Federal Reserve Bank of Minneapolis
Research Department Staff Report 413

September 2008

Temporary Price Changes and the Real Effects of Monetary Policy*

Patrick J. Kehoe

Federal Reserve Bank of Minneapolis,
University of Minnesota,
and National Bureau of Economic Research

Virgiliu Midrigan

New York University

ABSTRACT

In the data, prices change both temporarily and permanently. Standard Calvo models focus on permanent price changes and take one of two shortcuts when confronted with the data: drop temporary changes from the data or leave them in and treat them as permanent. We provide a menu cost model that includes motives for both types of price changes. Since this model accounts for the main regularities of price changes, its predictions for the real effects of monetary policy shocks are useful benchmarks against which to judge existing shortcuts. We find that neither shortcut comes close to these benchmarks. For monetary policy analysis, researchers should use a menu cost model like ours or at least a third, theory-based shortcut: set the Calvo model's parameters so that it generates the same real effects from monetary shocks as does the benchmark menu cost model. Following either suggestion will improve monetary policy analysis.

*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 or the Federal Reserve System.

At the heart of monetary policy analysis is the question, How large are the real effects of monetary shocks? The most popular class of models used to attempt to answer this question assumes that goods prices are sticky, or that they change relatively infrequently. This assumption is a key determinant of the answers these models get. If prices change infrequently, then the models predict that the real effects of monetary shocks will be large. If prices actually change frequently, however, then the models predict the policy effects will be small. The measured stickiness of prices is thus critical for anyone using these models for monetary policy analysis.

Measuring the frequency of price changes in the data is not as straightforward as it may seem. How sticky prices actually are depends on whether the data being measured include temporary price changes. For in the data, only a small fraction of price changes are long-lasting, or permanent. A much larger fraction of price changes are quickly reversed; not long after a change, the price returns to its original level. When temporary price changes are included in a data set, therefore, prices naturally look fairly flexible, and without them, prices look quite sticky. This can be seen clearly in Figure 1, which displays a fairly typical price series for goods in our data set. When we include both types of price changes in the data (as in the dashed line), the good's price changes frequently; but when we include only permanent changes (as in the solid line), the price changes rarely.

Despite the critical nature of this distinction in the data, researchers generally make no attempt to model it, by explicitly building into their models motives for firms to make temporary price changes. Instead, when confronted with data in which a large fraction of price changes are temporary, researchers generally take one of two shortcuts. The most popular shortcut is to exclude temporary price changes from the data, construct a model without a motive for temporary price changes, and then choose parameters to match the frequency of price changes in the data with the temporary price changes excluded. We refer to this approach as the *temporary-changes-out* approach. An alternative shortcut, used less often, is to construct a model without a motive for temporary price changes and then choose parameters to match the frequency of price changes in the data with temporary price changes included. We refer to this approach as the *temporary-changes-in* approach.

Here we use theory to evaluate the adequacy of these two approaches for analyzing the real effects of monetary shocks. We find both approaches inadequate and offer two alternatives. Our theory is a simple menu cost model that explicitly includes a motive for temporary price changes and, hence, itself is an alternative to the common shortcuts. We document the regularities in the U.S. data concerning temporary price changes and then demonstrate that our model can account for them well. Because of that performance, we then use the model as a benchmark against which to judge the existing shortcuts used with standard Calvo (1983) sticky price models. We have the menu cost model predict the real effects of monetary shocks and compare its predictions to those of a standard model using each of the two common approaches. Neither approach performs well. We find that if we take the temporary changes out of the data, prices change infrequently, only every 50 weeks, and the Calvo model overestimates the real effects of monetary shocks by almost 70%. If we leave the temporary changes in the data, prices change much more often, every 3 weeks, and the Calvo model predicts only 11% of the real effects of monetary shocks as does our benchmark model.

Some researchers may find our first suggested alternative to their shortcuts—using a version of our benchmark menu cost model—computationally difficult. For them, we offer another alternative: use a simpler model that approximates the benchmark model's real effects. One way to do that is to set the frequency of price adjustment in the standard Calvo model so that it reproduces the real effects in the benchmark menu cost model. We demonstrate here that to do so, the Calvo model's parameters must be set so that, on average, prices change every 17 weeks. This second alternative is a theory-based shortcut that is preferable to the existing shortcuts.

Before we describe our benchmark menu cost model in detail, we attempt to describe its simple analytics, to help provide intuition for our results. We build the simplest possible model of temporary price changes that can be solved using pen and paper. The model is a Calvo sticky price model of price-setting modified to have temporary as well as permanent price changes. Since in this model the only aggregate shocks are shocks to the money supply, we measure the real effects of these shocks by the variance of output. We treat the model as the data-generating process and solve it in closed form for the law of motion for output and its variance. We then solve for similar closed-form expressions for output under the temporary-changes-out and -in approaches. A comparison of the expressions proves that the temporary-changes-out approach overstates the real effects and the temporary-changes-in approach understates them.

We then turn to our quantitative analysis. We start by documenting six regularities (or *facts*) about temporary and permanent price changes that we use to quantify the patterns of these changes. Among these regularities are that overall prices change frequently (every 3 weeks), 94% of price changes are temporary, about 90% of temporary price changes are cuts and about 10% are increases, temporary price changes revert to their preexisting level 80% of the time, and permanent price changes last for about a year. Unlike much of the previous research, we also document the comovements of quantities and prices. In particular, we find that periods in which a temporary price is charged account for a disproportionate amount of goods sold.

We then turn to our benchmark model, which is purposely chosen to be a parsimonious extension of the standard menu cost model of, say, Golosov and Lucas (2007) and Midrigan (2007). Indeed, we add to that model only one parameter on the technology of price adjustment. Nonetheless, our simple extension allows the model to produce patterns of both temporary and permanent price changes that are similar to those in the data.

In our model, firms are subject to two types of idiosyncratic disturbances: persistent productivity shocks and transitory shocks to 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.

To understand the technology for changing prices in our model, consider the problem of an individual firm. Such a firm enters each period with a preexisting *regular* price. This is the price the firm can charge in the current period at no extra cost. If the firm wants to charge a different price in the current period, then it has two options: change its current regular price to a new regular price, or change its price temporarily. To change its regular price, the firm must pay a fixed cost, or *menu cost*, which gives it the right to charge this price both today and in all future periods with no extra costs. We think of this option as akin to *buying* a permanent regular price change. To instead make a temporary price change, the firm must pay a smaller fixed cost, which gives it the right to charge a price that differs from its existing regular price for the current period only and keep its regular price unchanged. We think of this option as akin to *renting* rather than buying a price change. (Of course, the firm can rent a price change for several consecutive periods if it pays the rental cost each period.) We show that, essentially, the optimal choice in this environment is for firms to use a temporary price change to respond to a transitory demand shock. In contrast, their optimal choice when faced with much more persistent monetary and productivity shocks is to use a regular price change.

We show that our model can generate the salient features of the micro price data, including the frequency of temporary and permanent price changes, which we document here. We then use the model as a laboratory in which to study how well the two existing shortcuts approximate the real effects of monetary shocks. With our extended menu cost predictions as the benchmark, we demonstrate that the existing shortcuts that are meant to deal with temporary price changes are likely to be inadequate in applied settings. That should not be true of our alternative, theory-based approaches.

Our work is related to a recent debate in the literature between Golosov and Lucas (2007) and Midrigan (2007), focusing on how good an approximation a simple Calvo model of price changes is to a menu cost model. This literature assumes that the true data-generating process is a menu cost model and that researchers, for convenience, fit a Calvo model instead. Golosov and Lucas have found that the approximation is not good because the Calvo model greatly overstates the real effects of monetary shocks; Midrigan, however, argues that if a researcher matches more details of the micro data on prices, including the fat tails of the distribution of prices, such a researcher would overstate the real effects of monetary shocks only slightly. Our work extends this debate to environments with both temporary and permanent price changes. To match the details of the micro data on prices, we follow Midrigan (2007) and Gertler and Leahy (2008) in assuming fat-tailed shocks.

Our work is also related to a growing literature that documents features of micro price data in panel data sets. Two influential works in this literature are those of Bils and Klenow (2004) and Nakamura and Steinsson (forthcoming). When these researchers approach the data, they focus on temporary price declines, referred to as *sales*, rather than all temporary price changes, which are the sum of temporary price declines and temporary price increases.

These researchers have found, as we do, that the frequency of price changes measured in the data depends sensitively on the treatment of temporary price declines.

Our use of time-varying demand elasticities attempts to capture, in a very simple way, the spirit of an industrial organization literature that explains sales as arising from intertemporal price discrimination. (See, for example, the work of Varian (1980) and Sobel (1984), among others.) A critical distinction between our model and those in the earlier literature is that we have nominal frictions, menu costs for either temporarily or permanently deviating from an existing regular price. These frictions make it optimal for firms to return their price to the preexisting level after temporary price cuts. Without such menu costs, money would be neutral and the presence or absence of sales would be irrelevant for the real effects of monetary policy shocks.

The focus of our work is, however, quite different from that in the industrial organization literature. We want to understand how the presence of temporary price changes (including sales) alters a model's predictions about the size of the real effects of monetary policy shocks. The industrial organization literature aims to explain why temporary price changes (especially sales) ever arise at all. Because of our focus, we adopt a simple model of temporary price changes that is purposefully chosen to be similar to the existing sticky price models. We do not build an elaborate model of intertemporal price discrimination that has layered onto it nominal frictions that make money non-neutral. Building such an elaborate model is an interesting area for future research beyond the scope of this work.

Finally, Rotemberg (2004) offers an alternative explanation for why temporary prices return to their previous level. One view of Rotemberg's work is that it shows how costs to the firm of changing prices that act similarly to menu costs can arise from the preferences of consumers.

1. An Analytic Exercise

Before getting into the details of our quantitative model, we attempt to provide some of the intuition behind our argument that the common shortcuts to modeling temporary price changes are inadequate. We describe a simple Calvo model of price-setting and extend it to include both temporary and permanent price changes. We solve the extended Calvo model for a closed-form expression for the real effects of monetary shocks and then use it to analytically evaluate the two shortcuts. We find that both approaches are poor predictors of the real effects of monetary shocks. The temporary-changes-out approach overestimates them, and the temporary-changes-in approach underestimates them.

A. Extending the Simple Calvo Model

In our extended Calvo model, the only aggregate shock is to the money supply. Hence, aggregate real variables in this economy fluctuate only because money is not neutral. We measure the magnitude of the real effects of monetary shocks by the variance of aggregate consumption. We begin by briefly describing the economy and then solve for this variance as a function of the primitive parameters in the economy.

We borrow the formulation of the consumer problem from the menu cost model we will describe in detail later. That is a standard cash-in-advance model with a consumer who has the choice of a continuum of differentiated consumption goods. Here we describe just the key elements of the consumer problem that we need to illustrate our points.

In particular, the consumer's preferences in this Calvo model are defined over leisure and a continuum of consumption goods such that, given that the price of good i is P_{it} in any time period t, the consumer demand for each good i is

$$(1) \quad c_{it} = \left(\frac{P_{it}}{P_t}\right)^{-\theta} C_t,$$

where $C_t = \left(\int_0^1 c_{it}^{\frac{\theta-1}{\theta}} di\right)^{\frac{\theta}{\theta-1}}$ is the composite consumption good and θ the elasticity of substitution between goods. The corresponding aggregate price index is

$$(2) \quad P_t = \left(\int_0^1 P_{it}^{1-\theta} di\right)^{\frac{1}{1-\theta}}.$$

Moreover, the consumer utility function is such that the first-order condition for labor is

$$(3) \quad \frac{W_t}{P_t} = \psi C_t,$$

where W_t is the nominal wage and ψ is a parameter governing the disutility of working.

Finally, the cash-in-advance constraint binds:

(4) $P_t C_t = M_t,$

where the supply of money M_t is given by an exogenous stochastic process that follows $M_t = \mu_t M_{t-1}$. Here $\log \mu_t$ is independently and identically distributed (i.i.d.) with mean 0 and variance σ_μ^2.

