aggregate fluctuations. Working Paper 2007/13, Center for Financial Studies, University of Frankfurt.

Nakamura, Emi, and Jón Steinsson. 2008. Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics* 123 (4, November): 1415–64.

Nakamura, Emi, and Jón Steinsson. Forthcoming. Monetary non-neutrality in a multi-sector menu cost model. *Quarterly Journal of Economics*.

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

Peltzman, Sam. 2000. Prices rise faster than they fall. *Journal of Political Economy* 108 (3, June): 466–502.

Romer, Christina D., and David H. Romer. 2004. A new measure of monetary shocks: Derivation and implications. *American Economic Review* 94 (4, September): 1055–84.

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

Sobel, Joel. 1984. The timing of sales. *Review of Economic Studies* 51 (3, July): 353–68.

Varian, Hal R. 1980. A model of sales. *American Economic Review* 70 (4, September): 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 (2, May): 321–52.

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 (2): 514–33.

Zbaracki, Mark J., Mark Bergen, and Daniel Levy. 2007. The anatomy of a price cut: Discovering the organizational sources of the costs of price adjustment. Mimeo. Univ.W.O.

Table 1: Facts about price changes from two data sets

Statistic	BLS	Dominick's	
	monthly	weekly	monthly
Frequency of all price changes, %	22	33	36
Frequency of regular price changes, %	6.9	2.9	10.6
Fraction of price changes that are temporary, %	72	94	88
Fraction of returns to old regular price, %	50	80	66
Probability that temporary price spell ends, %	53	46	67
Fraction of periods with temp. prices, %	10	24	22
Fraction of periods with price temp. down, %	6	20	15
Fraction of prices at annual mode, %	75	58	67
Fraction of prices below annual mode, %	13	30	23
Fraction of output sold when temp prices, %	n.a.	39	40
Fraction of output sold when price temp. down, %	n.a.	35	35

n.a. = not available

Table 2: Parameterization of Economy with Temporary Price Changes

A. Moments			B. Parameter Values	
	BLS Data	Model	<i>Calibrated</i>	
Frequency of all price changes, %	22	22	Menu cost of regular price change, κ , % SS profits	0.81
Frequency of regular price changes, %	6.9	6.9	Cost of temp. price deviation, ϕ , % SS profits	0.68
Fraction of price changes that are temporary, %	72	76	Arrival rate of permanent shock, λ_a	0.078
Fraction of returns to old regular price, %	50	70	Upper bound of permanent shock, v_bar	0.184
Probability that temporary price spell ends, %	53	54	Size of transitory cost shock, z_bar	0.143
Fraction of periods with temp. prices, %	10	11	Probability of return to medium state, $1-\rho_s$	0.480
Fraction of periods with price temp. down, %	6	6	Probability of entering low state, ρ_l	0.026
Fraction of prices at annual mode, %	75	73	Probability of entering high state, ρ_h	0.038
Fraction of prices below annual mode, %	13	17		
Mean size of price changes, %	11	11	<i>Assigned</i>	
Mean size of regular price changes, %	11	11	Period length	1 month
IQR of all price changes, %	9	9	Probability of exit	0.018
IQR of regular price changes, %	8	8	Annual discount factor	0.96
			AR(1) growth rate of M	0.61
			S.D. of shocks to growth rate of M , %	0.18
			Elasticity of substitution	3

IQR = interquartile range

Table 3: Impulse responses to monetary shock

A. Original BLS Economies

Statistic	Our model: With temporary changes	Standard model: Without temporary changes
Micro price stickiness, months	4.5	11.8
Aggregate price stickiness, %	40.2	40.2
Average output response, %	3.01	3.01
Half-life of output response, months	8.3	7.9

B. BLS Economies with no strategic complementarities

Statistic	Our model: With temporary changes	Standard model: Without temporary changes
Micro price stickiness, months	4.5	11.8
Aggregate price stickiness, %	20.3	20.3
Average output response, %	0.95	0.95
Half-life of output response, months	19.7	18.9

Note: The table reports the impulse response to a cumulative 5% increase in the money stock.

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.

Table 4: Parameter values for models using Dominick's data

A. Productivity shocks		B. Demand shocks	
<i>Calibrated</i>		<i>Calibrated</i>	
Cost of regular price change, % SS profits	1.89	Menu cost of regular price change, % SS profits	1.73
Cost of temp. price change, % SS profits	1.01	Cost of temp. price deviation, % SS profits	0.86
Arrival rate of permanent shock	0.043	Arrival rate of permanent shock	0.045
Upper bound of permanent shock	0.110	Upper bound of permanent shock	0.130
Size of transitory cost shock	0.235	Weight of type <i>B</i> consumers	0.110
Probability of returns to median state	0.450	Probability of returning to 0 demand state	0.450
Probability of entering low state	0.030	Probability of entering 1 demand state	0.147
Probability of entering high state	0.125	Demand elasticity of type <i>A</i> consumers	2.000
<i>Assigned</i>		<i>Assigned</i>	
Period length	1 week	Period length	1 week
Probability of exit	0.0045	Probability of exit	0.0045
Annual discount factor	0.96	Annual discount factor	0.96
AR(1) growth rate of <i>M</i>	0.88	AR(1) growth rate of <i>M</i>	0.88
S.D. of shocks to growth rate of <i>M</i> , %	0.032	S.D. of shocks to growth rate of <i>M</i> , %	0.032
Elasticity of substitution	3	Elasticity substitution of type <i>B</i> agents	6

Table 5: Moments in economies using Dominick's weekly data

Statistic	Data	Model	
		Productivity shocks	Demand shocks
Frequency of all price changes, %	33	33	33
Frequency of regular price changes, %	2.9	2.9	2.9
Fraction of price changes that are temporary, %	94	95	95
Fraction of returns to regular price, %	80	88	88
Probability that temporary price spell ends, %	46	46	46
Fraction of periods with temp. prices, %	24	25	25
Fraction of periods when price temp. down, %	20	20	24
Fraction of prices at annual mode, %	58	55	56
Fraction of prices below annual mode, %	30	30	34
Mean size of price changes, %	17	17	17
Mean size of regular price changes, %	11	11	11
IQR of price changes, %	15	18	19
IQR of regular price changes, %	8	6	5
Fraction of output sold when temp. prices, %	39	36	39
Fraction of output sold when price temp. down, %	35	33	38

Table 6: Impulse responses to monetary shock using Dominick's data


Statistic	Productivity shocks		Demand shocks	
	Our model: With temporary changes	Standard model: Without temporary changes	Our model: With temporary changes	Standard model: Without temporary changes
Micro price stickiness, months	0.7	6.2	0.7	7.3
Aggregate price stickiness, %	18.4	18.4	22.5	22.5
Average output response, %	1.60	1.62	1.82	1.88
Half-life of output response, months	3.9	3.4	5.2	4.0

Note: The table reports the impulse response to a cumulative 5% increase in the money stock.

