Federal Reserve Bank of Minneapolis
Research Department

Temporary Price Changes and the
Real Effects of Monetary Policy[*]

Patrick J. Kehoe and Virgiliu Midrigan

Working Paper 661

May 2008

ABSTRACT _____

In the data, a large fraction of price changes are temporary. We provide a simple menu cost model which explicitly includes a motive for temporary price changes. We show that this simple model can account for the main regularities concerning temporary and permanent price changes. We use the model as a benchmark to evaluate existing shortcuts that do not explicitly model temporary price changes. One shortcut is to take the temporary changes out of the data and fit a simple Calvo model to it. If we do so prices change only every 50 weeks and the Calvo model overestimates the real effects of monetary shocks by almost 70%. A second shortcut is to leave the temporary changes in the data. If we do so prices change every 3 weeks and the Calvo model produces only 1/9 of the real effects of money as in our benchmark. We show that a simple Calvo model can generate the same real effects as our benchmark model if we set parameters so that prices change every 17 weeks.

_____

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 quantify the real effects of monetary shocks assume that goods prices are sticky. A key ingredient in these models that determines the size of the real effects of monetary shocks is the frequency of price changes: if firms change prices frequently, then monetary shocks have small real effects; if they change them infrequently, monetary shocks have large real effects.

In the data a large fraction of price changes are temporary and a small fraction are permanent. Hence, the answer to the question of how frequently do prices change in the data crucially depends on how researchers treat temporary price changes. If temporary price changes are included in the data, prices are fairly flexible; if they are excluded, prices are fairly sticky. To see just how big a difference these two approaches makes to the data consider Figure 1, which is a fairly typical price series in our data set. When we leave temporary changes in we obtain the dashed line, in which prices change very frequently, while if we take them out we obtain the solid line in which prices change rarely.

Existing models make no attempt to distinguish between temporary and permanent price changes. Hence, when confronted with data in which a large fraction of price changes are temporary, the models take one of two shortcuts. The most popular approach is to exclude temporary price changes from the data, write down 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 approach is to include temporary price changes in the data, write down 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. How the temporary price changes are treated matters greatly for how large are the real effects of monetary shocks and existing theory provides little guidance as to how to proceed. We provide some.

We provide a simple menu cost model which explicitly includes a motive for temporary price changes and hence provides an alternative to either of the existing shortcuts. We show that this simple model can account for the main regularities concerning temporary and

permanent price changes. We also use the model as a benchmark to evaluate the existing shortcuts taken in Calvo models. We find that if we take the temporary changes out of the data prices change 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 every 3 weeks and the Calvo model produces only 1/9 of the real effects of money as in our benchmark. We show that a simple Calvo model can generate the same real effects as our benchmark model if we set parameters so that prices change every 17 weeks.

We start by documenting six regularities concerning temporary and permanent price changes that we use in our quantitative analysis. (We discuss the algorithm we use to identify temporary and permanent price changes in the appendix.) Before we turn to our benchmark model, which is a menu cost model with a motive for temporary price changes, we begin with a motivating example. This example is the simplest possible model of temporary and permanent price changes that we can solve using pen and paper in order to build some intuition.

In the motivating example we use a Calvo model of price setting modified to have temporary as well as permanent price changes. We use this model to illustrate the gist of our argument. We treat the model as the data-generating process and solve in closed form for an expression governing the effects of money. We then solve for similar closed form expressions for these real effects under the temporary-changes-out approach and the temporary-changes-in approach. We prove that the temporary changes out approach overstates the real effects of money while the temporary in approach understates the real effects of money.

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

In the model firms are subject to two types of disturbances: persistent productivity shocks as well as transitory shocks to the elasticity of demand for the firm's product. The latter shocks are meant to capture the idea that firms face demand for their products with time-varying elasticity.

In each period, the firm enters the period with a pre-existing regular price. This price is the price it can charge in the current period at no extra cost. If the firm wants to charge a different price in the current period, it has two options: change its regular price or have a temporary price change. To change its regular price, the firm pays 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 price change. To have a temporary price change the firm pays 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 periods in a row if it pays the rental cost each period.) We show that, essentially, it is optimal in this environment for firms to use a temporary price change to respond to a transitory taste shock. In contrast, it is optimal to use a regular price change to respond to the much more persistent monetary and productivity shocks.

We show the model can generate the salient features of the micro-price data, including the frequency of permanent and temporary price changes. We then use this model as a laboratory to study how well the two existing shortcuts approximate the real effects of money in our benchmark setup. We argue that the existing shortcuts to dealing with temporary price changes are likely to be inadequate in interesting applied settings.

We provide two alternatives. One alternative, is to use our simple extension of the price-setting technology in Golosov and Lucas model that explicitly includes a motive for temporary price changes. This extension can either be used as is or as the core of a richer one with more rigidities and shocks that may be useful in helping the model to mimic other features of the business cycle. A cruder but simpler alternative is to stick with a simple Calvo model and set the frequency of price adjustment so as to mimic the real effects in a model with temporary and permanent changes. We work out these formulas by hand when the model with both types of price changes is an extended Calvo model and we provide a quantitative benchmark when the model with both types of price changes in an extended menu cost model.

Our work is related to a growing literature that documents features of micro price data

in panel data sets. Two influential papers in this literature are Bils and Klenow (2004) and Nakamura and Steinsson (2007). When these researchers approach the data they focus on temporary price declines, referred to as *sales*, rather than all temporary price changes which is the sum of temporary price declines and temporary price increases. These researchers have found, as we do, that the frequency of price changes depends sensitively on the treatment of temporary price declines.

1. Price Changes in the Data

We begin by documenting six facts about price changes in the U.S. data that we will use both to calibrate and evaluate our model.

The source of our data 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 (GSB) (the James M. Kilts Center for Marketing). The data consist of nine years (1989 to 1997) of weekly store-level data from 86 stores in the Chicago area on the prices of more than 4,500 individual products which are organized into 29 product categories[1]. The products available in this data base range from non-perishable foodstuffs (some of which are represented by the categories of frozen and canned food, cookies, crackers, juices, sodas, and beer) to various household supplies (some of which are represented by the categories of detergents, softeners, and bathroom tissue) as well as pharmaceutical and hygienic products. (For details see the appendix.)

We use a simple algorithm (described in the appendix) to categorize all price changes as either *temporary* or *permanent*. To do so we define an artificial series called a *regular* price series, denoted $\{P_t^R\}$, which is constructed and used mainly to define which periods are periods of temporary price changes. The intuitive way to think about our analysis is to imagine that at any point in time there is an existing regular price and that there are two types of price changes: temporary price changes that tend to return to the regular price and permanent price changes in which the regular price itself is changed.