The firm side of the model is more interesting. Each firm is the monopolistic supplier of a single good. Each firm enters a period with a preexisting regular price for its good, $P_{R,t-1}$. The firm must charge its existing regular price $P_{R,t-1}$ in the current period unless one of two events occurs. One event, referred to as a *permanent price change,* occurs with probability α_R and allows the firm to change this existing regular price to some new regular price P_{Rt}. The other event, referred to as a *temporary price change*, occurs with probability α_T and allows the firm to charge a price P_{Tt} that differs from its existing regular price $P_{R,t-1}$, but only for one period. That is, a firm that experiences a temporary price change will charge P_{Tt} in the current period and $P_{R,t-1}$ in the subsequent period unless in the subsequent period that firm again experiences one of the two price-changing events. (Note that this feature of the model is consistent with the observation that most temporary price changes revert to the preexisting regular price.)

Consider the problem facing a firm that is allowed a temporary price change P_{Tt} in period t. Clearly, the choice of this price has no influence on the firm's profits in any future period or state. Thus, the firm simply solves the static problem of maximizing current profits,

$$(P_{T,t} - W_t)\left(\frac{P_{T,t}}{P_t}\right)^{-\theta} C_t.$$

Here the optimal price is $P_{T,t} = \theta W_t/(\theta - 1)$. Note from (3) and (4) that in equilibrium $W_t = \psi M_t$. For convenience, normalize all nominal variables by the money supply. Doing so and then log-linearizing gives that

(5) $p_{T,t} = 0,$

where $p_{T,t}$ is the log deviation of $P_{T,t}/M_t$ from its steady state.

Consider next the problem facing a firm that is allowed a permanent price change. That is, in period t the firm can reset its regular price $P_{R,t}$. Clearly, in choosing its new price, that firm need consider only the states for which that price will be in effect. (This price has no effect on either future periods in which the firm can choose a temporary price or future periods in which a new regular price will be in effect.)

The firm will want to maximize the value of profits during those periods and states in which the price chosen today will be in effect. Letting $\lambda = 1 - \alpha_R$ be the probability that the firm doesn't make a permanent change, the objective is to maximize this expression:

$$(P_{R,t} - W_t)\left(\frac{P_{R,t}}{P_t}\right)^{-\theta} C_t + E_t \sum_{s=t}^{\infty} \lambda^{s-t} \frac{1-\alpha_T-\alpha_R}{\lambda} \tilde{Q}_{t,s}\left[(P_{R,t} - W_s)\left(\frac{P_{R,t}}{P_s}\right)^{-\theta} C_s\right],$$

where $Q_{t,s}$ is the price of a dollar in period s in units of dollars in t, normalized by the conditional probability of the state in s given the state in t. To understand this objective, note that in t the prevailing price is $P_{R,t}$, in $t+1$ the prevailing price is $P_{R,t}$ with probability $1-\alpha_T-\alpha_R$, in $t+2$ the prevailing price is $P_{R,t}$ with probability $\lambda(1-\alpha_T-\alpha_R)$, and so on. Letting $p_{R,t}$ denote the log deviation of $P_{R,t}/M_t$ from its steady state, we can easily see that the log-linearized first-order condition for this problem is

$$p_{R,t}\left[1 + \sum_{j=1}^{\infty} (\lambda\beta)^j \frac{1-\alpha_T-\alpha_R}{\lambda}\right] = w_t + E_t \sum_{j=1}^{\infty} (\lambda\beta)^j \frac{1-\alpha_T-\alpha_R}{\lambda} w_{t+j}.$$

As we have already noted, (3) and (4) imply that in equilibrium $W_t = \psi M_t$. Letting w_t denote the log deviation of W_t/M_t from its steady state, we have that $w_s = 0$ for all s, so that

$$(6) \quad p_{R,t} = 0.$$

The intuition for (6) is simple. The firm chooses its new regular price as a markup over the discounted value of its expected future marginal costs—here, future nominal wages. Since wages are proportional to the nominal money supply and the money supply is a random walk, the mean of future wages is equal to current wages and, hence, proportional to the current money supply. Hence, the firm sets its new price proportional to the current money supply, which in normalized log-deviation terms means the firm sets it equal to zero.

Now we describe how aggregate consumption in this Calvo economy evolves.

Proposition 1. Aggregate consumption in log-linearized form for this economy evolves according to

$$c_t = (1 - \alpha_R)c_{t-1} + (1 - \alpha_R - \alpha_T)\mu_t. \tag{7}$$

Proof. We establish Proposition 1 using the cash-in-advance constraint (4). Log-linearizing this constraint gives that

$$c_t = -p_t. \tag{8}$$

Thus, to get an expression for the evolution of aggregate consumption, we need only solve for the law of motion of the price index. From (2) we know that this index is given by

$$p_t = \int_0^1 p_{it}\, di. \tag{9}$$

To compute the right side of (9), we note that the fraction α_R of firms in t charge $p_{Rt} = 0$, the fraction α_T of firms in t charge $p_{Tt} = 0$, and the rest charge whatever is their existing regular price. Let $\bar{p}_{R,t-1}$ denote the average of existing regular prices in $t-1$ normalized by the money supply in $t-1$ and expressed in log-deviation form. Then we can write the price index as

$$p_t = \alpha_R p_{R,t} + \alpha_T p_{T,t} + (1 - \alpha_R - \alpha_T)(\bar{p}_{R,t-1} - \mu_t). \tag{10}$$

To prove the proposition, we must also describe the law of motion for the average existing regular price $\bar{p}_{R,t}$. Given that α_R firms reset prices in t to $p_{R,t}$ and that $1 - \alpha_R$ do not, but instead use whatever their regular price was in $t-1$, we can write the law of motion for $\bar{p}_{R,t}$ recursively as

(11) $\bar{p}_{Rt} = \alpha p_{R,t} + (1 - \alpha_R)(\bar{p}_{R,t-1} - \mu_t),$

where, from (6), we know that $p_{Rt} = 0$. Combining (10) and (11) gives that

$$p_t = (1 - \alpha_R)p_{t-1} - (1 - \alpha_R - \alpha_T)\mu_t.$$

Substituting from (8) gives our result (7). *Q.E.D.*

B. Evaluating the Two Common Shortcuts

Now we use this extended Calvo model to evaluate the two common shortcuts to dealing with temporary price changes. Consider a researcher who studies the data generated by our extended Calvo model, with both temporary and permanent price changes, through the lens of a simple standard Calvo model, with only permanent price changes and with a frequency of price change α. The researcher using such a simple model follows one of the two common approaches we have described to calibrate the frequency of price changes in this model. In the *temporary-changes-out* approach, we imagine that the researcher is able to isolate the permanent price changes and thus concludes that the frequency of price changes is $\alpha = \alpha_R$. In the *temporary-changes-in* approach, we imagine that the researcher uses the raw data that include the temporary price changes, concluding that the frequency of price changes is approximately $\alpha = \alpha_R + 2\alpha_T$. To see where this last expression comes from, recall that every temporary price change involves two price changes, one *to* and one *from* the temporary price.

To set up evaluation of the two approaches, note that our derivation above implies that the standard Calvo pricing, in which a fraction α of firms reset prices in any given period, has a law of motion for consumption of

$$c_t = (1 - \alpha)c_{t-1} + (1 - \alpha)\mu_t,$$

and the unconditional variance of c_t is, therefore,

$$(12)\quad \mathrm{var}(c_t) = \frac{(1-\alpha)^2}{1-(1-\alpha)^2}\sigma_\mu^2.$$

Now set $\alpha = \alpha_R$ and $\alpha = \alpha_R + 2\alpha_T$ in (12). Let c_t^{Out} and c_t^{In} denote the stochastic processes for consumption generated under the two approaches, and let c_t denote the stochastic process for the data-generating process. Then we can say this:

Proposition 2. The temporary-changes-out approach overstates the real effects of monetary shocks, whereas the temporary-changes-in approach understates those effects. In particular, the temporary-changes-out approach predicts that

$$\mathrm{var}(c_t^{\mathrm{Out}}) > \mathrm{var}(c_t) > \mathrm{var}(c_t^{\mathrm{In}}).$$

Proof. Evaluating (12) at $\alpha = \alpha_R$ and $\alpha = \alpha_R + 2\alpha_T$ gives that

$$(13)\quad \frac{(1-\alpha_R)^2}{1-(1-\alpha_R)^2}\sigma_\mu^2 > \frac{(1-\alpha_R-\alpha_T)^2}{1-(1-\alpha_R)^2}\sigma_\mu^2 > \frac{(1-\alpha_R-2\alpha_T)^2}{1-(1-\alpha_R-2\alpha_T)^2}\sigma_\mu^2.$$

Clearly, the left-most term in (13) is (12) evaluated at $\alpha = \alpha_R$, the middle term follows from (7), and the right-most term is (12) evaluated at $\alpha = \alpha_R + 2\alpha_T$. *Q.E.D.*

The intuition for this result is as follows. The temporary-changes-out approach correctly predicts the persistence of consumption, but it overstates the volatility of shocks to the consumption process because it ignores the fact that a fraction α_T of firms change prices in any given period and thus offset the monetary shock. In contrast, the temporary-changes-in approach understates the persistence of consumption because it fails to recognize that some of the prices change only temporarily and revert to their previous value. Moreover, that approach counts the returns from the temporary price to the permanent price as a price change that is useful in responding to the monetary shock, but in fact it is not, since the price returns to a preexisting level.

This simple setup can also be used to answer a question that can help researchers improve their models' predictions: To what frequency of price changes should a researcher calibrate a simple Calvo model with no temporary price changes in order to predict the real effects of monetary shocks in the model with a fraction α_R of permanent price changes and a fraction α_T of temporary price changes? Using the results above, we know that the frequency of price changes, α, equates to

$$\frac{(1-\alpha)^2}{1-(1-\alpha)^2} = \frac{(1-\alpha_R-\alpha_T)^2}{1-(1-\alpha_R)^2}.$$

We thus have the following corollary to Proposition 2:

Corollary: If the data are generated by our extended Calvo model, which includes both permanent and temporary price change parameters, α_R and α_T, then a simple Calvo model with parameter

$$(14) \quad 1-\alpha = \frac{1-\alpha_R-\alpha_T}{[1-(1-\alpha_R)^2+(1-\alpha_R-\alpha_T)^2]^{\frac{1}{2}}}$$

will predict the same real effects of monetary shocks.

2. Price Changes in the Data: Six Facts

We turn now to documenting how prices have changed in the U.S. data. We here describe six regularities, or *facts*, that we see in these data. We will later use the data to both calibrate and evaluate our model.

Our data base is a by-product of a randomized pricing experiment conducted by the Dominick's Finer Foods retail chain in cooperation with the University of Chicago Graduate School of Business (the James M. Kilts Center for Marketing). The data base 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.[1] 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).)

We use a simple algorithm (described in Appendix A) to categorize all price changes in this data base as either *temporary* or *permanent*. To do so, we define for each product an artificial series called a *regular* price series, denoted $\{P_t^R\}$, which we construct and use mainly to define which periods are periods of temporary price changes. An intuitive way to think about our analysis is to imagine that at any point in time every product has an existing regular price and may experience two types of price changes: *temporary* changes in which the price briefly moves away from the regular price and *permanent* changes which are changes in the regular price itself.

In Figure 2, we illustrate the results of our algorithm for several particular price series. On each of the four graphs, for each of the four products, the dashed lines are the raw data (the original prices), and the solid lines are the regular price series constructed with our algorithm. On each graph, every price change that is a deviation from the regular price line is defined as a *temporary* price change, whereas every price change that coincides with a change in the regular price is defined as a *permanent* price change. Perusal of these graphs makes some facts about price changes clear: across the board, price changes are frequent and large, most of them are temporary, and most temporary prices return to the preexisting regular price.

We turn now to a more formal description of the data that we will use in our theoretical model. In Table 1 and Figure 3, we report a variety of general facts about price changes that our data reveal. (All statistics are computed by weighting each good by its sales share.)

FACT 1: *Prices change frequently, but most price changes are temporary, and after temporary changes, prices tend to return to the regular price.* Notice from line 1 in Table 1 that the frequency of weekly price changes in these data is 33%, so prices change on average once every three weeks. However, most of these price changes—indeed, 94% of them—are temporary (line 2). Regular prices, therefore, change infrequently, with a weekly frequency rate of 2%, or about once a year. The temporary price changes are short-lived; on average they last two weeks, so the probability that a temporary price change reverts to the preexisting regular price is 46% (line 4). Moreover, 80% of the time (line 3), temporary price changes return to the preexisting regular price.

Fact 2: *Most temporary price changes are cuts, not increases.* Of all the periods in the data when the store charges a temporary price (24.3%, line 5), most of the time the price moves temporarily down (20.3%, line 6) rather than up (2.1%, line 7).