Aggregate price stickiness is measured as the average difference between *M* and *P* responses, relative to the *M* response.

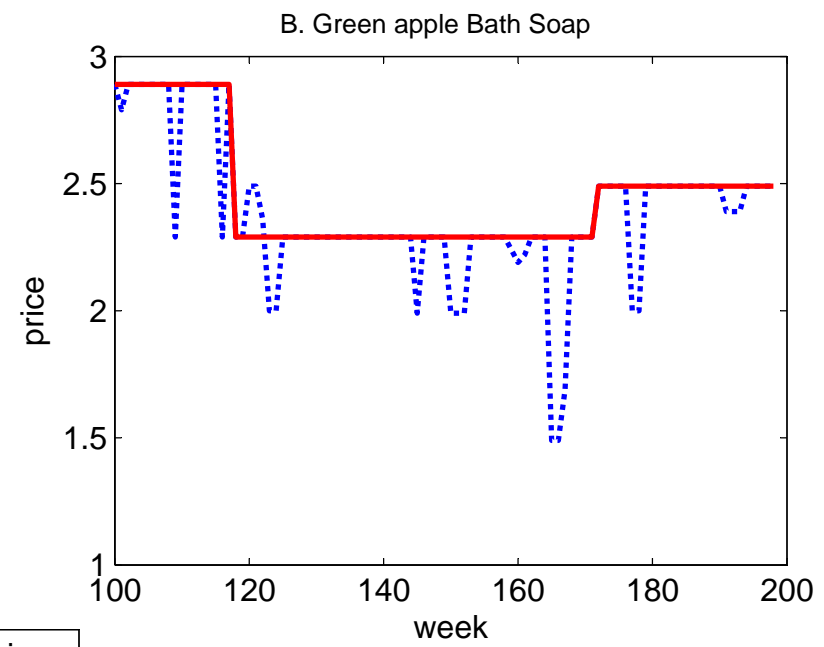

Responses are computed for the first 2 years after the shock.

Figure 1: Example of individual good price series
(VT Bath Tissue, 2 ply)