In Figure 2 we illustrate the results of our algorithm for price series. The raw data are represented by the dashed lines and the solid lines are the regular price series. Note that every data point is represented by a solid dot. Every price change that is a deviation from

the regular price line is defined as a temporary price change while every price change that is accompanied by a change in the regular price is defined as a permanent price change. Perusal of these pictures makes some facts about price changes clear: price changes are frequent and large, most of these changes are temporary, and most temporary prices return the pre-existing regular price. We turn now to a more formal description of the data that we will use in our theoretical model.

In Table 1 we report a variety of general facts about price changes that our data reveal. For each of the 29 product categories, we first computed category-level statistics by weighting each product by its share in total sales in each category. In Table 1 we report a weighted average of these category-level statistics, where the weights are each category's share in total sales. We summarize these features as follows

Fact 1: Prices change frequently, but most price changes are temporary and tend to return to the regular price.

To see this, notice from line 1 in Table 1 that the frequency of weekly price changes is 33% , so prices change on average every 3 weeks. However, most of these price changes are temporary, indeed 94% of them are temporary (line 4). Regular prices therefore change infrequently with a weekly frequency of 2%, so regular prices change only once a year. The temporary price changes are very short-lived, as they last for 2 weeks on average in that the probability that a temporary price change reverts is 46 % (line 8). Moreover, 80% of the time (line 7) they return to the pre-existing regular price.

Fact 2: Most temporary price changes are price cuts not price increases.

Of all the periods when the store charges a temporary price (24.3%, line 9), most of the time the price is temporarily down (20.3%, line 12), rather then up, 2.1%, line 11).

Fact 3. During a year, prices spend most of their time at their modal value and when prices are not at the mode they are much more likely to be below their annual mode than above it.

Table 1 also shows that, on average during a 50-week period, prices tend to be at their modal value 58% of the time. To see this, notice that when prices are not at their annual mode, they are most likely below it (30%, line 6). Table 1 shows that prices are below their

annual modal value 30% of the time and above it 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 a price change is 17% (line 2), and the interquartile range is 15% (line 13). The mean of regular price changes is 11%. Temporary price increases and decreases are also large and dispersed. The mean deviation of the temporary price from the regular price is -22%, (line 14) when the price is temporarily down and 13% (line 16) when the price is temporarily up. The interquartile range of these deviations is 21% (line 15) and 12% (line 16), respectively.

Fact 5. Temporary price cuts account for a disproportionate amount of goods sold. Quantities sold are more sensitive to prices in periods of temporary price declines than they are during periods of permanent price declines.

In the data 38% of output is sold in periods with temporary prices (line 10), 35.4% (line 19) when the price is temporarily down, and 1.2% when the price is temporarily up (line 18), even though the fraction of weeks accounted by these episodes is 24.3%, 20.3% and 2.1 % (see fact 3). To put another way, in periods of temporary price declines more than twice as many goods are sold relative to a period with regular prices. A regression of changes in quantities on changes in prices during regular price changes yields a slope coefficient of -2.08. A similar regression during episodes when the price change is from a regular price to a temporary decline yields a slope coefficient of -2.93. (Of course, the slope coefficient in our simple regression is not a true structural measure of demand elasticity. Nonetheless, we find it instructive to note that in static monopolistic competition setting a increase in a 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 $t + k$ given that 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, holiday and seasonal dummies, as well as by modelling age-dependence non-parametrically. Each

product is weighted according to its share in Dominick's total revenue in constructing the likelihood function. The figure reports the effect of varying the age of the price spell 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 reports the hazard for all price changes, including temporary and permanent. The panel shows that the hazard at one week is 53%. That is, conditional on a store changing the price of a given product last week, the store changes that price this week 53% of the time. More generally, we see that the hazard is sharply declining 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, so that there tend to be periods with many price changes followed by prolonged periods with no price changes. The right panel of the figure presents the hazards for the regular price changes. Here the hazard is flat and somewhat increasing in the first few weeks.

2. A Motivating Example

We take the simple Calvo model of price setting and extend it to have temporary and permanent price changes. We then solve for a closed form expression for the real effects of money and we analytically evaluate the two shortcuts to dealing with temporary price changes.

In the 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 money by the variance of aggregate consumption. We begin by briefly describing the economy and then solving for this variance as a function of the primitive parameters in the economy.

We borrow the consumer's side of the problem from the model of the next section. It is a standard cash-in-advance model with a consumer who has the choice of a continuum of differentiated consumption goods. In that section we lay out this problem in detail, here we just describe the key elements that we need to illustrate our points. In particular, consumer's preferences are defined over leisure and continuum of consumption goods such that given the

price of good i is P_{it} the 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

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

is the aggregate price index. Moreover, the utility function is such that the first order condition for labor is

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

and the cash-in-advance constraint binds:

$$(4) \quad P_t C_t = M_t$$

The supply of money is given by an exogenous stochastic process that follows $M_t = \mu_t M_{t-1}$, where $\log \mu_t$ is iid with mean 0 and variance σ_μ^2.

The firm side is more interesting. Each firm is the monopolistic supplier of a single good. Each firm enters a period with a preexisting regular price denote $P_{R,t-1}$. Absent one of two events the firm must charge its existing regular price $P_{R,t-1}$ in the current period. The first event, referred to as *permanent price change* occurs with probablity α_R and allows the firm to change this regular price to some new price P_{Rt}. The second event, referred to as a *temporary price change*, occurs with probablity α_T and allows the firm to charge a price P_{Tt} in this period that differs from its existing regular price $P_{R,t-1}$.

Note that a firm that experiences a temporary price change in the current period will charge P_{Tt} in the current period and price $P_{R,t-1}$ in the subsequent period unless that firm experiences again one of the two price changing events in the subsequent period.

Consider the problem of a firm that is allowed a temporary price change P_{Tt} at time

t. Clearly, the choice of this price has no influence on a firm's profits at any future date or state. Thus the firm simply solves the static problem of maximizing current profits

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

Here the optimal price is $P_{T,t} = \theta W_t / (\theta - 1)$. Note from (3) and (4) that $W_t = \psi M_t$ in equilibrium. It is convenient to normalize all nominal variables by the money stock. 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 of a firm facing a permanent price change. That is, at t the firm is allowed to reset its regular price $P_{R,t}$. Clearly, that firm needs only to consider the states for which the regular price it chooses today will be in effect. (This price has no effect either on future dates in which the firm can choose a temporary price or in future dates in which a new regular price will be in effect.)