Fact 3: *During a year, prices stay at their annual modal value most of the time. When prices are not at their annual mode, they are much more likely to be below it than above it.* Table 1 shows that, on average during a 50-week period, prices tend to be at their annual modal value 58% of the time (line 8). When prices are not at their annual mode, they are most likely below it (30%, line 9). That leaves prices above their mode only 12% of the time. Thus, prices are about 2.5 times as likely to be below the annual mode than above it.

Fact 4: *Price changes are large and dispersed.* The mean size of all price changes in these data is 17% (line 10), and their interquartile range is 15% (line 13). The mean of regular price changes is 11%. Also large and dispersed are temporary price changes. The mean deviation of the temporary price from the regular price is −22% (line 11) when the price is temporarily down and 13% (line 12) when it is temporarily up. The interquartile ranges of these temporarily down and up deviations are 21% (line 14) and 12% (line 15), respectively.

Fact 5: *Periods of temporary price cuts account for a disproportionately large share of goods sold. Quantities sold are more sensitive to prices when prices decline temporarily than when they decline permanently.* In the data, 38% of output is sold in periods with temporary prices (line 16), 35.4% when the price is temporarily down (line 17), and 1.2% when the price is temporarily up (line 18), even though the fraction of weeks with temporary prices is relatively small: 24.3%, 20.3%, and 2.1%, respectively. (See Fact 2.) Put another way, in periods of temporary price declines, more than twice as many goods are sold as in periods of regular prices. A regression of changes in quantities on changes in prices during regular price change periods yields a slope coefficient of −2.08 (line 19). A similar regression during periods when the price change is a temporary decline from a regular price yields a slope coefficient of −2.93 (line 20). (Of course, the slope coefficient in our simple regression is not a true structural measure of demand elasticity. Nonetheless, note that in static monopolistic competition, setting an increase in demand elasticity from 2.08 to 2.93 would lower the monopolist's markup from 92% to 52%. In this metric, the change in the slope coefficient is large.)

FACT 6: *Price changes are clustered in time.* In Figure 3 we display the *hazard of price changes*, defined as the probability that prices change in period $t + k$ when the last price change occurred in period t. We computed this hazard by assuming a log-log functional form for the hazard of price adjustment and estimating the resulting model by allowing for good-specific random effects and holiday and seasonal dummies, as well as by modeling age dependence nonparametrically. In constructing the likelihood function, we weight each product according to its share in Dominick's total revenue.

Figure 3 shows the effect of varying the age of the *price spell*, or how long the new price lasts, while holding all other covariates constant at their mean.[2] Note that this procedure implicitly accounts for ex ante heterogeneity in the frequency of price changes across products by use of good-specific random effects.

The left panel of the figure displays the hazard for all price changes, both temporary and permanent. The panel shows that the hazard for a price change at one week after a change is 38%. That is, if a store has changed the price of a given product last week, then the store changes that price again this week 38% of the time. More generally, we see that the hazard sharply declines in the first two weeks after a price change and follows a declining trend thereafter. This implies that price changes tend to come in clusters: overall, the data include periods with many price changes followed by prolonged periods with none.

The right panel of the figure displays the hazard for just regular price changes. Without temporary price changes included, the hazard is low and flat, though slightly increasing in the first few weeks.

3. A Model of Temporary and Permanent Price Changes

Now we attempt to build a model that can reasonably well approximate the facts about price changes that we have just documented.

Our model explicitly allows temporary as well as permanent price changes, yet is a parsimonious extension of a standard menu cost model. Indeed, our model includes only one parameter on the firm side that is not part of that standard model. Here, as there, firms can pay a fixed cost and change their regular price. Our simple innovation is to allow firms the option in any period of paying a different and smaller fixed cost in order to change their price temporarily, for only one period, leaving their regular price unchanged. Our one new parameter is the size of the fixed cost for a temporary price change. At an intuitive level, we think of the standard model as requiring that the only way a price can change is that the firm buys a potentially permanent price change. In this way, we think of our model as adding an option of renting a price change for one period.

The standard menu cost model of Golosov and Lucas (2007) has only technology shocks, but we allow both technology shocks and demand shocks. Our motivation is from both theory and data.

Our theoretical motivation 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 the product happen to have high elasticities.

Our empirical motivation comes from two observations. First, as we have shown, quantities sold seem to be more sensitive to price changes during periods of temporary price declines than during other periods (Fact 5). Second, as several researchers have shown, temporary price cuts are associated with reductions in price-cost margins. (See, for example, the work of Chevalier, Kashyap, and Rossi (2003).) 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 argue that our extended menu cost model is a useful laboratory for evaluating the common approaches to treating temporary price changes in the data. We do this by showing that the model can fit what we think are the key aspects of the micro data.

A. Setup

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

In the model, we have aggregate shocks to money supply and idiosyncratic shocks to a firm's productivity and the demand for each good. In terms of the money supply shocks, we assume that the (log of) money growth follows an autoregressive process of the form

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

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.

Technology and Consumers

In each period t, the commodities in this economy are labor, money, and a continuum of consumption goods indexed by $i \in [0, 1]$. Good i is produced using the technology

$$y_i(s^t) = a_i(s^t) l_i(s^t),$$

where $y_i(s^t)$ is the output of good i, $l_i(s^t)$ the labor input to the production process, and $a_i(s^t)$ the good-specific productivity shock that evolves according to

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

where ρ_a is the persistence of the productivity process and $\varepsilon_i(s^t)$ the persistent shock to productivity.

The economy has two types of consumers, differentiated by how much their demand responds to price changes: measure $1 - \omega$ of *low elasticity* consumers and measure ω of *high* *elasticity consumers.* The stand-in consumer for the low elasticity consumers, a consumer of type A, has preferences of the form

$$(17)\quad \sum \beta^t \pi(s^t)[\log c_A(s^t) - \psi l_A(s^t)],$$

where β is the discount factor, $c_A(s^t)$ is a composite of goods consumed given by $\left(\int_0^1 c_{Ai}(s^t)^{\frac{\theta-1}{\theta}} di\right)^{\frac{\theta}{\theta-1}}$, $l_A(s^t)$ is labor supplied by this consumer, and ψ is a parameter governing the disutility of work. The stand-in consumer for the high elasticity consumers, a consumer of type B, has preferences of the form

$$(18)\quad \sum \beta^t \pi(s^t)[\log c_B(s^t) - \psi l_B(s^t)],$$

where $c_B(s^t)$ is a composite of goods given by $\left(\int_0^1 z_i(s^t)^{\frac{1}{\gamma}} c_{Bi}(s^t)^{\frac{\gamma-1}{\gamma}} di\right)^{\frac{\gamma}{\gamma-1}}$, $l_B(s^t)$ is labor supplied by this consumer, and $z_i(s^t)$ is a type of preference shock for individual goods or, more simply, *demand shocks.* Note that all high elasticity consumers receive the same realization of the demand shock for a specific good. In this way, variations in this shock represent demand variation at the level of each good but induce no aggregate uncertainty because there is a continuum of goods. Note also that on the labor side we follow Hansen (1985) by assuming that indivisible labor decisions are implemented with lotteries.

In this economy, the markets for state-contingent money claims are complete. We represent the asset structure by having 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 $t+1$. One unit of this bond pays one unit of money in period $t+1$ if the particular state s^{t+1} occurs and 0 otherwise. Let $Q(s^{t+1}|s^t)$ denote the price of this bond in period t and state s^t. Clearly, $Q(s^{t+1}|s^t) = Q(s^{t+1})/Q(s^t)$.

Consider the constraints facing the consumer of type A (with low elasticity). The purchases of goods by this consumer must satisfy the following cash-in-advance constraint:

$$\int P_i(s^t) c_{Ai}(s^t)\; di \leq M(s^t),$$

where $p_i(s^t)$ is the price of good i and $M(s^t)$ is nominal money balances. The budget constraint of this consumer is

$$\text{(19)}\quad M(s^t) + \sum_{s_{t+1}} Q\left(s^{t+1}|s^t\right) B(s^{t+1})$$

$$= R(s^{t-1})W(s^{t-1})l_A(s^{t-1}) + B(s^t) + \left[M(s^{t-1}) - \int P_i(s^t)c_{Ai}(s^t)\ di\right] + T(s^t) + \Pi(s^t),$$

where $1/R(s^t) = \sum_{s_{t+1}} Q\left(s^{t+1}|s^t\right)$ is the uncontingent nominal interest rates, $W(s^t)$ is the nominal wage rate, $l(s^t)$ is labor supplied, $T\left(s^t\right)$ is transfers of money, and $\Pi(s^t)$ are profits. The left side of (19) is the nominal value of assets held at the end of bond market trading. The terms on the right side are the returns to last period's labor market activity, the value of nominal debt bought in the preceding period, the consumer's unspent money, the transfers of money, and the profits from the firms. The cash-in-advance constraint and the budget constraint for a consumer of type B (with high elasticity) are analogous.

Notice that in (19) we are assuming that firms pay consumers $W(s^{t-1})l_A(s^{t-1})$ at the end of period $t-1$ and that the government transfers to consumers $[R(s^{t-1})-1]W(s^{t-1})l_A(s^{t-1})$ and pays for those transfers with lump-sum taxes implicit in $T(s^t)$. Having the government make such transfers is a simple device that eliminates the standard distortion in the labor-leisure decision that arises in cash-in-advance models because consumers get paid in cash at the end of one period and must save that cash at zero interest until the next period. These distortions are not present in the recent literature on sticky prices, so we abstract from them here in order to retain comparability.

Solving the consumers' problem in two stages is convenient. In the first stage, we solve for the optimal choice of expenditure on each variety of good, given the composite demands. Consider, again, a consumer of type A. For composite demand $c_A(s^t)$, we solve

$$\min \int_0^1 P_i(s^t)c_A(s^t)\ di$$

subject to $c_A(s^t) = \left(\int_0^1 c_{Ai}(s^t)^{\frac{\theta-1}{\theta}} di\right)^{\frac{\theta}{\theta-1}}$, and we define the resulting price index as $P_A(s^t) = \left(\int_0^1 p_i^{1-\theta}(s^t)\ di\right)^{\frac{1}{1-\theta}}$. We solve an analogous problem for the composite demand of a consumer of type B, $c_B(s^t) = \left(\int_0^1 z_i(s^t)^{\frac{1}{\gamma}} c_{Bi}(s^t)^{\frac{\gamma-1}{\gamma}} di\right)^{\frac{\gamma}{\gamma-1}}$, and define the resulting price index as

$P_B(s^t) = \left(\int_0^1 z_i(s^t) p_i^{1-\gamma}(s^t)\, di\right)^{\frac{1}{1-\gamma}}$. The resulting total demand for good i is thus given by

$$(20)\quad q_i(s^t) = (1-\omega)\left(\frac{P_i(s^t)}{P_A(s^t)}\right)^{-\theta} c_A(s^t) + \omega \left(\frac{P_i(s^t)}{P_B(s^t)}\right)^{-\gamma} z_i(s^t) c_B(s^t).$$

Notice that (20) makes clear the precise sense in which the shocks $z_i(s^t)$ represent a type of demand shock: if $z_i(s^t)$ is relatively high, then at a given set of prices and composite demands $c_A(s^t)$ and $c_B(s^t)$, the total demand for good i is relatively high. The expression in (20) also makes clear that our model generates time-varying elasticities of demand in a simple way. In periods when z_i is relatively high, a large fraction of goods are demanded by consumers with a high demand elasticity (γ), and when z_i is relatively low, a large fraction of goods are demanded by consumers with a low demand elasticity (θ).

In the second stage of the consumer's problem, we solve, in the standard way, the intertemporal problem for the composite demands $c_A(s^t)$ and $c_B(s^t)$ as well as the rest of the allocations.

Firms

Consider now the problem of a firm in this economy. The firm has menu costs, measured in units of labor, of changing its prices. Let $P_R(s^{t-1})$ denote the firm's regular price from the previous period that is a state variable for the firm at the subsequent s^t. The firm has three options for the price it sets after the history s^t: pay nothing and charge the regular price $P_R(s^{t-1})$; pay a fixed cost κ and change the regular price to $P_R(s^t)$; or pay a fixed cost ϕ and have a temporary price change in the current period. Having a temporary price change at s^t entitles a firm to charge a price different from its inherited regular price $P_R(s^{t-1})M$ for that one period t only. If the firm wants to continue that temporary price change into the next period, it must again pay ϕ. In the period after the period of a temporary price change, the firm inherits the existing regular price $P_R(s^{t-1})$.