Source: Dominick's Finer Food Database

Figure 2: Examples of algorithm applied to price series in the Dominick's data



..... Actual price
—— Regular price





Figure 3: Firm's optimal decision rules in economy with temporary changes

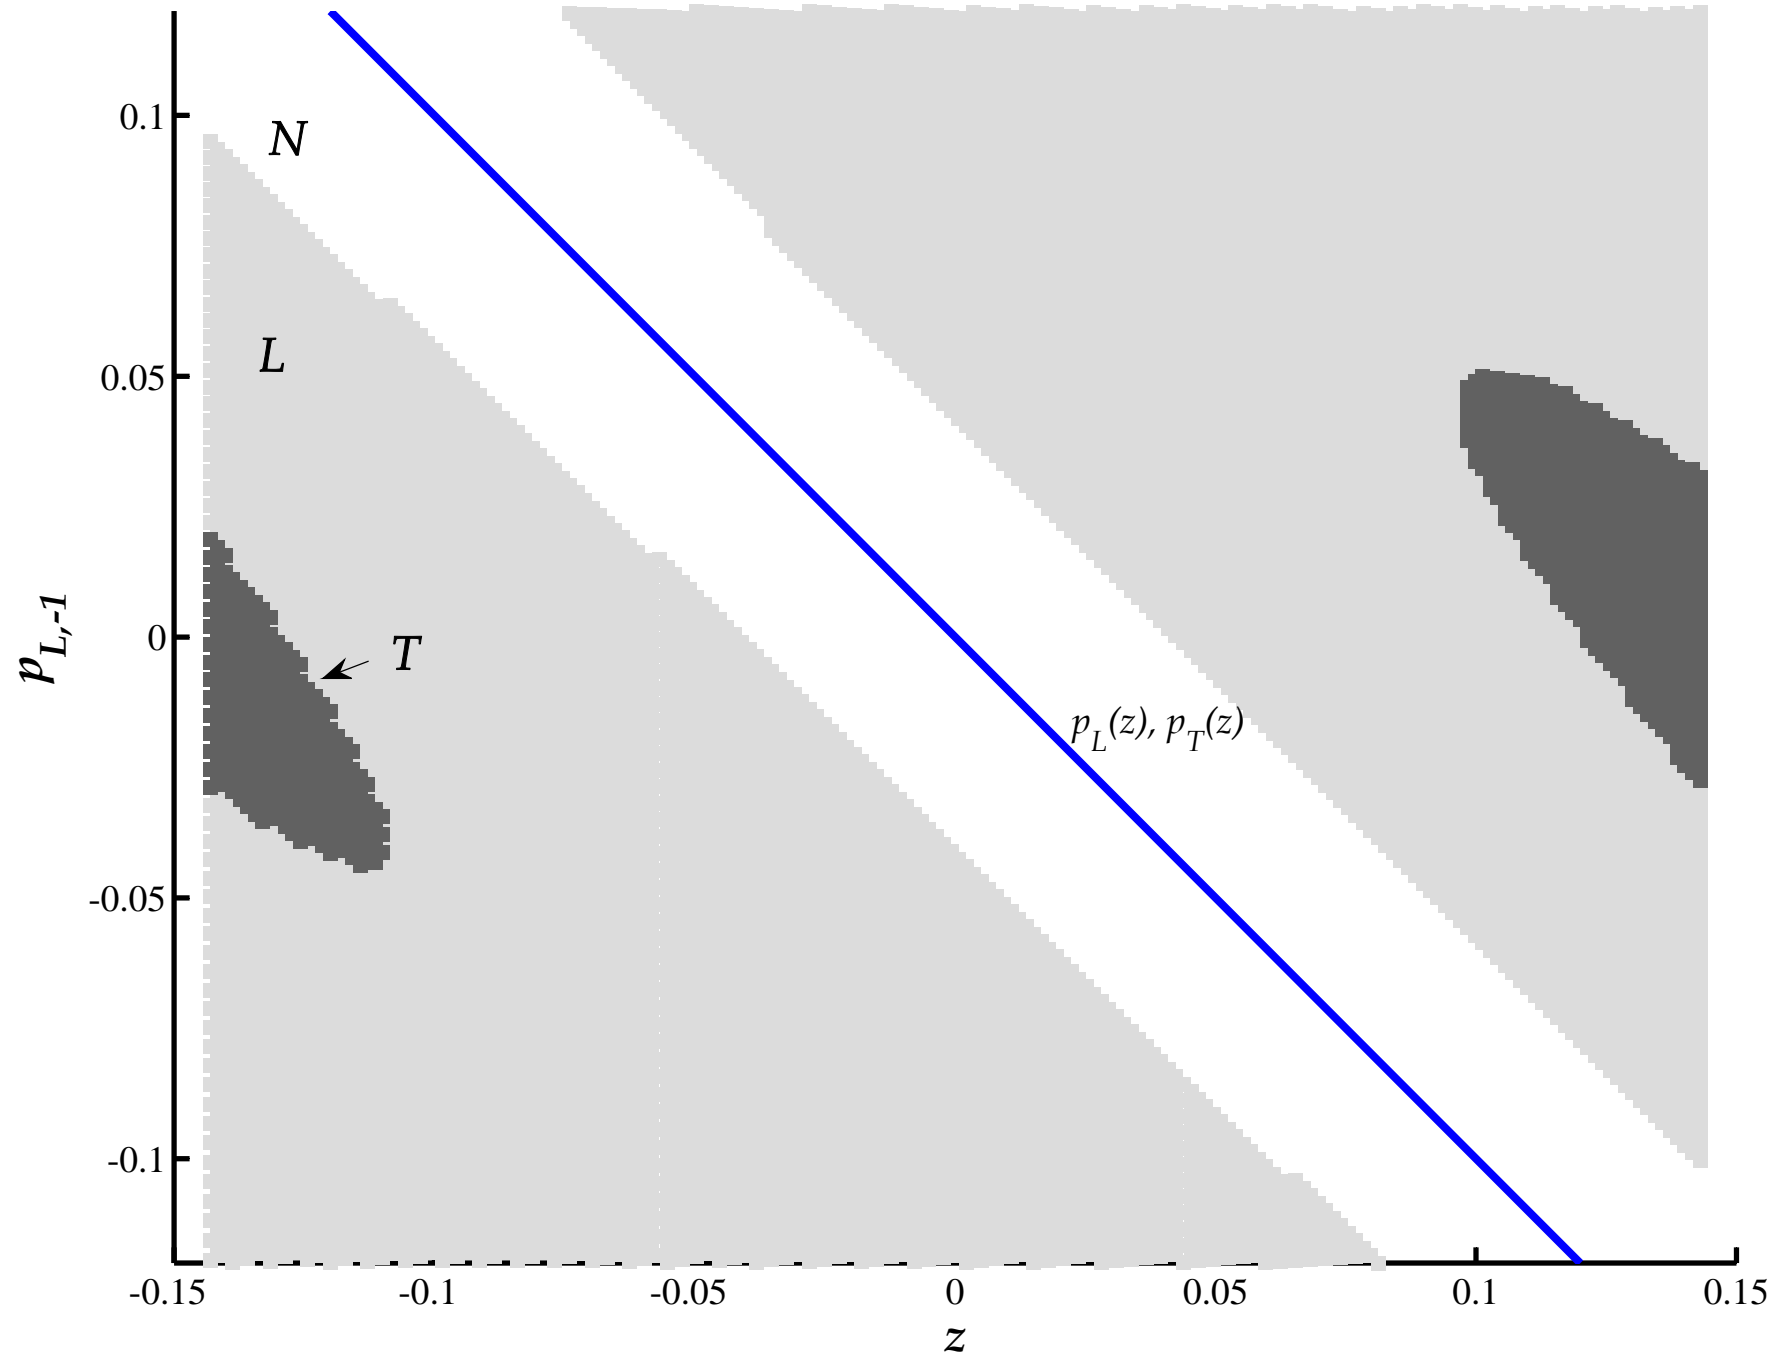


Figure 4: Simulation of prices and shocks

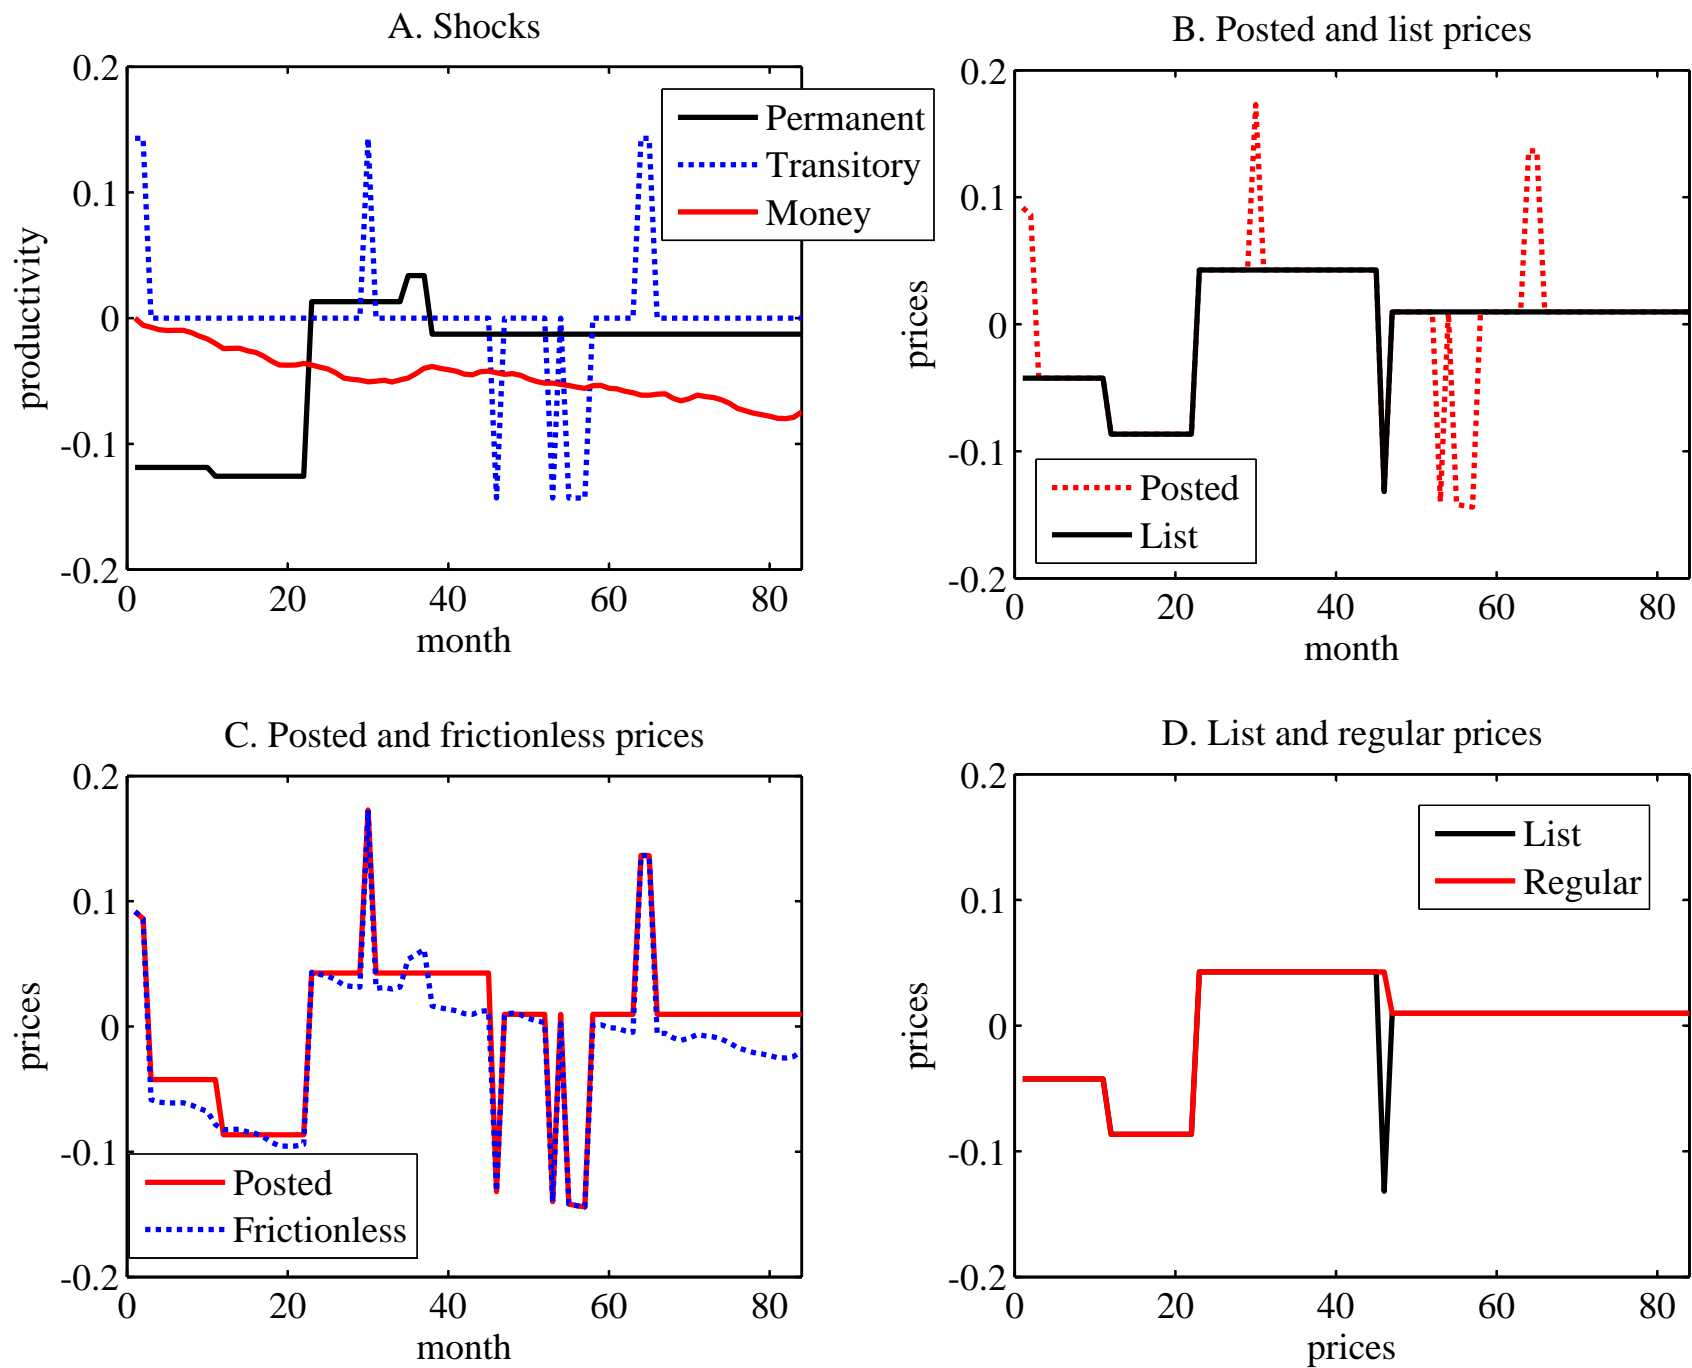


Figure 5: Relationship between high- and low-frequency stickiness in standard model.