Consider then the value of profits during those dates 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 get a permanent change, the objective is:

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

where here $Q_{t,s}$ is price of a dollar at s in units of dollars at t, normalized by the conditional probability of the state at s given the state at t. To understand this objective note that at t the prevailing price is $P_{R,t}$, at $t+1$ the prevailing price is $P_{R,t}$ with probability $1 - \alpha_T - \alpha_R$, at $t+2$ 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 it is easy to show that the log-linearized first order condition for this problem is

9

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

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

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

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 its sets it equal to zero.

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

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

Proof. We establish Proposition 1 as follows. To do so we use the cash-in-advance constraint (4). Log-linearizing this constraint gives that

(8) $c_t = -p_t.$

Thus we need only solve for the law of motion for the price index. From (2) this index is given by

(9) $p_t = \int_0^1 p_{it}di$

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

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

The average of existing regular prices at $t-1$ can be written recursively as first compute the law of motion for the average regular price

Consider next the law of motion for the average existing regular price $\bar{p}_{R,t}$. Given that α_R firms reset prices at t to $p_{R,t}$ and $(1 - \alpha_R)$ do not but instead use whatever their regular price was a $t-1$, with some manipulations we can write the law of motion for $\bar{p}_{R,t}$ recursively as

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

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

A. Evaluation of the two common approaches

Consider next our evaluation of the two common approaches. Consider a researcher who studies the data generated by our model with both temporary and permanent price changes through the lens of a simple Calvo model with only permanent price changes with a frequency of price change α. The researcher follows one of the two common approaches to calibrate the frequency of price changes in his model. In the *temporary- changes-out* approach we imagine that the researcher is able to isolate the permanent prices changes and thus concludes that the frequency of price changes, $\alpha = \alpha_R$. In the *temporary-changes-in* approach we imagine that researcher uses the raw data with the temporary price changes in and conclude that the frequency of price changes is, approximately, $\alpha = \alpha_R + 2\alpha_T$. The last expression is due to the fact that every time the firm changes its price temporarily (α_T of the time) is accompanied by 2 price changes, to, and from the temporary price.

To set up our proposition note that our derivation above implies that the standard

11

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

(11) $\quad var(c_t) = \dfrac{(1 - \alpha)^2}{1 - (1 - \alpha)^2}\sigma_\mu^2.$

Setting $\alpha = \alpha_R$ and $\alpha = \alpha_R + 2\alpha_T$ in (11) then yields the following proposition. Letting 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.

$\quad$ *Proposition 2.* The temporary-changes-out approach overstates the real effects of money while the temporary-changes-in understates the real effects of money. In particular, the temporary-changes-out approach predict

$$var(c_t^{Out}) > var(c_t) > var(c_t^{In})$$

$\quad$ *Proof.* Evaluating (11) at $\alpha = \alpha_R$ and $\alpha = \alpha_R + 2\alpha_T$ gives that

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

Clearly, the left-hand term in (12) is (11) evaluated at $\alpha = \alpha_R$, the middle term follows from (7) and the right-hand term is (11) evaluated at $\alpha = \alpha_R + 2\alpha_T$. *Q.E.D.*

$\quad$ Intuitively, the temporary-changes-out approach correctly predicts the persistence of consumption, but it overstates the volatility of innovations to the consumption process as it ignores the fact that a fraction α_T of firms change prices in any given period and thus offset the money change. In contrast, the temporary-changes-in approach understates the persistence of consumption as it fails to recognize that some of the price changes are temporary and revert to their previous value. Moreover, it counts the returns from the temporary to the permanent price as a price change that is useful in responding to the monetary disturbance,

whereas in fact it is not since the price returns to a pre-existing level.

We can also use this simple setup to ask the following question. What frequency of price changes should a research calibrate a simple Calvo model with no temporary price changes in order to predict the real effects from money of the model with α_R fraction of permanent and α_T fraction of temporary price changes? Using the results above, this is the frequency of price changes, α, that equates

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

We thus have the following corollary to Proposition 2.

Corollary: If the data is generated by a model permanent and permanent price changes, α_R and α_T, then a simple Calvo model that predicts the same real effects of money is one with

$$(13) \quad 1-\alpha = \frac{1-\alpha_R-\alpha_T}{[1-(1-\alpha_R)^2+(1-\alpha_R-\alpha_T)^2]^{\frac{1}{2}}}.$$

3. A Model of Temporary and Permanent Price Changes

The recent debate in the literature between Golosov and Lucas (2007) and Midrigan (2007) has focused on how good an approximation is a simple Calvo model of price changes to a menu cost model. That literature presumes that the true data generating process is a menu cost model and that a researcher, for simplicity's sake, fits a simple Calvo model. Golosov and Lucas found that such a researcher would overstate the real effect of money by about 500%. 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 by only 25%.

We take up an unexplored aspect of this general debate. We ask, suppose the true data generating process is a menu cost model with temporary and permanent price changes and a researcher fits the data with a simple Calvo model, would the researcher overstate the real effects of money or understate them and by how much. The answer, as one would antici-

pate, depends on whether one follows the temporary-changes-out approach or the temporary-changes-in approach. We find that a researcher that followed the temporary changes out approach would overstate the real effects of money by almost 70% and the a researcher that followed the temporary-changes-in approach would understate them and find only about 1/9th as large real effects as there are in the data generating process.

The model we use is a parsimonious extension of a standard menu cost model. Indeed, it adds only one parameter on the firm side relative to the standard model. In our model, as in the standard menu cost model, 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. 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) only has technology shocks. We allow for both technology shocks and a demand shocks. Our motivation is both from theory and from the 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 when a large number of buyers of the product happen to have high elasticities.

Our motivation from the data comes from two observations. First, as we have shown quantities seems to be more sensitive to price changes in periods of temporary price declines than in other periods. Second, as several authors have shown temporary price cuts are associated with reductions in price-cost margins. (See, for example, Chevalier, Kashyap, and Rossi 2003.) Taken together these features suggest that in the data the demand elasticity faced by firms are time-varying and this feature leads them to have time-varying markups. Motivated by theory and data we introduce time-varying elasticities by having consumers with differing demand elasticities and have good-specific shocks to preferences.

We argue that our model is a useful laboratory for evaluating the common approaches

14

by showing that it can fit what we feel are the key aspects of the micro data. In this sense, we follow much of the spirit of the recent debate.

A. Setup

Formally, we study a a monetary economy populated by a large number of identical, infinitely lived consumers, firms, and a government. In each period t, the 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 zero, 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 growth and idiosyncratic shocks to the productivity and demand for each good. In terms of the money growth shocks we assume that the (log of) money growth follow an autoregressive process of the form

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

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.