In this simple model, the only role of temporary price changes is to economize on the costs of changing prices. Firms face a mixture of shocks—some more permanent and some more temporary. Given this mixture of shocks, firms sometimes choose to change their prices temporarily and sometimes choose to change them permanently.

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

$$R\big(P_i(s^t);s^t\big) = \big(P_i(s^t) - W(s^t)\big)q_i(s^t),$$

where we have used the demand function (20). The present discounted value of profits of the firm, expressed in units of period 0 money, is given by

$$(21)\quad \sum_t \sum_{s^t} Q(s^t)\big[R_i\big(P_i(s^t);s^t\big) - W(s^t)\big(\kappa\delta_{R,i}(s^t) + \phi\delta_{T,i}(s^t)\big)\big],$$

where $\delta_{R,i}(s^t)$ is an indicator variable that equals one when the firm changes its regular price and zero otherwise, and $\delta_{T,i}(s^t)$ is an indicator variable that equals one when the firm has a temporary price change and is zero otherwise. In expression (21), the term

$$W(s^t)\big(\kappa\delta_{R,i}(s^t) + \phi\delta_{T,i}(s^t)\big)$$

is the labor cost of changing prices, or the *menu cost*. The constraints are that $P_i(s^t) = P_R(s^{t-1})$ if $\delta_{R,i}(s^t) = \delta_{T,i}(s^t) = 0$, that there is neither a regular nor a temporary price change, and that $P_i(s^t) = P_R(s^t)$ if $\delta_{R,i}(s^t) = 1$, so that there is a regular price change.

Equilibrium

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

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

requires that the sum of the labor used in production and the menu costs (measured in units of labor) of making both regular and temporary price changes equals total labor. The market-clearing condition on bonds is $B(s^t) = 0$.

An *equilibrium* for this economy is a collection of allocations for consumers $\{c_i(s^t)\}_i$, $M(s^t)$, $B(s^{t+1})$, and $l(s^t)$; prices and allocations for firms $\{P_i(s^t), y_i(s^t)\}_{i,}$; and aggregate prices $W(s^t)$, $P_A(s^t), P_B(s^t)$, 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 and transfers satisfy the specifications above.

Writing the equilibrium problem recursively will be convenient. At the beginning of s^t, after the realization of the current monetary, productivity, and demand shocks, the state of an individual firm i is characterized by its regular price in the preceding period, $P_{Ri}(s^{t-1})$; its idiosyncratic productivity level, $a_i(s^t)$; and the idiosyncratic demand for its good, $z_i(s^t)$. Normalizing all of the nominal prices and wages by the current money supply is convenient. For real values, we let $p_{R-1,i}(s^t) = P_{Ri}(s^{t-1})/M(s^t)$ and $w(s^t) = W(s^t)/M(s^t)$ and use similar notation for other prices. With this normalization, we can write the state of an individual firm i in s^t as $[p_{R-1,i}(s^t), a_i(s^t), z_i(s^t)]$.

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 $x = (p_{R,-1}, a, z)$ and the aggregate state variables as $S = (\mu, \lambda)$. Let

$$(22) \quad R(p_i, a, z, S) = \left(p_i - \frac{w(S)}{a}\right) q(p_i, z, S),$$

where real wage $w(S)$ and quantity demanded of good i $q(p_i, z, S)$ are known functions of the aggregate state. The function R is the static gross profit function, normalized by the current money supply M. Let $\lambda' = \Lambda(\lambda, S)$ denote the *transition law* on the measure over the firms' state variables.

In any period, the value of a firm that does nothing (N)—does not change its price and instead uses its existing regular price—is

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

(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 $p_T \neq p_{R,-1}$ is

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

and that of a firm that changes its regular (R) price is

$$V^R(p_{R,-1}, a, z; \mu, \lambda) = \max_{p_R} [R(p_R, a, S) - \kappa w(S)] + E\left[\sum_{S'} Q(S', S) V(p_R, a'; \mu', \lambda') | a, z\right].$$

Recall the intuitive way to think about the difference between a temporary and a regular price change. A temporary price change corresponds to renting a new price today, for just one period, whereas a regular price change corresponds to buying a new price that can be used for more than one period in the future; hence, the new regular price has a capital-like feature. As the state variables drift away from the current state, the investment in a new regular price depreciates in value.

Inspection of the value function V^T makes clear that, conditional on having a price change, the optimal pricing decision for p_T is static, and the optimal temporary price sets the marginal gross profit $R_p(p, a, z, S) = 0$. Note that the optimal temporary price is

$$(23) \quad p_T = \frac{\varepsilon(p, z, S)}{\varepsilon(p, z, S) - 1} \left(\frac{1}{a}\right) w(S),$$

where $\varepsilon(p, z, S)$ is the demand elasticity of $q(p, z, S)$ derived from (20). Note that this price is a simple markup over the nominal unit cost of production and is exactly what a flexible price firm would charge when faced with such a unit cost. In contrast, if the regular price is changed, then the optimal pricing decision for the new regular price, p_R, is dynamic. (In particular, p_R will not typically equal p_T. This feature of our quantitative model differs from the corresponding one in our analytic exercise.)

As (23) makes clear, if a price is changed temporarily, then the inherited regular price

$p_{R,-1}$ is irrelevant, so we can write $p_T(a, S)$. Similarly, as inspection of the value function V^R makes clear, conditional on having a regular price change, the inherited regular price $p_{R,-1}$ is also irrelevant, so we can write $p_R(a, S)$.

B. Quantification and Prediction

We want to use the facts described above as the basis for our model and its evaluation. Now 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. We also go on to determine the model's implications for the real effects of a monetary shock, which we will later use to judge other models.

Functional Forms and Parameters

We set the length of the period as one week and therefore choose a discount factor of $\beta = .96^{1/52}$. We choose the value of ψ, the disutility of labor parameters, to ensure that without aggregate shocks, consumers supply one-third of their time to the labor market. We set γ, the elasticity of substitution for the high elasticity types, to be 6. This is at the high end of the substitution elasticities estimated for grocery stores.

Since our model is weekly, so is the process for money growth (15) in our numerical experiments. However, the highest frequency at which the U.S. Federal Reserve's monetary aggregate data are available is monthly. Thus, we pin down the model's autocorrelation ρ_μ and variance σ_μ^2 of weekly money growth by requiring the model to generate a monthly growth rate of money that has the same autocorrelation and variance as the Fed's measure of currency and checking accounts (M1) during 1989–97, the years for which the Dominick's price data are available.

The rest of the parameters are calibrated so that the model can closely reproduce the facts we have described which are based on those price data: κ, the (menu) cost the firm incurs when changing its regular price; ϕ, the cost of having a temporary price change; as well as the specifications of the productivity and demand shocks. We will discuss the values of these critical parameters after we display the model's predictions.

Now consider our specification of the two idiosyncratic shocks in our model. We begin with the productivity process. As (16) indicates, this process has persistence ρ_a. The distribution of the shocks $\varepsilon_i(s^t)$ requires special attention. Midrigan (2007) shows that when $\varepsilon_i(s^t)$ is normally distributed, a model like ours generates counterfactually low dispersion in the size of price changes. Midrigan argues that a fat-tailed distribution is necessary in order for the model to account for the distribution of the size of price changes in the data. 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 and are, conditional on arrival, uniformly distributed on the interval $[-\bar{\nu}, \bar{\nu}]$. This is the approach we take in our numerical experiments:

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

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

Paying special attention to the distribution of the productivity shocks is useful 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 demand shocks. To keep the model simple, we assume that the demand shock, z_t, follows a Markov chain, with $z_t \in \{z_l, z_m, z_h\}$ and transition probabilities

$$\begin{bmatrix} \lambda_s & 1-\lambda_s & 0 \\ \lambda_l & \rho_v & 1-\lambda_l-\rho_v \\ 0 & 1-\lambda_s & \lambda_s \end{bmatrix}.$$

Here the subscripts l, m, and h indicate the low, medium, and high values. Hence, ρ_v is the probability of staying in the medium demand state z_m, λ_s is the probability of staying in either the low demand state z_l or the high demand state z_h, and λ_l is the probability of transiting from the medium demand state to the low demand state. We normalize $z_l = 0$. Our parameterization of these shocks thus has five parameters $\{z_m, z_h, \lambda_s, \lambda_l, \rho_v\}$.

Predictions

We now show that the parameters of our parsimonious extension of a standard menu cost model can be chosen so that the model can account for the six facts about price changes we have documented. We detail those parameters as well as the predictions. We then give some intuition for how the model works.

The Facts

The particular parameters that matter for the facts about price changes are the parameters governing the costs of changing prices and the productivity and demand shocks. In setting these parameters, we target the 12 moments in the data checked in the last column of Table 1. These moments include the frequency of weekly price changes (including and excluding temporary price changes), the fraction of temporary price changes, the proportion of returns to the old regular price, the probability of a temporary price spell ending, the fraction of periods and goods sold in periods when prices are temporarily up and down, the fraction of prices at the annual mode, the fraction of prices below that mode, as well as the size and dispersion of price changes (including and excluding temporary price changes).

In Table 1 we see that with a particular set of parameters, our parsimonious model does a remarkably good job of reproducing the first five facts about prices that we have documented. The frequency of weekly price changes is high: 33% in the data and 31% in the model, with all prices included (but much lower both in the data, 2.0%, and in the model, 1.9%, when temporary price changes are excluded). The mean size of all price changes is high in both the data (17% for all price changes and 11% for regular price changes) and the model (16% and 11%), and the dispersion is high in both as well. The portion of price changes that are at the annual mode is also high: 58% in both the data and 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. Specifically, prices spend 30% of the time below the annual mode in the data and about 28% in the model. Most price changes are temporary: 94% in both the data and the model. Most temporary prices tend to return to the regular price that existed before the temporary change: 80% in the data and 90% in the model. We also see that temporary price changes are transitory: the fraction of weeks with temporary price changes that are followed immediately by weeks without temporary price changes is 46% in the data and 59% in the model. Finally, our model predicts, as in the data, that the fraction of goods sold in periods when firms charge temporary prices is disproportionately high. Even though these periods account for 24.3% in the data and 18.4% in the model, the fraction of output sold in these periods is close: 38% in the data and 34.5% in the model.

We also investigate our model's implications for some other moments that we have not directly used to parameterize the model. Recall the sixth fact from the data, displayed in Figure 3, that price changes are clustered in time, in the sense that for all prices the hazard of price changes drops sharply in the first two weeks after a price change and declines thereafter. Figure 4 reproduces the curves from Figure 3 and adds to them the hazard predicted by our model. Clearly, our model generates a pattern similar to that in the data.

In Table 1 we also consider statistics about the mean and interquartile range of deviations of the temporary prices from the regular prices, as well as the relative fraction of goods sold in periods with prices temporarily up and down. We see that for most of these, the model produces values similar to those in the data. Finally, as in the data, the model's slope coefficient of a regression of changes in quantities on changes in prices for regular price changes is smaller in absolute value for periods with regular price changes (-2.1 in the data vs. -2.2 in the model) than in periods with temporary price cuts (-2.9 in the data vs. 4.4 in the model).

Table 2 lists the parameter values that have allowed the menu cost model to best match the moments in the data. The menu cost of changing regular prices κ is .90% of a firm's steady-state labor expense. In contrast, the cost of a temporary price change ϕ is .44% of a firm's steady-state labor expense, or about 50% of the cost of changing the regular price. Productivity shocks arrive with probability $\lambda = .061$ and have an upper bound of $\bar{\nu} = .095$. Moreover, the productivity process is highly transitory; its persistence is $\rho_a = .991$.

The fraction of high elasticity consumers, ω, is .08. The distribution of demand shocks is $\{z_l, z_m, z_h\} = \{0, .047, .197\}$, and the parameters governing the Markov transition matrix are $\lambda_s = .369, \rho_v = .803$, and $\lambda_l = .072$. Thus, the medium demand state is most persistent, whereas firms that are in the low or high demand states expect to return to the medium with high probability $1 - \lambda_s = .631$. Finally, the low elasticity of type A consumers is equal to $\theta = 1.984$.

Now consider our model's prediction for the main point of all this analysis: the real effects of monetary shocks. Our summary measure of the real effect is the standard deviation (or volatility) of output, which is .72%. In this model, if a monetary shock has no real effect, then this standard deviation should be zero; and the larger is the real effect, the larger should be the standard deviation. We find this predicted value useful; it is a benchmark against which to compare the sizes of the real effects of monetary shocks predicted by other models.