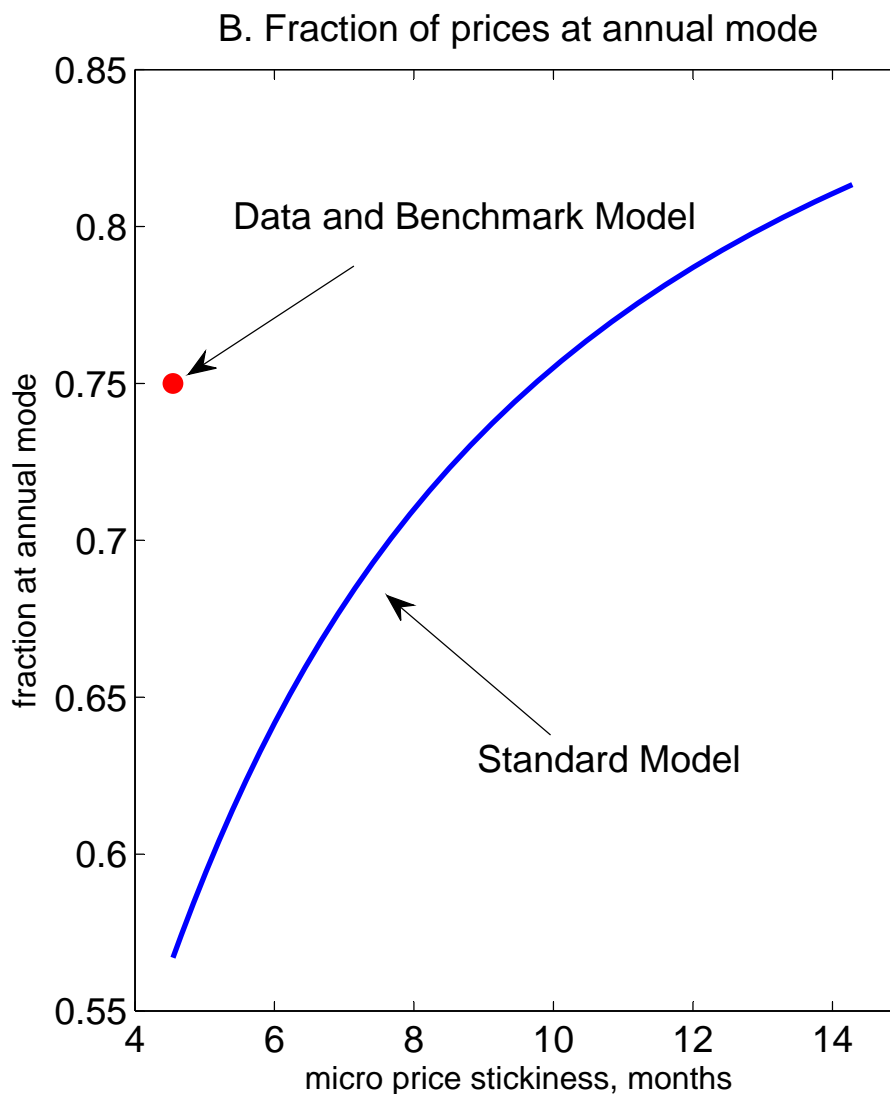
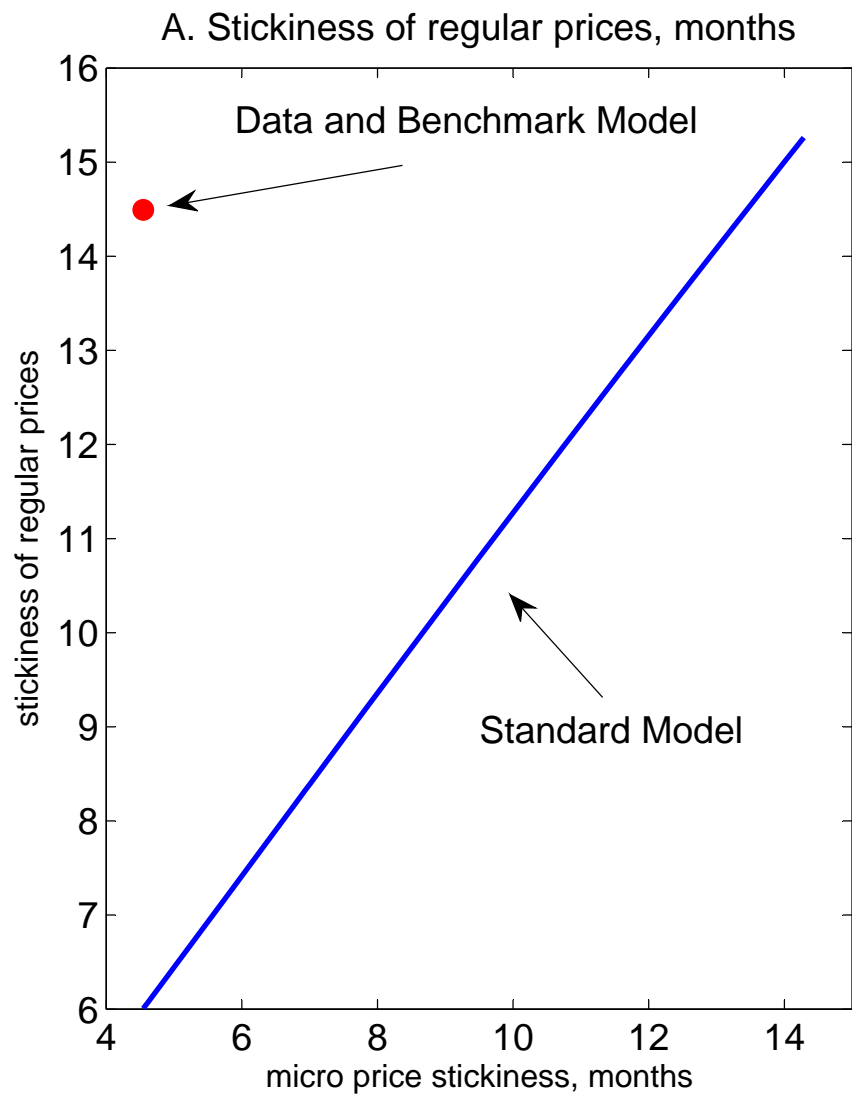