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

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

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

The economy has two types of consumers: measure $1 - \omega$ of *low elasticity* consumers and measure ω of *high elasticity consumers*. (CHANGE TABLES) The stand-in consumer for the low elasticity consumers, consumer of type A, has preferences of the form

$$(16) \quad \sum \beta^t \pi(s^t)[\log c_A(s^t) - \psi l_A(s^t)]$$

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

$$(17) \quad \sum \beta^t \pi(s^t)[\log c_B(s^t) - \psi l_B(s^t)]$$

where $c_B(s^t)$ is a composite of goods given by $\left(\int_0^1 z_i(s^t)^{\frac{1}{\gamma}} c_{Bi}(s^t)^{\frac{\gamma-1}{\gamma}} di \right)^{\frac{\gamma}{\gamma-1}}$ and $l_B(s^t)$ is labor supplied by this consumer and $z_i(s^t)$ are a type of preference shocks 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 are a continuum of goods. Note also that on the labor side we follow Hansen (1985) by assuming indivisible labor decisions 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 household of type A. The purchases of goods by this household must satisfy the following cash-in-advance constraint:

$$\int p_i(s^t) c_{Ai}(s^t) di \leq M(s^t).$$

16

The budget constraint of this household is

$$(18) \quad M(s^t) + \sum_{s_{t+1}} Q\left(s^{t+1}|s^t\right) B(s^{t+1})$$

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

where $1/R(s^t) = \sum_{s_{t+1}} Q\left(s^{t+1}|s^t\right)$ is the uncontingent nominal interest rates, $M_t(s^t)$ is nominal money balances, $W(s^t)$ is the nominal wage rate, $l(s^t)$ is labor supplied, and $T\left(s^t\right)$ is transfers of currency and $\Pi(s^t)$ are profits. The left side of (18) is the nominal value of assets held at the end of bond market trading. The right hand side terms 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 currency, and the profits from the firms. The cash-in-advance constraint and the budget constraint for consumers of type B is analogou.

Notice that in (18) 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 recent literature on sticky prices so we abstract from them here as well to retain comparability.

It is convenient to solve the households' problem in two stages. 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 $P_A(s^t) = \left(\int_0^1 p_i^{1-\theta}(s^t)di\right)^{\frac{1}{1-\theta}}$. We solve an analogous problem for the composite demand $c_B(s^t) = \left(\int_0^1 z_i(s^t)^{\frac{1}{\gamma}}c_{Bi}(s^t)^{\frac{\gamma-1}{\gamma}} di\right)^{\frac{\gamma}{\gamma-1}}$ and define the resulting price index $P_B(s^t) = \left(\int_0^1 z_i(s^t)p_i^{1-\gamma}(s^t)di\right)^{\frac{1}{1-\gamma}}$.

17

The resulting total demand for good i is given by

$$(19) \quad q_i(s^t) = (1 - \omega) \left(\frac{P_i(s^t)}{P_A(s^t)} \right)^{-\theta} c_A(s^t) + \omega \left(\frac{P_i(s^t)}{P_B(s^t)} \right)^{-\gamma} z_i(s^t) c_B(s^t)$$

Notice that (19) makes clear the precise sense in which the shocks $z_i(s^t)$ represent a type of demand shock: when $z_i(s^t)$ is relative 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 (19) also makes clear that our model generates time-varying elasticties of demand in a simple way. In periods when z_i is relatively high, a high fraction of goods are demanded by consumers with a high demand elasticity (γ) and when z_i is relatively low, a high fraction of goods are demanded by consumers with a low demand elasticity (θ).

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

Firms

Consider now the problem of a firm. 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 for that one period t to charge a price different from its inherited regular price $P_R(s^{t-1})$. If the firm wants to continue that temporary price change into the next period, it must again pay ϕ. In period after the period of temporary price changes ends, the firm inherits the existing regular price $P_R(s^{t-1})$.

In our simple model, the only role of temporary price changes is to economize on the costs of changing prices. In our model 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 their regular prices.

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

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

where we have used the demand function (19). The present discounted value of profits of the firm, expressed in units of period 0 money, is given by

$$(20) \quad \sum_t \sum_{s^t} Q(s^t)[R_i(P_i(s^t); s^t) - W(s^t)(\kappa\delta_{R,i}(s^t) + \phi\delta_{T,i}(s^t))].$$

where the variable $\delta_{R,i}(s^t)$ is an indicator variable that equals one when the firm changes its regular price and zero otherwise, while $\delta_{T,i}(s^t)$ is an indicator variable that equals one when a firm has a temporary price change and is zero otherwise. In expression (20), the term

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

is the labor cost of changing prices. 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 is there is neither a regular price change 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 the market-clearing conditions and the definition of *equilibrium*. The market-clearing condition on labor,

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

requires that the labor used in production as well as the menu costs (measured in units of labor) of making both regular price changes and temporary changes adds up to 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 $\{P_i(s^t), y_i(s^t)\}_i$; 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.

It will be convenient to 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 regular price in the last period, $P_{Ri}(s^{t-1})$, its idiosyncratic productivity level, $a_i(s^t)$ and the idiosyncratic taste for its good, $z_i(s^t)$ It is convenient to normalize all of the nominal prices and wages by the current money supply. 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 $[p_{R-1,i}(s^t), z_i(s^t), a_i(s^t)]$.

Let $\lambda(s^t)$ denote the measure over firms of these state variables. Since the only aggregate uncertainty is money growth and the process for money growth is autoregressive, it follows that 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

$$(21) \quad R(p_i, a, z, S) = \left(p_i - \frac{w(S)}{a} \right) q(p_i, z, S),$$

where $w(S), p(S)$, and $q(p_i, \cdot, S)$ are all known functions of the aggregate state. The function 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.

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 actuarily-fair manner. Of course, our formalization is equivalent to having an intertemporal price

defined over idiosyncratic shocks and aggregate shocks and then simply summing over both idiosyncratic shocks and aggregate shocks.)