The Workings of the Menu Cost Model

Our model works differently from existing menu cost models because of a firm's ability to use temporary price changes to respond to shocks. To understand our model's predictions, we 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 permanent price change. Briefly, we find that firms use temporary price changes primarily to respond to temporary shocks and use permanent price changes to respond to (more) permanent shocks.

Consider the firm's optimal decision rules in the quantitative menu cost model. These rules are a function of the individual states, namely, the normalized regular price $p_{R-1} = P_{R,-1}/M$, the current productivity level a, and the current demand shock 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 5. Since the demand shock takes on three values, we report the firm's decision rules for each of the three demand states: low demand z_l, medium demand z_m, and high demand z_h. Figure 5 shows two decision rules for each of these three demand states: the regular price $p_R(a)$, conditional on the firm's choice to change the regular price, and the temporary price $p_T(a)$, conditional on the firm's choice to set a temporary price.

Figure 5 also shows the regions of the state space in which the firm optimally chooses to make a regular price change (R), to make a temporary price change (T), or to not change its price (N). All three panels share a standard feature: if the current price $p_{R,-1}$ is close enough to both $p_R(a)$ and $p_T(a)$ (that is, if the price lies in the regions labeled N), then the firm finds it optimal to forgo paying any costs and just charge the regular price.

As we have noted above when discussing the value functions, the temporary price $p_T(a)$ is a constant markup over marginal cost given by (23), and it does not equal the regular price $p_R(a)$. The temporary price $p_T(a)$ in log space falls one-for-one with a for all values of a because the log of marginal costs falls one-for-one with a. In contrast, the regular price $p_R(a)$ differs from $p_T(a)$ because its choice reflects the dynamic considerations.

More interesting is the difference in behavior across different states of demand. In the medium demand state, if the firm does choose a price different from its existing regular price, then it always chooses a new regular price, never a temporary price. The productivity and monetary shocks are highly persistent, so the firm expects its new regular price to be close to what is statically optimal for a long period of time. Hence, the firm is willing to pay the large fixed costs to change its regular price.

In contrast, in the high demand state, if the firm chooses a price different from its existing regular price, then it always chooses a new temporary price. Here, the firm knows that the state of high demand is temporary and significantly different from the medium demand state. Therefore, the firm is better off paying the relatively small fixed cost in order to use a temporary price than paying a large fixed cost and have to change its regular price twice, since the firm knows the current state will not last long. Of course, if the state of high demand lasts for a second period, then the firm will again choose to have a temporary price change. In this sense, two periods of high demand can generate two periods of temporary price declines or increases.

The firm's optimal decision rules in the low demand state are somewhat subtler. The key difference between the low and high demand states is that the low demand state is fairly close to the medium demand state, whereas the high demand state is not. (That is, z_l is only about 5% lower than z_m, but z_h is about 20% higher than z_m.) If the firm's existing regular price is very far from what is currently statically optimal, then the firm changes its regular price. Its new regular price is essentially what it would charge if it were in the medium demand state today. In this sense, the firm realizes that the temporary state of low demand will pass quickly and makes a once-and-for-all adjustment to have a new price—a strategy that works well when the medium demand state resumes. If the firm's existing regular price is such that tomorrow if the medium state resumes it would be essentially in the inaction region, then the firm decides to have a temporary price increase today. The final subtle part is that when demand is low, the costs of having a price that differs from the statically optimal one are lower than when the demand state is medium because the lost profits are low when demand is low. Hence, the inaction region in the low demand state is wider.

4. Experiments

We have shown that our menu cost model with permanent and temporary price changes can reproduce the main features of Dominick's micro price data. Thus, we view our model as a reasonable laboratory in which to evaluate the two common approaches to dealing with temporary price changes. In our experiments, we focus on the common approaches that use the simple Calvo model of pricing with only permanent price changes because this model is most popular in the applied literature and is viewed as a simple approximation to an underlying menu cost model. When we compare the predictions of our benchmark menu cost model to those of the Calvo model using the two existing shortcuts, we find the same qualitative results as in our previous comparison: the temporary-changes-in shortcut understates the real effects of monetary shocks, whereas the temporary-changes-out shortcut overstates them. Then we offer a third shortcut, based on our benchmark menu cost model, which should provide better results than the current two. (In Appendix B, we evaluate existing shortcuts that use a menu cost model without a motive for temporary price changes.)

A. The Standard Shortcuts

The Calvo models we consider are similar to the menu cost model described above except that the Calvo models have time-dependent sticky prices and no temporary price changes. The consumers in this type of model are identical to those in our benchmark menu cost model. Firms are allowed to adjust their prices in an exogenous, costless, and random fashion as in the analytic exercise discussed earlier. Specifically, in a given period, with probability α a firm can change prices, and with probability $1-\alpha$ the firm cannot change prices. We refer to α as the *frequency of price changes.*

In such a model, the problem of a firm that is allowed to change prices in state s^t is

$$\max_{p_i(s^t)} \sum_{r=t}^{\infty} \sum_{s^r} (1-\alpha)^{r-t} Q(s^r|s^t) R\big(P_i(s^t); s^r\big),$$

where

$$R\big(P_i(s^t); s^r\big) = \big(P_i(s^t) - W\left(s^r\right)\big) \left(\frac{P_i(s^t)}{P(s^r)}\right)^{-\theta} c(s^r).$$

Since changing prices is costless, the resource constraint is simply

$$l(s^t) = \int_i l_i(s^t)\, di.$$

In the Calvo models, the parameters of technology, preferences, and stochastic processes are set to be equal to those in our benchmark model. (See Table 2.) The additional parameter that needs to be set is α. We consider two parameterizations corresponding to the two shortcuts discussed above.

In the *temporary-changes-out* approach, we filter the data using the same algorithm as before, in order to remove temporary price changes, and treat the resulting regular price series as the relevant data. We then choose the frequency of price changes in the Calvo model, $1-\alpha$, so as to reproduce the frequency of regular price changes. In the temporary-changes-in approach, we choose the frequency of price changes so as to reproduce the frequency of all price changes.

In the Calvo temporary-changes-out model, we set $\alpha = .02$ so that the average duration of prices is 50 weeks. In the Calvo temporary-changes-in model, we set $\alpha = .33$ so that the average duration of prices is 16.7 weeks. We leave all other parameters unchanged.

In Table 3 we report these Calvo models' predictions for the size of the real effects of monetary shocks and compare them to our menu cost model's. We see that neither Calvo model predicts effects close to those of the benchmark model. The standard deviation of real output from the temporary-changes-out approach is 68% higher than the benchmark model's, and the standard deviation of output from the temporary-changes-in approach is only about 11% of the benchmark model's. Neither shortcut to modeling temporary price changes thus appears adequate for modeling or evaluating monetary policy.

B. An Alternative Shortcut

So far the only alternative we have offered to the inadequate existing shortcuts in the literature is to build an extended menu cost model that explicitly includes motives for temporary as well as permanent price changes. However, we acknowledge that some researchers may find implementing this alternative computationally difficult. For such researchers, we suggest an alternative, theory-based shortcut: use the simple Calvo model, but adjust the model's duration of price changes so that the Calvo model mimics the real effects of our menu cost model.

Figure 6 gives a sense of what such an adjustment entails and how well it may work. In that figure, we plot the standard deviation of output in a simple Calvo model with duration of prices $T = 1/\alpha$ in weeks relative to the standard deviation of output in the menu cost model. We see that when the duration of prices in the simple Calvo model is set equal to 16.7 weeks, the real effects in the two models are equal. This, then, is the duration price we recommend simple Calvo models use for monetary policy analysis.

5. Conclusion

In the data, a sizable fraction of price changes are temporary. Existing sticky price models abstract from explicitly modeling these changes. Should they? We have demonstrated here that they should not. Neither of the existing approaches to handling temporary price changes in the data provides predictions of the real effects of monetary policy that are near those of a menu cost model with an explicit motive for temporary price changes which is consistent with the price data. The temporary-changes-out approach leads to much larger effects than those of menu cost model, and the temporary-changes-in approach leads to much smaller effects.

A key insight to explain these results has to do with the nature of the monetary shocks: they are permanent. Their effects, therefore, cannot be expected to be offset, or diminished to a great degree, by temporary price changes alone. Despite their high frequency temporary price changes cannot offset the effects of monetary shocks as much as would equally frequent permanent price changes. In this sense, temporary and permanent price changes act quite differently in an economy. Models that don't treat these two types of price changes differently or that ignore one of them completely will thus naturally provide poor predictions of monetary policy effects.

We have offered two theory-based alternatives to the common approaches to handling temporary price changes in the data. One alternative, of course, is to build and use a menu cost model like ours, which explicitly includes a motive for temporary price changes. We have shown that this parsimonious extension of the standard menu cost model can be made to account for many of the patterns of price changes in the data. A cruder but simpler alternative to using a new model is to stick with the simple Calvo model but instead use our analysis to set the model's frequency of price adjustment. Even this crude theory-based approach is likely to produce better monetary policy analysis than do the approaches commonly used today.

6. Appendixes

A. The Algorithm to Construct the Regular Price

Here we describe, intuitively and precisely, our algorithm for constructing a regular price series for each product in the Dominick's data base.

The 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.[3] 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.[4] 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 actual price does not change, but only if the actual 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.

Examples of regular price series constructed using this algorithm are displayed in Figures 1–4.

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

1. Choose parameters: $l = 5$ (= *lag*, or size of the window: the number of weeks before or after the current period used to compute the modal price), $c = 1/3$ (= *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 $(t-l, ..., t+l)$ is $\geq 2al$, then
     - Let $p_t^M = mode(p_{t-l}, ..., p_{t+l})$.
     - 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 $t = 2+l, ...T$
     - 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 $\mathcal{R} = \{t : p_t^R \neq p_{t-1}^R \text{ \& } p_{t-1}^R \neq 0 \text{ \& } p_t^R \neq 0\}$ be the set of periods with regular price changes.
   - Let $\mathcal{C} = \{t : p_t^R = p_t \text{ \& } p_t^R \neq 0 \text{ \& } p_t \neq 0\}$ be the set of periods in which a store charges the regular price.
   - Let $\mathcal{P} = \{t : p_{t-1}^R = p_{t-1} \text{ \& } p_{t-1}^R \neq 0 \text{ \& } p_{t-1} \neq 0\}$ 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 $p_{\{\mathcal{R}\cap\mathcal{P}\}}^R = p_{\{\mathcal{R}\cap\mathcal{P}\}-1}$.

B. An Alternative Menu Cost Model

Since Calvo models are by far the most popular in applied work related to monetary policy analysis, we have focused our attention on them. For completeness, however, we also performed analogous experiments with simple menu cost models using two common approaches. Our results are qualitatively similar.

Table A1 reports our choice of parameter values for the simple menu cost model without a motive for price changes. This type of model abstracts from demand shocks and assumes that the measure of type B (high elasticity) consumers is constant at 0. We calibrate the frequency and size of price changes by choosing the arrival rate of productivity shocks and the upper bound of the support of their distribution.

Table A2 displays the real effects of monetary shocks predicted by these models. We find again that the temporary-changes-out approach overstates the real effects—now by about 40%. Similarly, the temporary-changes-in approach again understates the real effects: it predicts real effects that are only about 20% of those predicted by our menu cost model with a motive for temporary price changes.

In Table A1 we can identify one discrepancy between the menu cost model without a motive for temporary price changes and the data (and thus our benchmark model): the simple menu cost model misses the fraction of prices at the annual mode. In the data, that fraction is 58%. The temporary-changes-in approach underpredicts the fraction as 22%; the temporary-changes-out approach overpredicts it as 77%.

Finally, we show that an alternative superior approach to shortcuts that include or exclude temporary price changes in the data is to choose parameters in models without a motive for temporary price changes by matching the fraction of annual prices at the mode. When we do that, the implied frequency of price changes, reported in the last column of Table A1, is .051, or about once every 20 weeks. In Table A2 we see that this alternative parameterization predicts real effects of monetary shocks that are similar to those of the benchmark model: it overstates those effects by only 8%.

Notes

[1]The data used by Bils and Klenow (2004) and Nakamura and Steinsson (forthcoming) have a much wider set of products than the grocery store data, but the data are only collected as point-in-time prices at the monthly frequency. These monthly data thus provide no direct evidence about the critical issue of how many temporary price changes happen within a month. To see how much of a quantitative difference the use of weekly versus monthly data makes, note that in the weekly Dominick's data, prices change every three weeks, whereas in the monthly data, prices change every nine weeks.