Figure 6: Impulse responses to a monetary shock of 5% in our model





Figure 7: Aggregate price stickiness vs. Micro price stickiness

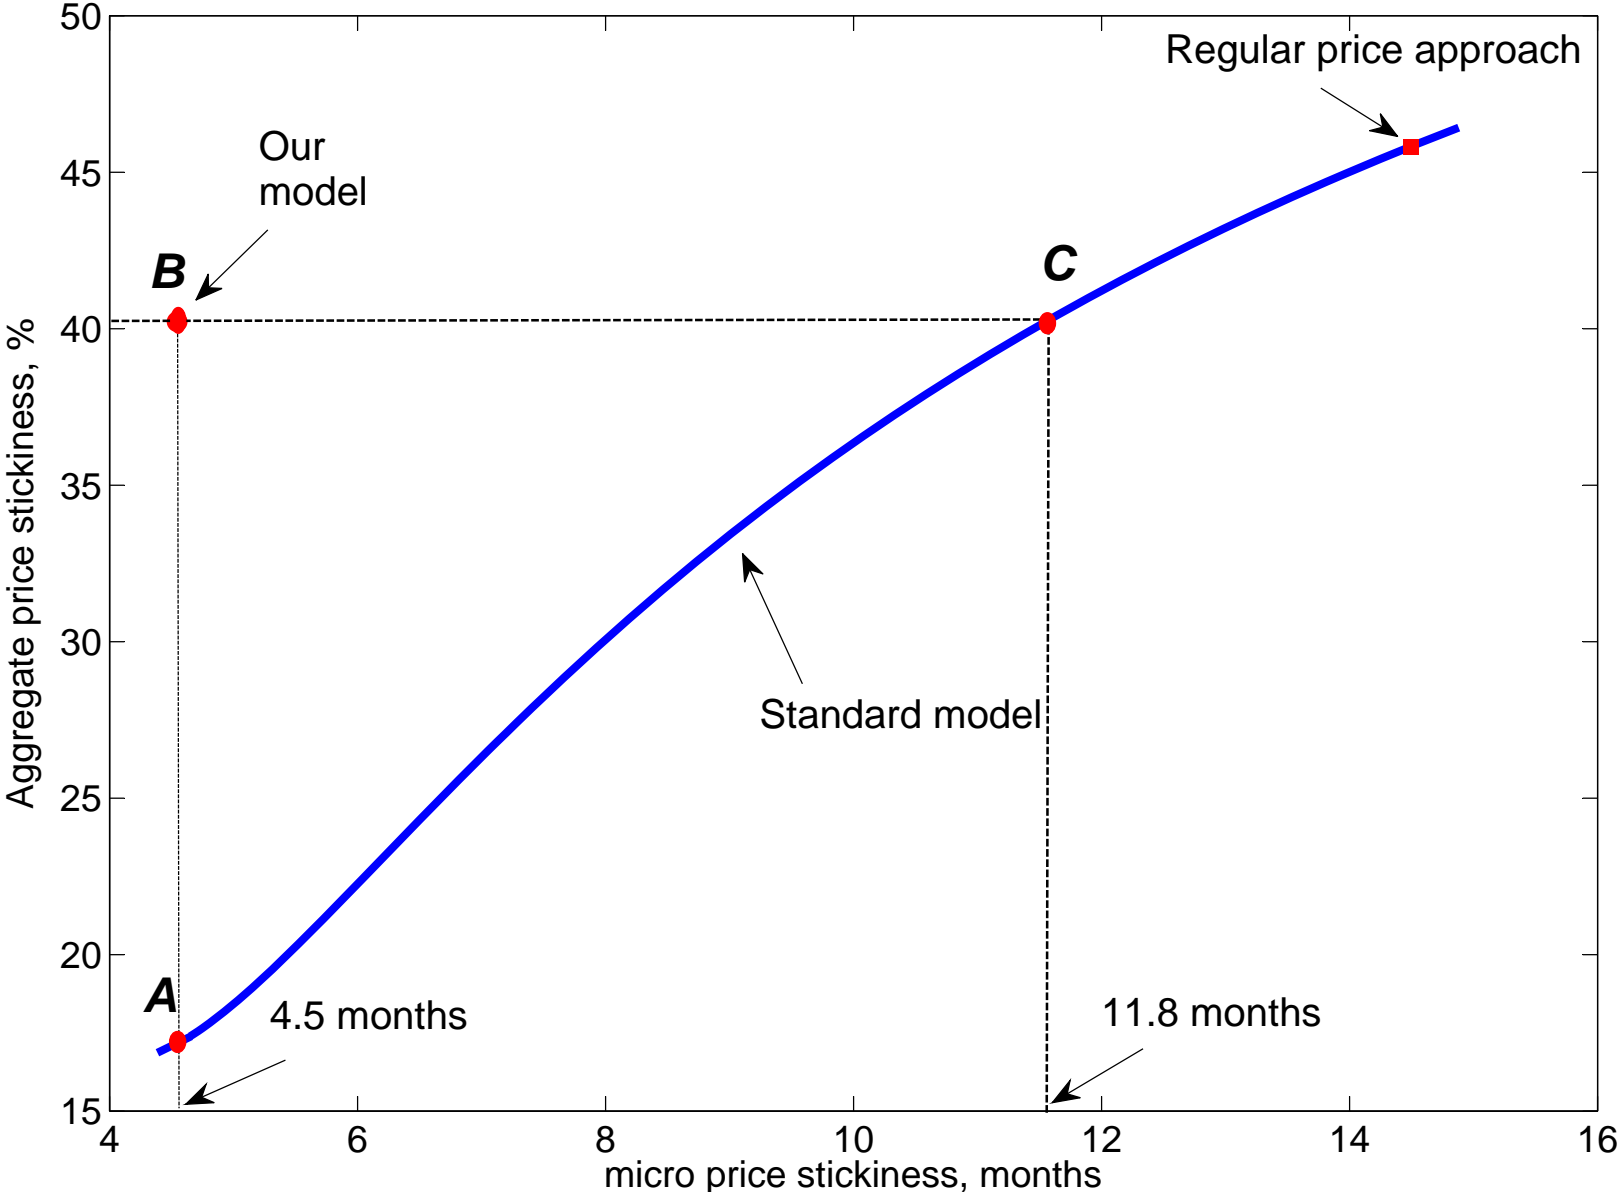


Figure 8: Impulse responses in our