The value of a firm that charges a temporary price $p_T \neq p_{R,-1}$

(22)

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

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

(23)

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

An intuitive way to think about the difference between a temporary and a regular price change is as follows. A temporary price change corresponds to renting a new price for today for one period, while a regular price change corresponds to buying a new price that can be used for a number of periods 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 (22) makes it 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

(24) $\quad p_T = \dfrac{\varepsilon(p, z, S)}{\varepsilon(p, z, S) - 1} \dfrac{1}{a} w(S).$

where $\varepsilon(p, z, S)$ is the demand elasticity of $q(p, z, S)$ derived from (19). Note that this price is a simple markup over the nominal unit cost of production and that this price is exactly what a flexible price firm would charge when faced with such a unit cost. In contrast, conditional on changing the regular price, the optimal pricing decision for the new regular price, p_R, is dynamic. (In particular, p_R will not typically equal p_T. Note that this feature of our quantitative model differs from the corresponding one in our motivating example)

21

As (24) makes clear, conditional on having a temporary price change, the inherited regular price $p_{R,-1}$ is irrelevant, so we can write $p_T(a, S)$. Likewise, as inspection of (23) 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

In this section, we describe how we choose functional forms and benchmark parameter values for our model. We then investigate whether or not 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 real quantitative reactions to a monetary shock, which we will later use as a benchmark for judging other models.

Functional Forms and Parameters

We set the length of the period to one week and therefore choose a discount factor of $\beta = .96^{1/52}$. We choose ψ to ensure that in the absence of aggregate shocks, consumers supply one-third of their time to the labor market. We set γ, the elasticity of subsitution for the high elasticity types, to be 6. This is at the high end of the substitution elasticities estimated in grocery stores.

Our model is weekly, so the process for money growth (14) in our numerical experiments is weekly as well. Given that the highest frequency at which the U.S. Federal Reserve's monetary aggregate data are available is monthly, we pin down the model's serial correlation 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 serial correlation and variance as the Fed's measure of currency and checking accounts (M1) during 1989–97, the years for which the micro-price data used to calibrate the model are available.

The rest of the parameters are calibrated so that the model can closely reproduce the facts described earlier: κ, the cost the firm incurs when changing its regular price; ϕ, the cost of having a temporary sale; as well as the specification of the productivity shocks and the demand shocks.

Consider first the productivity process. As (15) indicates this process has persistence

ρ_a. The distribution of the innovations $\varepsilon_i(s^t)$ requires special attention. Midrigan (2006) shows that when $\varepsilon_i(s^t)$ is normally distributed, the model generates counterfactually low dispersion in the size of price changes. He argues that a fat-tailed distribution is necessary 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 (2006) 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 that has 3 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. For example, Golosov and Lucas (2007) show that monetary shocks are approximately neutral when productivity shocks are normally distributed. But as Midrigan (2006) shows, with a fat-tailed distribution of productivity shocks, changes in money have much larger real effects. The reason is that as the kurtosis of the distribution of productivity shocks increases, changes in the identity of adjusting firms are muted.

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 $z_t \in \{z_l, z_m, z_h\}$, with a transition probabilities

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

Hence ρ_v is the probability of staying in median 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 median demand state to the low demand state. We normalize $z_l = 0$. Our parameterization of these shocks has 5 parameters $\{z_m, z_h, \lambda_s, \lambda_l, \rho_v\}$.

Predictions

The Facts

We show now that our parsimonious model can be made to account for the five facts about prices we have documented. We then give some intuition for how the model works.

We ask, can the parameters governing the costs of changing prices and the productivity and demand shocks be jointly chosen to mimic well the patterns of prices and sales in the data as described by the facts? In setting these parameters, we target the 15 moments in the data indicated in Table 2 (which we have seen in Table 1). These moments include two on the frequency of price changes (including and excluding temporary price changes), the size and dispersion of price changes (including and excluding temporary price changes), the fraction of prices at the annual mode, the fraction of annual prices below the mode, the fraction of temporary price changes, the proportion of returns to the old regular price, the probability of a temporary price spell ending, as well as the fraction of periods and good sold in periods when prices are temporarily up and temporarily down.

In Table 2 we see that with a particular set of parameters, our parsimonious model does a remarkably good job at reproducing these facts. The prevalence of weekly price changes is high: .33 in the data and .31 in the model with all prices included (and much lower both in the data, .020, and the model, .019, when sale price changes are excluded). The mean size of price changes is high in both the data (.17 for all price changes and 0.11 for regular price changes) and the model (.16 and 0.11), and the dispersion is high in both as well. The proportion 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 spend more time below the annual mode than above it. Specifically, in the data, prices spend 30% of their time below the annual mode and in the model, about 28%. Most price changes are temporary: 94% in the data and in the model. Most temporary prices tend to return to the regular price

24

that existed before the temporary change: 80% do in the data and 90% do in the model. We also see that temporary price chagnes are very transitory: in the data, 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 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. Consider, first, the last fact from the data— namely, that price changes are clustered, in the sense that the hazard of price changes is sharply decreasing in the first two weeks after a price change. Figure 3 shows that our model generates a sharply declining hazard in the first two weeks, just as in the data. After those weeks, the hazard in the data continues to decline somewhat, while in our model the hazard is essentially flat.

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 the relative fraction of goods sold in periods with prices temporarily up and down. We see that for most of these, the model produces numbers similar to those in the data. Finally, as in the data, the 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 (-2.2) than in periods with temporary price cuts (-4.4).

Table 2 lists the parameter values that allow the model to best match the moments in the data. The menu cost of changing regular prices is 0.9% of a firm's steady-state labor expense. In contrast, the cost of a temporary price change is 0.44% of a firm's steady-state labor expense, or about 50% of the cost of changing the regular price. Productivity (or technology) shocks arrive with probability $\lambda = 0.061$ and have an upper bound of $\bar{\nu} = 0.095$. Moreover, the productivity process is highly transitory; its persistence is $\rho_a = 0.99$. The faction of high-elasticity consumers, ω, is 0.08. The distribution of taste shocks is $\{z_l, z_m, z_h\} = \{0, 0.047, 0.197\}$, and the parameters governing the Markov transition matrix

are $\lambda_s = 0.369, \rho_v = 0.803$, and $\lambda_l = 0.072$. Thus the median preference state is most persistent, while firms that are in the low or high preference states expect to return to the median with high probability $1 - \lambda_s = 0.631$. Finally, the elasticity of type A consumers is equal to $\theta = 1.984$.

Consider our model's prediction for the real effects of a monetary shock. Our summary measure of the real effect is the standard deviation (or volatility) of output, which we see from Table 4 is .72%. In this model, if a monetary shock has no real effect, then this standard deviation should be zero; and the larger are the real effects, the larger is the standard deviation. We find this number useful as setting the benchmark against which to compare the relative sizes of the real effects of monetary shocks in alternative versions of the model.

The Workings of the Model

Now consider the firm's optimal decision rules. These rules are a function of the individual states, namely the normalized regular price $p_{R-1} = P_{R,-1}/M$, and the current productivity level a and the current demand shock z, as well as the aggregate state variable, namely the money supply growth rate and the distribution of firms λ which we treated in our approximation as a constant.