[2]We obtain similar results if we compute a product-specific hazard and then a weighted average of each of the hazards using each product's share of total sales as the weight.

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

[4]If in the window around this price more than half of the data are missing, then we set the initial reference price equal to the actual price.

References

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

Calvo, Guillermo A. 1983. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12 (September): 383–98.

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 (March): 15–37.

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

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

Hansen, Gary D. 1985. Indivisible labor and the business cycle. *Journal of Monetary Economics* 16 (November): 309–27.

Hoch, Stephen J., Xavier Drèze, and Mary E. Purk. 1994. EDLP, hi-lo, and margin arithmetic. *Journal of Marketing* 58 (October): 16–27.

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

Nakamura, Emi, and Jón Steinsson. Forthcoming. Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics*.

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

Rotemberg, Julio J. 2004. Fair pricing. Working Paper 10915, National Bureau of Economic Research.

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

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

Table 1: Five facts about prices *

In the data and the benchmark menu cost model

	Line	Moments	Data: Including temporary changes	Data: Excluding temporary changes	Benchmark model: Including temporary changes	Benchmark model: Excluding temporary changes	Used for calibration?
FACT 1	1	Frequency of weekly price changes	0.33	0.020	0.31	0.019	√
	2	Fraction of temporary price changes	0.94	-	0.94	-	√
	3	Proportion of returns to regular price	0.80	-	0.90	-	√
	4	Probability temporary price spell ends	0.46	-	0.59	-	√
FACT 2	5	Fraction of periods with temp prices, %	24.3	-	18.4	-	√
	6	Fraction of periods when price temp. down, %	20.3	-	16.0	-	-
	7	Fraction of periods when price temp. up, %	2.1	-	2.1	-	-
FACT 3	8	Fraction of prices at annual mode	0.58	-	0.58	-	√
	9	Fraction of prices below annual mode	0.30	-	0.28	-	√
FACT 4	10	Mean size of price changes	0.17	0.11	0.16	0.11	√
	11	Mean $\log(p_T/p_R)$ if price temporarily down	-0.22	-	-0.23	-	√
	12	Mean $\log(p_T/p_R)$ if price temporarily up	0.13	-	0.08	-	√
	13	IQR of price changes	0.15	0.08	0.17	0.04	√
	14	IQR $\log(p_T/p_R)$ if price temporarily down	0.21	-	0.13	-	-
	15	IQR $\log(p_T/p_R)$ if price temporarily up	0.12	-	0.01	-	-
FACT 5	16	Fraction output sold when temp prices, %	38.0	-	34.5	-	√
	17	Fraction output sold when price temp. down, %	35.4	-	32.6	-	-
	18	Fraction output sold when price temp. up, %	1.2	-	1.7	-	-
	19	Price elasticity for regular price changes	-2.08	-	-2.19	-	-
	20	Price elasticity for temporary price cuts	-2.93	-	-4.40	-	-

Note: IQR=Interquartile range.

* For Fact 6, see Figures 3-4.

Table 2: Parameter values for menu cost and Calvo models

Parameters	Benchmark model	Calvo model without temporary prices: Temporary-changes-in	Calvo model without temporary prices: Temporary-changes-out
Calibrated			
Cost of changing regular price, % of SS labor bill	0.90		
Cost of temporary markdown, relative to menu cost	0.44		
Arrival rate of productivity shock	0.061		
Upper bound on productivity	0.095		
Persistence of productivity	0.991		
Value of demand shock in medium demand state	0.047		
Value of demand shock in high demand state	0.197		
Substitution elasticity of low elasticity consumers	1.984		
Measure of high elasticity consumers	0.08		
Probability of staying in non-medium demand state	0.369		
Probability of staying in medium demand state	0.803		
Probability of jumping from medium to low demand state	0.072		
Probability of Calvo price adjustment	-	0.33	0.02
Assigned			
Discount factor	$0.96^{1/52}$	$0.96^{1/52}$	$0.96^{1/52}$
Substitution elasticity of high elasticity consumers	6	-	-
Autocorr. of growth rate of money supply	0.90	0.90	0.90
Std. dev. of shocks to growth rate of money supply	$8.31*10^{-4}$	$8.31*10^{-4}$	$8.31*10^{-4}$

Table 3: Real effects of monetary shocks predicted by model economies

		Calvo model without temporary prices	
	Benchmark model	Temporary-changes-in	Temporary-changes-out
Std. dev. of chain-weighted real output, %	0.72	0.08	1.21
Relative to benchmark model	-	0.11	1.68

Table A1: Data moments and parameter values in menu cost models using shortcuts

	Data		Menu cost model without temporary prices		
	Including temporary changes	Excluding temporary changes	Temporary-changes-in	Temporary-changes-out	Match fraction of prices at annual mode
Moments					
Frequency of weekly price changes	0.33	0.020	0.31	0.019	0.051
Mean size of price changes	0.17	0.11	0.16	0.11	0.16
Fraction of prices at annual mode	0.58	-	0.22	0.77	0.58
Calibrated parameters					
Arrival rate of productivity shock			0.750	0.012	0.049
Upper bound on productivity			0.152	0.100	0.132
Assigned parameters					
Cost of changing regular price, % of SS labor bill			0.90	0.90	0.90
Persistence of productivity process			0.84	0.84	0.84
Discount factor			$0.96^{1/52}$	$0.96^{1/52}$	$0.96^{1/52}$
Substitution elasticity of low elasticity consumers			3	3	3
Substitution elasticity of high elasticity consumers			-	-	-
Autocorr. of growth rate of money supply			0.90	0.90	0.90
Std. dev. of shocks to growth rate of money supply			$8.31*10^{-4}$	$8.31*10^{-4}$	$8.31*10^{-4}$

Table A2: Real effects of monetary shocks in menu cost models

	Benchmark model	Menu cost model without temporary prices		
		Temporary-changes-in	Temporary-changes-out	Match fraction of prices at annual mode
Std. dev. of chain-weighted real output	0.72	0.15	1.00	0.78
Relative to benchmark model	-	0.21	1.39	1.08

Figure 1: Illustration of temporary and permanent price changes

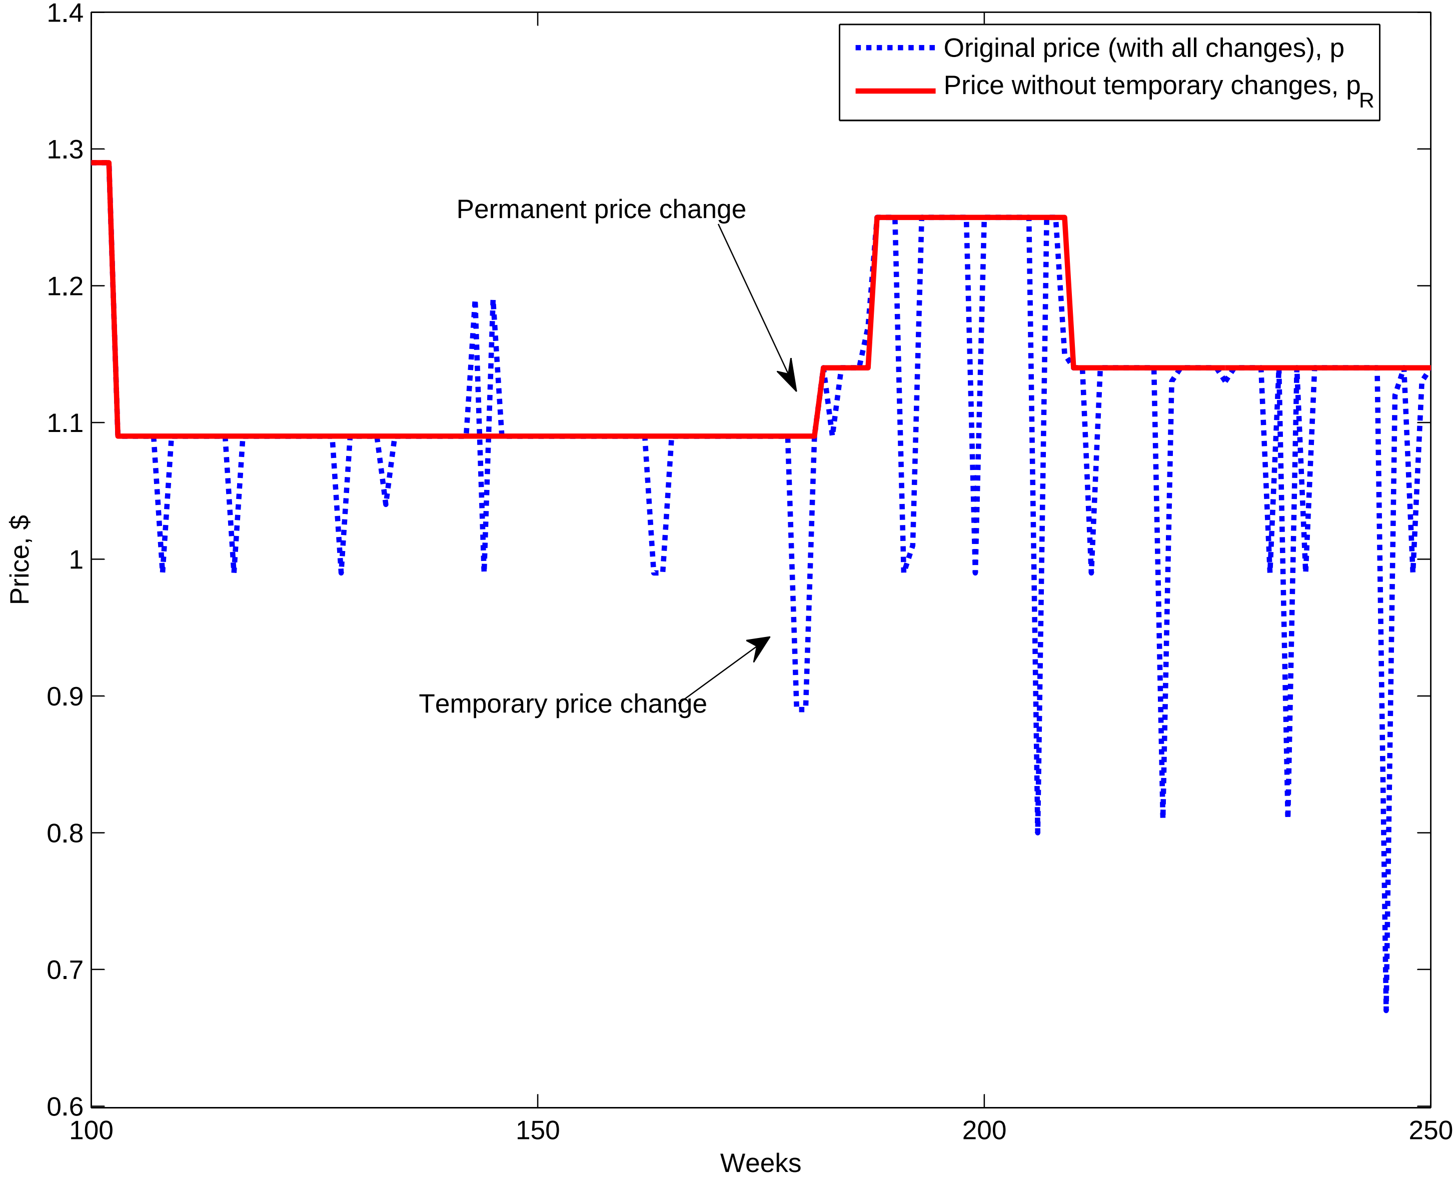

Figure 2: Examples of applying the algorithm to categorize price changes as temporary or permanent