model and in the standard model with 11.8-month stickiness



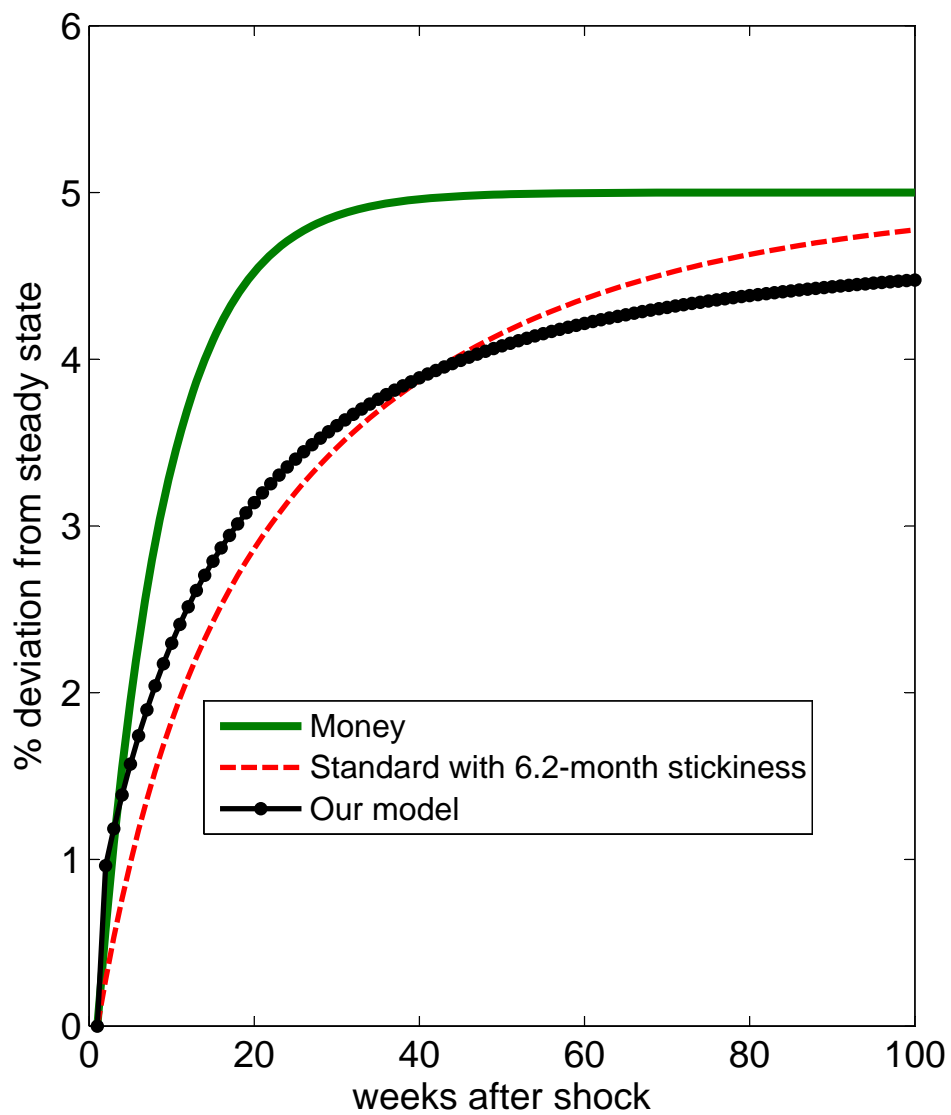


Figure 9: Impulse responses to a monetary shock in economy with productivity shocks using Dominick's data