We illustrate the optimal decision rules in Figure 4. Since the demand shock takes on three values we report the firm's decision rules for each of the three demand states. In particular, Figure 4 shows for each of the three demand states (low demand z_l, medium demand z_m, and high demand z_h) the regular price $p_R(a)$ the firm chooses conditional on the firm choosing to change the regular price and the temporary price the firm chooses $p_T(a)$ conditional on changing choosing to set a temporary price. Figure 4 also shows the regions of the state space in which the firm optimally chooses to have a regular price change (R), to have a temporary price change (T), or to do nothing (N).

As we have noted above when discussing (22) the price $p_T(a)$ is a constant markup over marginal cost given by (24), and it does not equal the 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, $p_R(a)$ differs from $p_T(a)$ because its choice reflects the dynamic considerations in (23),

Figure 4 also illustrates the regions of the state space for which the firm finds it optimal to have to change its regular price, have a temporary price change, or sell at the existing regular price (do nothing). All three panels have the standard feature that if the current regular price $p_{R,-1}$ is close enough to both $p_R(a)$ and $p_T(a)$ in the sense that it lies in the regions labelled N the firm finds it optimal to forgo paying any costs and just charging the regular price.

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 it always chooses a new regular price. The reason is the productivity and monetary shocks are very persistent so that the firm expects its new regular price to be close to what is statically optimal for a long period of time. Hence, it makes sense for the firm to pay the large fixed costs to change its regular price.

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

The firm's optimal decision rules in the low demand state are somewhat more subtle. The key difference between the low demand state and the high demand state is that the low demand state is fairly close to the medium state while the high demand state is much further (in the sense that z_l is only about 5 percent lower than z_m but z_h is about 20 percent higher than z_m.) If the firm's existing regular price is very far from what is currently statically optimal the firm changes it 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 the temporary state of low demand will pass quickly and makes a once-and-for-all adjustment to have a new price which 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

27

region the firm decides to have a temporary price increase today. The final subtle part is that when demand is low the costs of the having a price that differs from statically optimal one is lower than when the demand state is medium, because the lost profits or low when demand is low. Hence, the inaction region in the low demand state is somewhat wider.

4. Experiments

We have shown that our menu cost model with permanent and temporary price changes can reproduce the main features of the micro price data. Thus we view our model as an interesting laboratory to evaluate the two common approaches to dealing with temporary price changes. We focus primarily on the common approaches that use the simple Calvo model of pricing with only permanent price changes since it is most popular in the applied literature and it is viewed as an simple approximation to an underlying menu cost model. We find the same qualitative results with the Calvo model as in our previous comparison: the temporary-changes-in version understates the real effects of monetary policy, while the temporary-changes-out version overstates them. The temporary changes in approach produces real effects only about 20% as large as those in the menu cost model while the temporary changes out approach produces real effects that almost 70% larger than those of the menu cost model. We then ask: How would the duration of prices in the simple Calvo model have to be set so as to reproduce the same size real effects as in our menu cost model? We find the answer is 16.7 weeks.

In Appendix we also evaluated existing shortcuts that use a menu cost model without a motive for temporary price changes.

A. The Two Common Approaches with Calvo Models

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 the model are identical to those in the benchmark model. Firms are allowed to adjust their prices in an exogenous, costless, and random fashion as in the motivating example 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*.

The problem of a firm that is allowed to change prices in state s^t is

$$\max_{p_i(s^t)} \sum_{r=t}^{\infty} \sum_{s^r} (1-\alpha)^{r-t} Q(s^r|s^t) R\big(P_i(s^t); s^r\big) \left[P(s^t) c_{Hi}(s^r) - \frac{W(s^r)}{a(s^r)} c_{Hi}(s^r) \right],$$

where

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

Since there are no costs to changing prices, the resource constraint is simply

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

In the Calvo's models, the parameters of technology, preferences, and stochastic processes are set to be equal to those in our benchmark model. (See Table 3.) The additional parameter that needs to be set is α. We consider two parameterizations corresponding to the two alternative practices discussed above.

In the *temporary-changes-out* approach, we filter the data using the same algorithm as before 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 out model we set $\alpha = .33$ so that the average duration of prices is 16.7 weeks. We leave all other parameters unchanged.

In Table xx we report on our measure of the size of the measure of the real effects of money in the Calvo models relative to our menu cost model. We see that the standard deviation of output in the temporary-changes-out approach is 68% higher than in the menu cost model and the standard deviation of output in the temporary changes in approach is only about 1/9th of that in the menu cost model.

We now turn to a figure which gives some feel of how a researcher would have to adjust the duration of prices in the simple Calvo model to mimic the real effects of the menu cost

model. In Figure ?? we plot the standard deviation of output in a 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 similar.

5. Conclusion

In the data, a sizable fraction of price changes are related to temporary price changes. Existing sticky price models abstract from explicitly modeling these changes. Should they? We argue that they should not. We have argued that neither of the existing approaches provides a good approximation to one with an explicit motive for temporary price changes. The temporary-changes-out approach leads to much larger effects, and the temporary-changes-in approach leads to much smaller effects.

We have offered two alternatives to the existing approaches. One alternative is to use our simple extension of the price-setting technology in Golosov and Lucas model that explicitly includes a motive for temporary price changes. We have shown that this parsimonious extension of the standard menu cost model of sticky prices can be made to well account for many of the patterns of prices in the data. A cruder but simpler alternative is to stick with a simple Calvo model and use our analysis to set the frequency of price adjustment so as to mimic the real effects in a model with temporary and permanent changes.

Notes

[1]The data used by these Bils and Klenow (2004) and Nakamura and Steinsson (2008) has a much wider set of products than the grocery store data, but it is only collected as point-in-time data at the monthly frequency. Hence, this data gives 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 weekly vs. monthly data makes note that in the weekly Dominick's data prices change every three weeks and in the monthly data prices change every 9 weeks.

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

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, Anil Kashyap and Peter Rossi. 2003. Why don't prices rise during periods of peak demand? Evidence from scanner data. *American Economic Review* 93(1): 15–37.

Gertler, Mark, and John Leahy. 2006. A Phillips curve with an Ss foundation. Working Paper 11971, National Bureau of Economic Research.

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

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

Midrigan, Virgiliu. 2006. Menu costs, multi-product firms, and aggregate fluctuations. Manuscript, Federal Reserve Bank of Minneapolis.

Nakamura, Emi, and Jón Steinsson. 2007. Five facts about prices: A reevaluation of menu cost models. Manuscript, Harvard University.