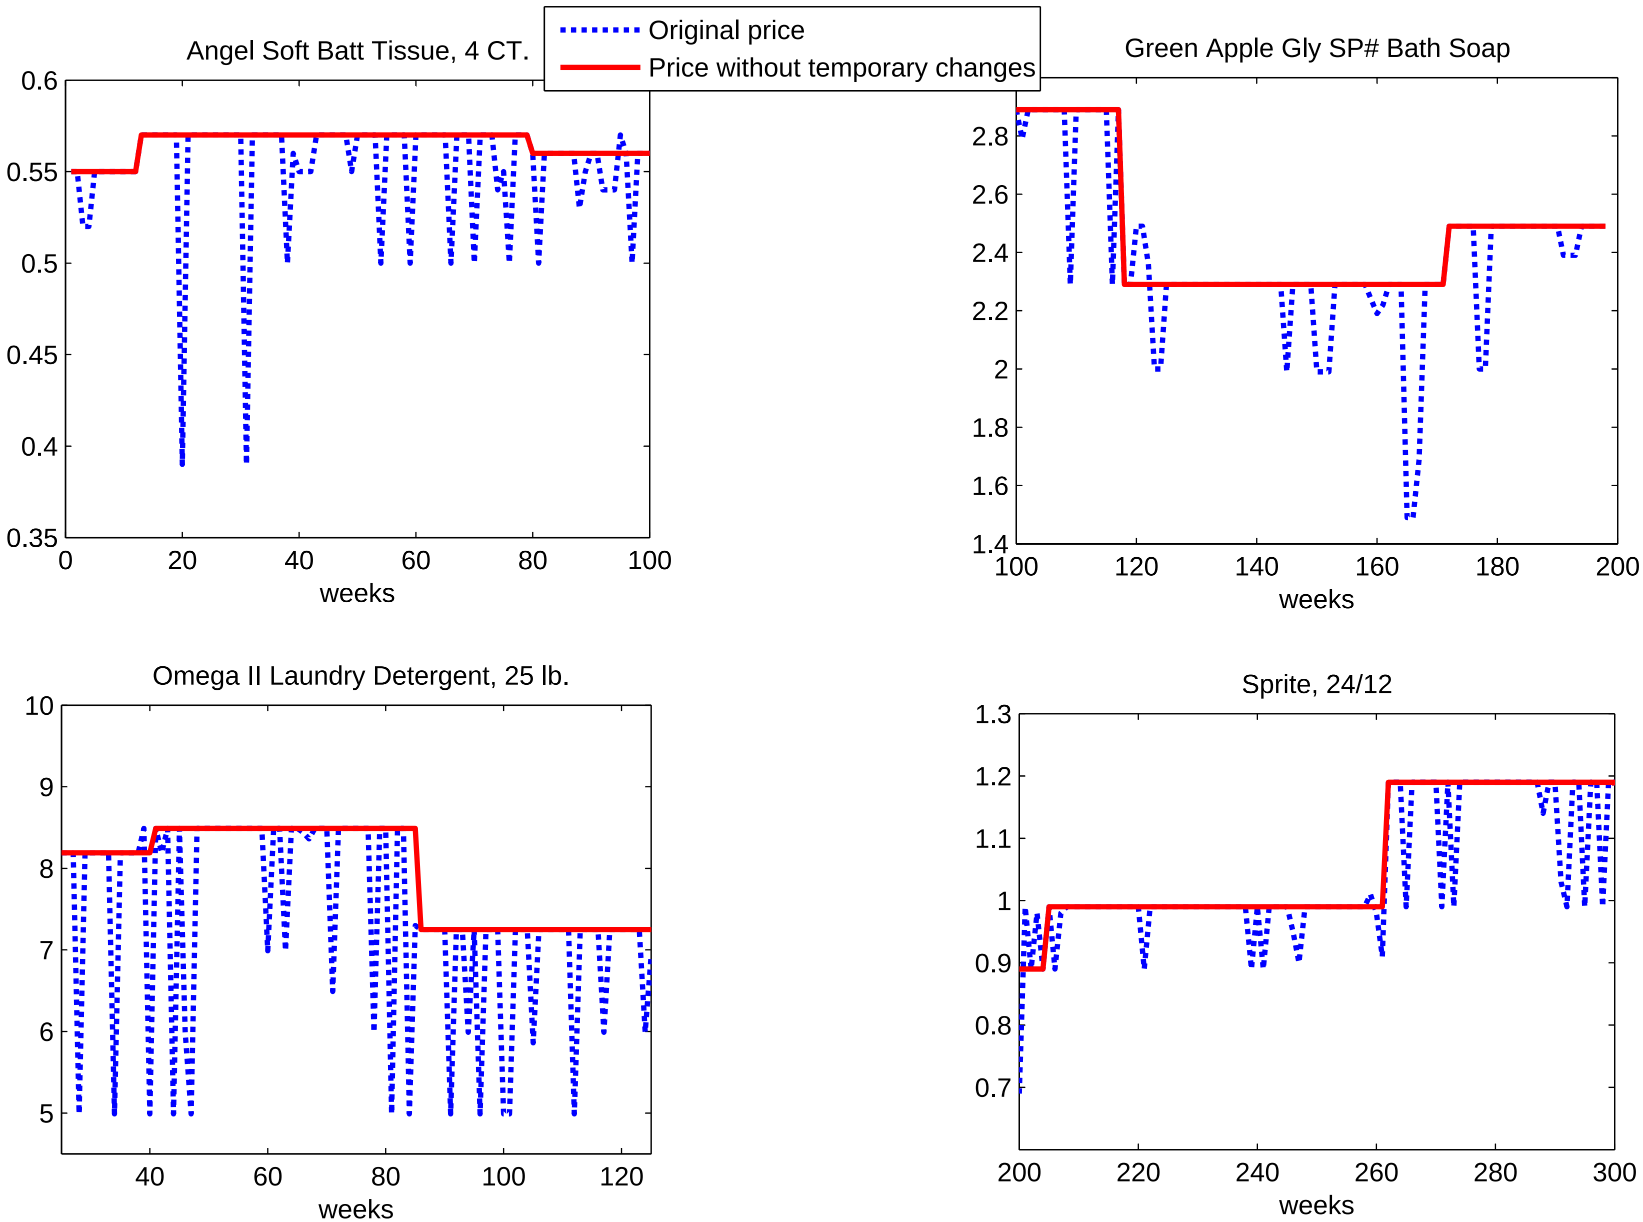

Figure 3: Fact 6 -- The hazard of price changes in Dominick's data

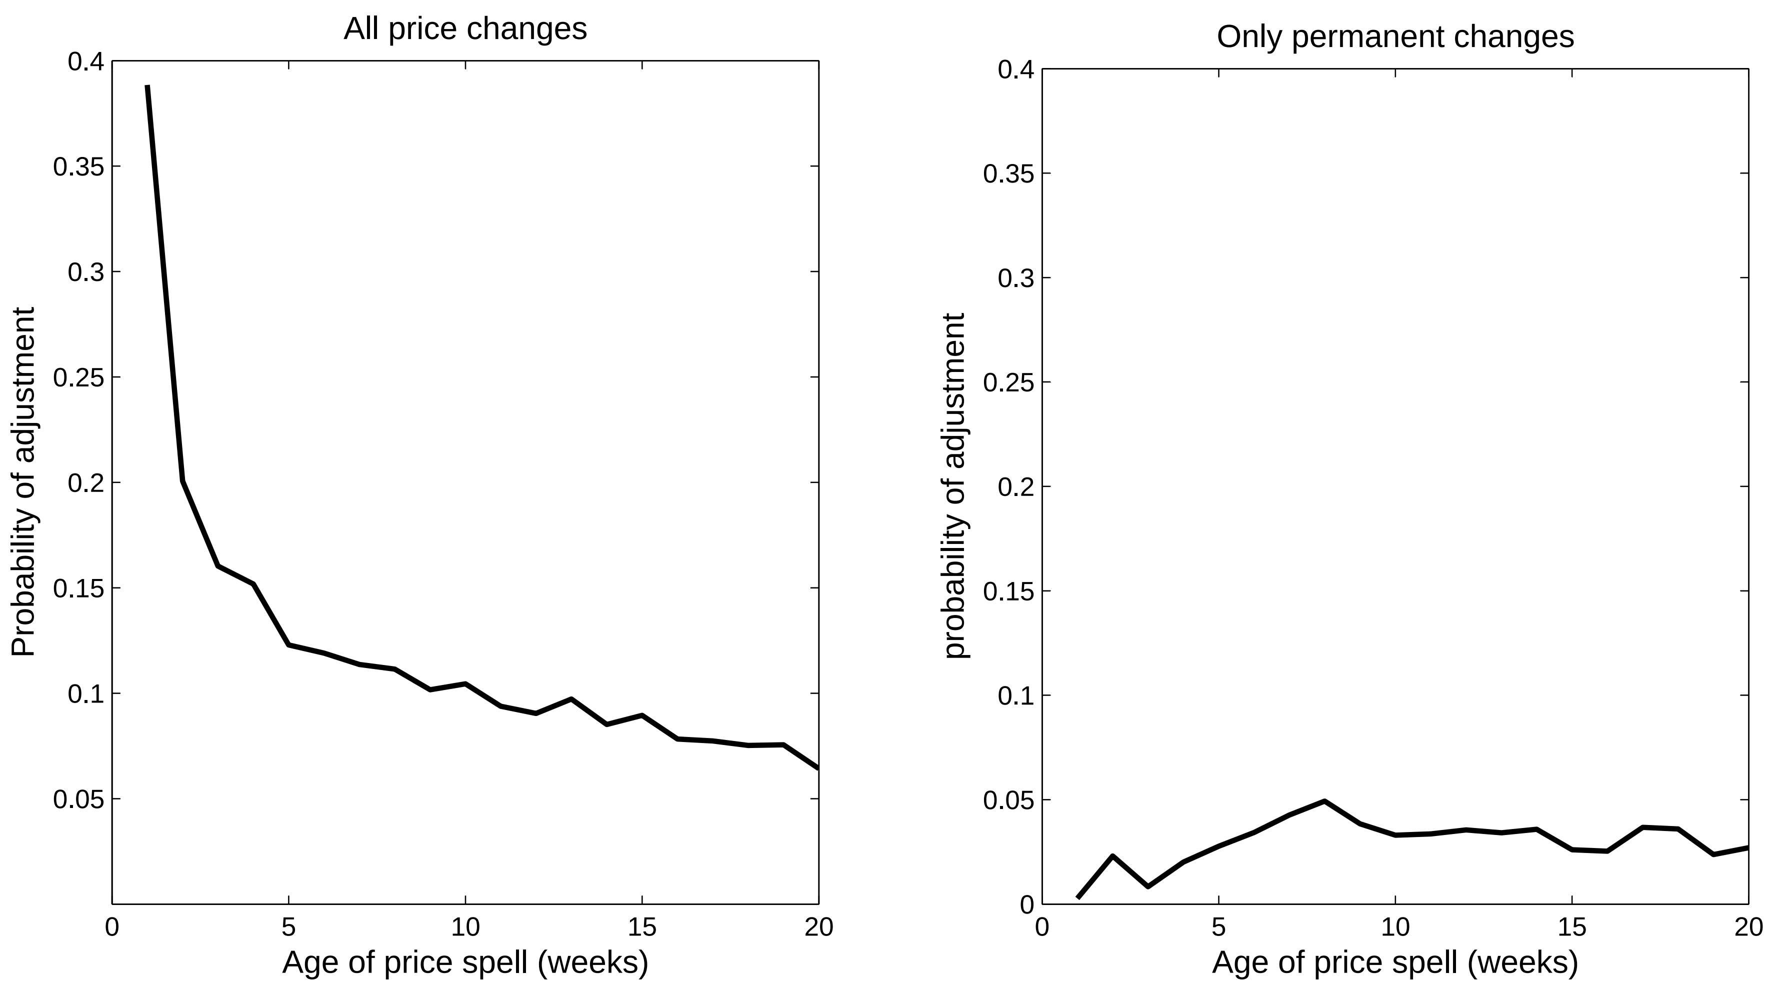

Figure 4: The hazard of price changes predicted by the benchmark model vs. Dominick's data