A. Prices



B. Output

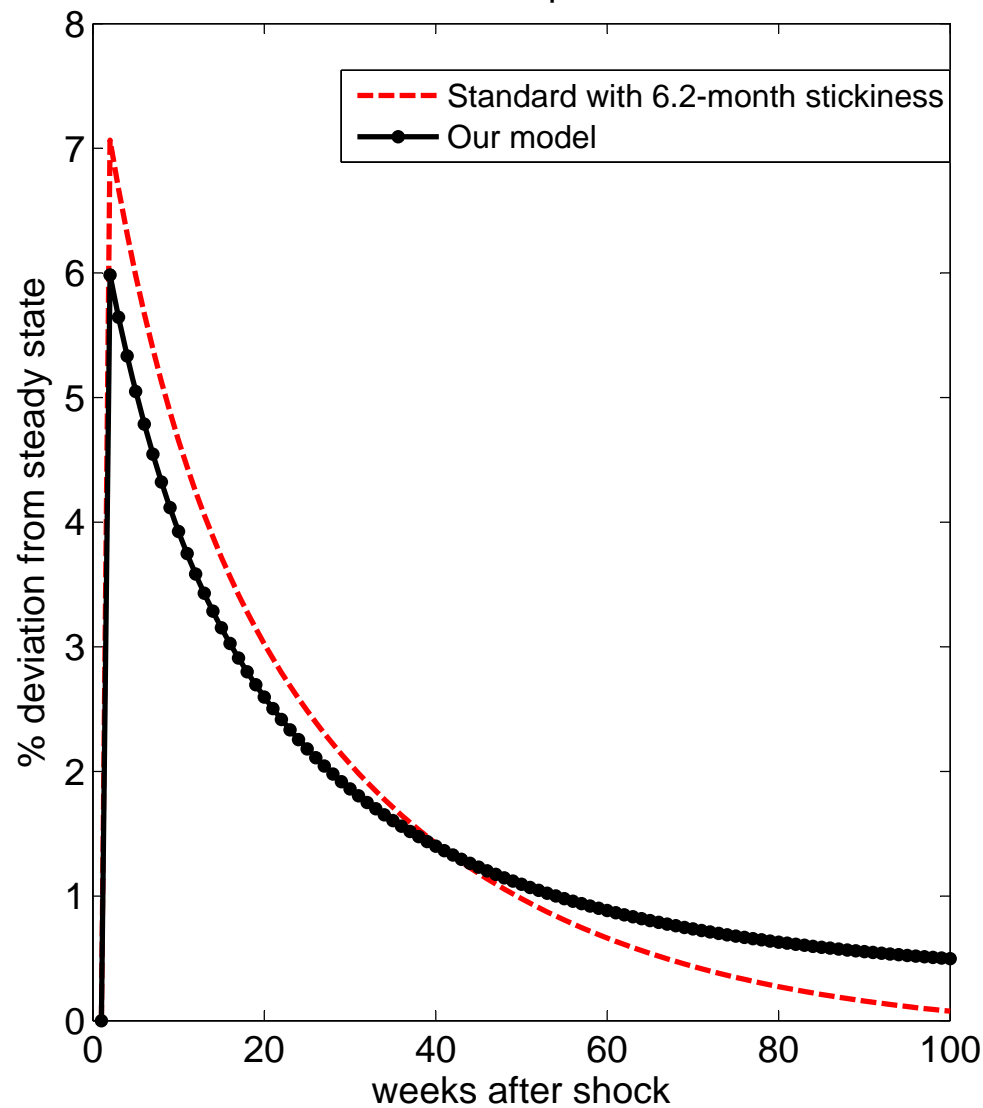
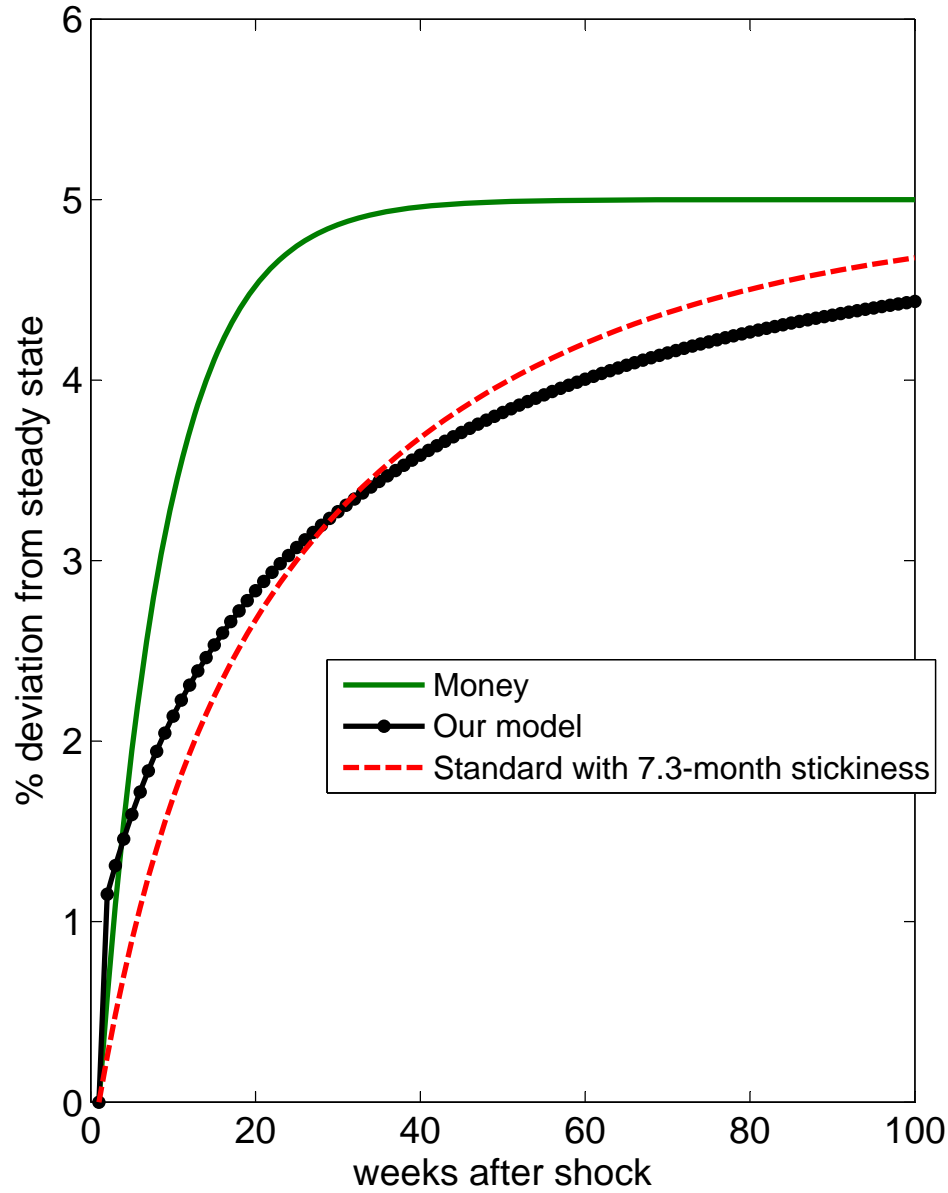


Figure 10: Impulse responses to a monetary shock in economy with demand shocks using Dominick's data

A. Prices



B. Output