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

6. Appendix: Algorithm to Construct the Regular Price

Our algorithm is based on the idea that a price is a *regular price* if the store charges it frequently in a window adjacent to that observation. For each period, compute the mode of prices p_t^M in a window that includes the previous 5 prices, the current price, and the next 5 periods[3]. Given the modal price in this window, the regular price is constructed recursively as follows. Set the initial period's 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, set the regular price equal to the modal price. Otherwise, set the regular price equal to last period's regular price. Finally, we would like to eliminate regular price changes that occur in the absence of changes in the store's actual price if the actual and regular price coincide in the period before or after the regular price change. Thus, if the initial algorithm generates a path for regular prices such that a change in the regular price occurs in the absence of a change in the actual price, we replace the last period's regular price with current period's actual price if the regular and actual prices coincide in the current period. 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. Figures 1-4 present examples of regular price series constructed using this algorithm.

So far we have described our algorithm intuitively. Here we provide the precise algorithm we used to compute the regular price.

1. Choose parameters: $l = 5$ (size of window: # weeks before/after current period used to compute modal price), $c = 1/3$ (cutoff used to determine whether a price is temporary), $a = .5$ (# periods in window with available price required in order to compute a modal price).

We apply this algorithm below for each good separately.

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

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

If # periods with available data in $(t - l, ...t + l)$ is $\geq 2al$,

    let $p_t^M = mode(p_{t-l}, ...p_{t+l})$

    let f_t = fraction of periods (with available data) in this window s.t. $p_t = p_t^M$.

Else, set $f_t, p_t^M = 0(missing)$

3 Define p_t^R using the following recursive algorithm.

If $p_{1+l}^M \neq 0$, 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 $(p_t^M \neq 0$ & $f_t > c$ & $p_t = p_t^M)$, set $p_t^R = p_t^M$

        else, set $p_t^R = p_{t-1}^R$

4. Repeat the following algorithm 5 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 set of periods with regular price change.

Let $\mathcal{C} = \{t : p_t^R = p_t$ & $p_t^R \neq 0$ & $p_t \neq 0\}$ be set of periods in which store charges 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 set of periods in which 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 $p_{\{\mathcal{R}\cap\mathcal{P}\}}^R = p_{\{\mathcal{R}\cap\mathcal{P}\}-1}$.

Appendix B

ince Calvo models are by far the most popular in applied work we have focused our attention on them. For completeness we also performed analogous experiments comparing how the real effects in simple menu cost models with no mechanism for temporary price cuts under the two common approaches compare to the real effects in our menu cost model.

The results we found were qualtitatively similar to those for the Calvo model and are reported in Tables A1and A2. Table A1 reports our choice of parameter values in the menu cost models without a motive for price changes. These models abstract from taste shocks and assume that the measure of type B consumers is constant at 0. We only calibrate the frequency and size of price changes by choosing the arrival rate of technology shocks and the upper bound of the support of their distribution. Table A2 illustrates the real effects of money in these economies. We find again that the temporary-changes-out approach leads to real effects overstates the real effects of money: now by about 40%. Similarly, the temporary-changes-in approach understates the real effects of money: it predicts real effects which are about 1/5th of those in the benchmark model.

In Table A1 we see that one discrepancy between the models without a motive for temporary price changes and the data (and thus our benchmark model) is that they miss the fraction of prices at the annual mode. The temporary-changes-in approach generates too few prices at the annual mode: 22%; the temporary-changes-out approach generates too many prices at the annual mode: 77% In contrast, in the data the fraction of prices at the annual mode is 58%.

We next show that an alternative superior approach to dealing with temporary price changes is to choose parameters in models without a motive for temporary price changes by matching the fraction of annual price at the mode. When we do so, the implied frequency of price changes reported in the last column of table A1 is 0.051, or about once every 20 weeks. In Table A2 we show that this alternative parametrization predicts real effects of money very similar to those of the Benchmark setup: it overstates the real effects by only 8%.

Table 1: Facts about prices

	Moments	Data		Benchmark model		
		Include temporary changes	Exclude temporary changes	Include temporary changes	Exclude temporary changes	Used for calibration?
Fact 1	Frequency of price changes	0.33	0.020	0.31	0.019	√
	Fraction of temporary price changes	0.94		0.94		√
	Proportion of returns to regular p	0.80		0.90		√
	Probability temporary price spell ends	0.46		0.59		√
Fact 2	Fraction pds with temp prices, %	24.3		18.4		√
	Fraction pds when p temp. up, %	2.1		2.1		
	Fraction pds when p temp. down, %	20.3		16.0		
Fact 3	Fraction of annual prices at mode	0.58		0.58		√
	Fraction of annual prices below mode	0.30		0.28		√
Fact 4	Mean size of price changes	0.17	0.11	0.16	0.11	√
	Mean $\log(p_T/p_R)$ if p. temporarily down	-0.22		-0.23		√
	Mean $\log(p_T/p_R)$ if p. temporarily up	0.13		0.08		√
	IQR of price changes	0.15	0.08	0.17	0.04	√
	IQR $\log(p_T/p_R)$ if p. temporarily down	0.21		0.13		
	IQR $\log(p_T/p_R)$ if p. temporarily up	0.12		0.01		
Fact 5	Fraction output when temp prices, %	38.0		34.5		√
	Fraction output when p temp. up, %	1.2		1.7		
	Fraction output when p temp. down, %	35.4		32.6		
	Price elasticity for regular price changes	-2.08		-2.19		
	Price elasticity for temporary price cuts	-2.93		-4.40		

Table 2: Parameter values

	Benchmark Model	Calvo model without temporary prices	
		Temporary-changes-out	Temporary-changes-in
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 technology shock	0.061		
Upper bound on technology	0.095		
Persistence of technology	0.991		
Median demand state	0.047		
Upper demand state	0.197		
Substitution elasticity of type-A consumers	1.984		
Measure of high elasticity consumers	0.08		
Probability of staying in non-mean demand state	0.369		
Probability of staying in mean demand state	0.803		
Probability of jumping to low demand state	0.072		
Probability of Calvo price adjustment		0.02	0.33
Assigned			
Discount factor	$0.96^{1/52}$	$0.96^{1/52}$	$0.96^{1/52}$
Substitution elasticity of type-B 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 money