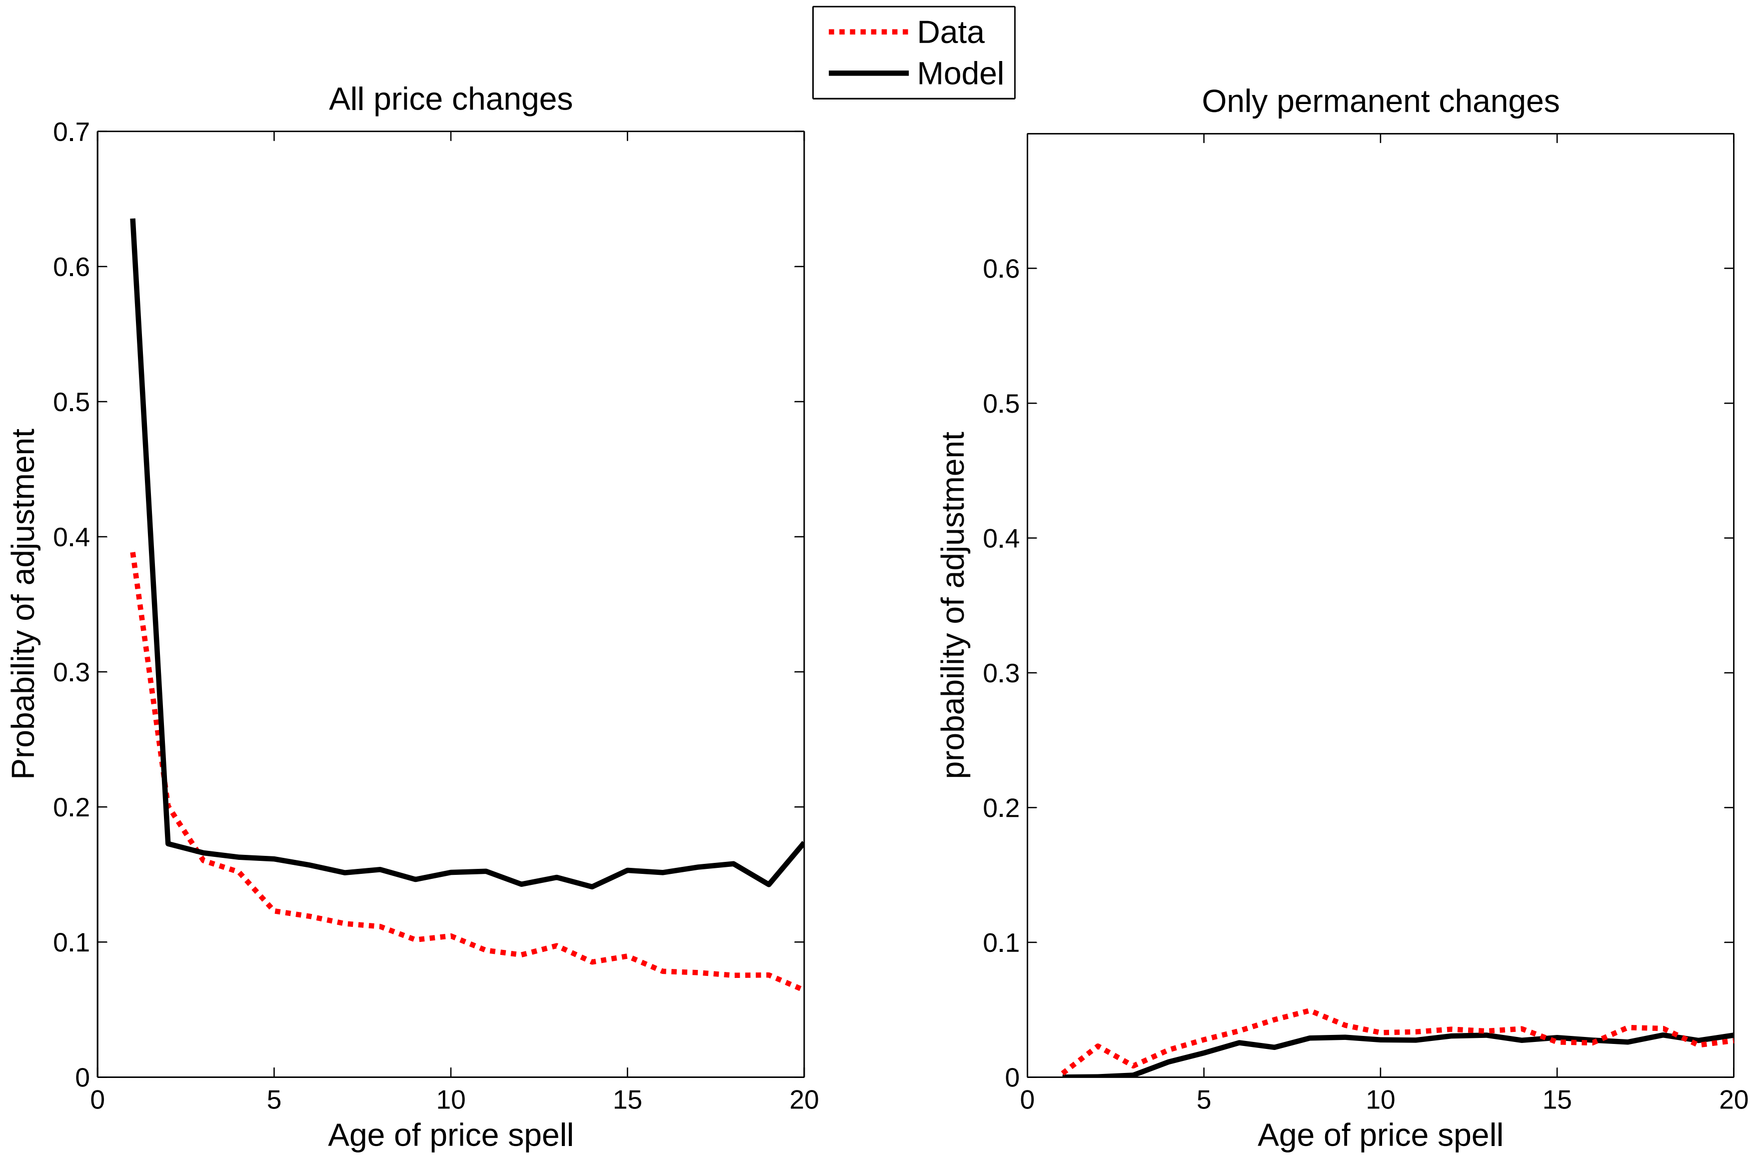

Figure 5: A firm's optimal decision rules in the benchmark model

Regions where firm's choice is R = Regular price change, T = Temporary price change, N = No price change

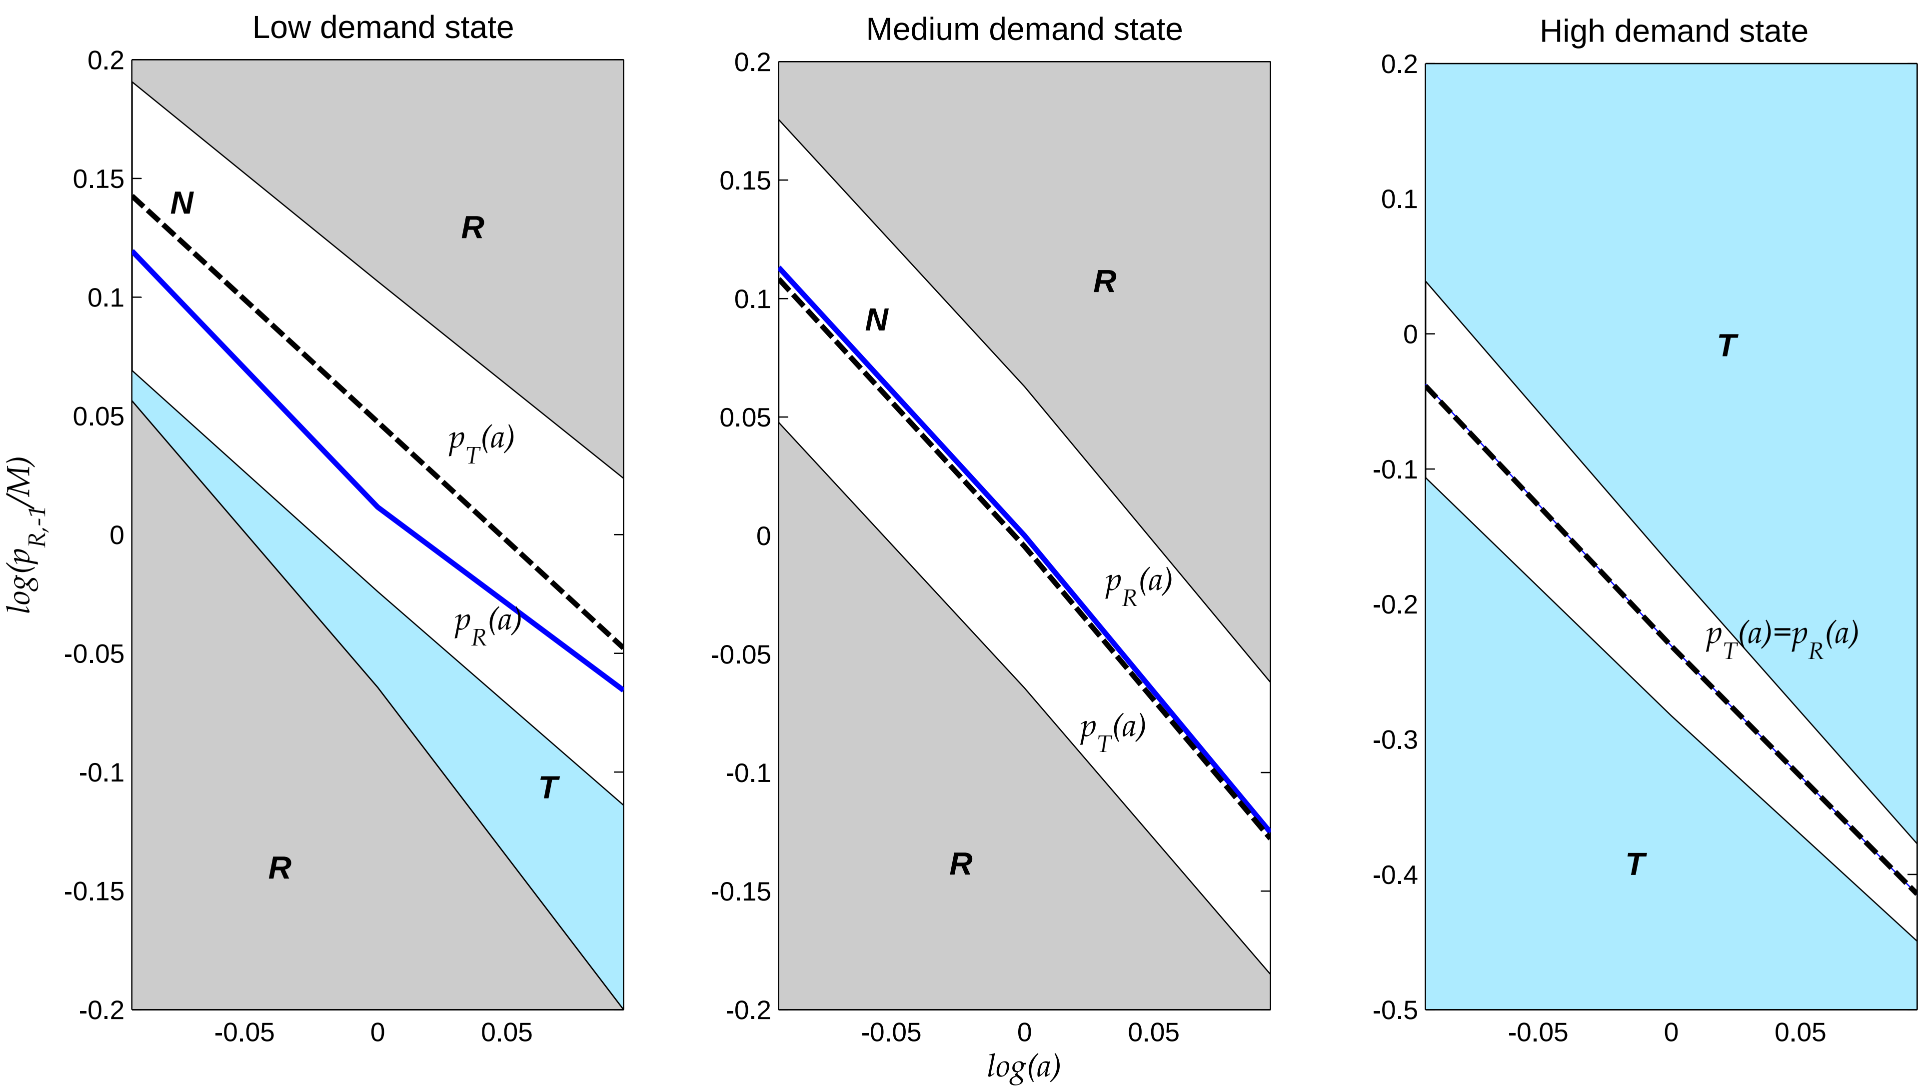

Figure 6: Real effects of monetary shocks in Calvo alternatives vs. the benchmark model