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

Table A1: Parameter values and moments in menu-cost model

	Data		Menu-cost model without temp prices		
	Include temporary changes	Exclude temporary changes	Temporary-changes-in	Temporary-changes-out	Match fraction of prices at annual mode
Moments					
Frequency of 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 technology shock			0.750	0.012	0.049
Upper bound on technology			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 technology			0.84	0.84	0.84
Discount factor			$0.96^{1/52}$	$0.96^{1/52}$	$0.96^{1/52}$
Substitution elasticity of type-A consumers			3	3	3
Substitution elasticity of type-B consumers			x	x	x
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 money in menu-cost models

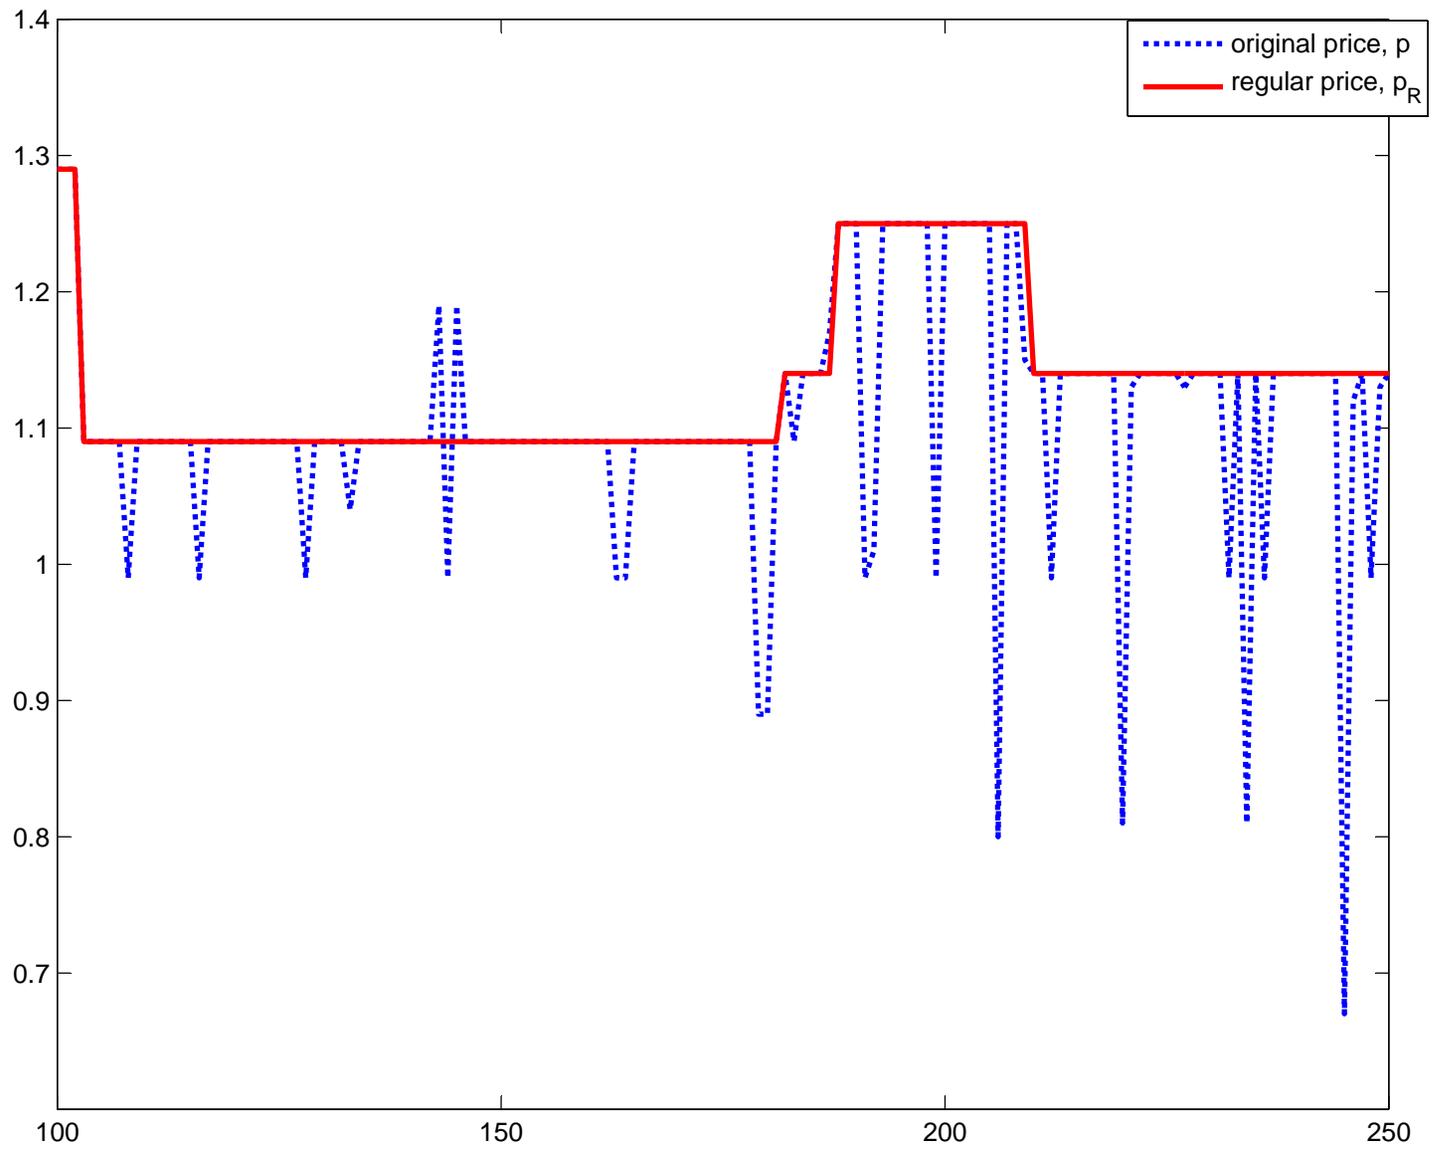|	Benchmark	Menu-cost model without temp prices		
		Temporary-changes-in	Temporary-changes-out	Match fraction of prices at annual mode
Std. dev. of chain-weighted real output	0.72	0.15	1.00	0.78
Relative to Benchmark model		0.21	1.39	1.08

Figure 1: Ilustrate temporary and permanent prices

Figure 2: Example of algorithm

Angel Soft Batt Tissue, 4 CT

GREEN APPLE GLY SP# Bath Soap

OMEGA II LAUNDRY DETERGENT, 25 LB

Sprite, 24/12

- - - - original price
——— regular price

Figure 3: Hazard of price changes, Dominick's data

Actual prices

probability of adjustment

age of price spell

Regular prices

probability of adjustment

age of price spell
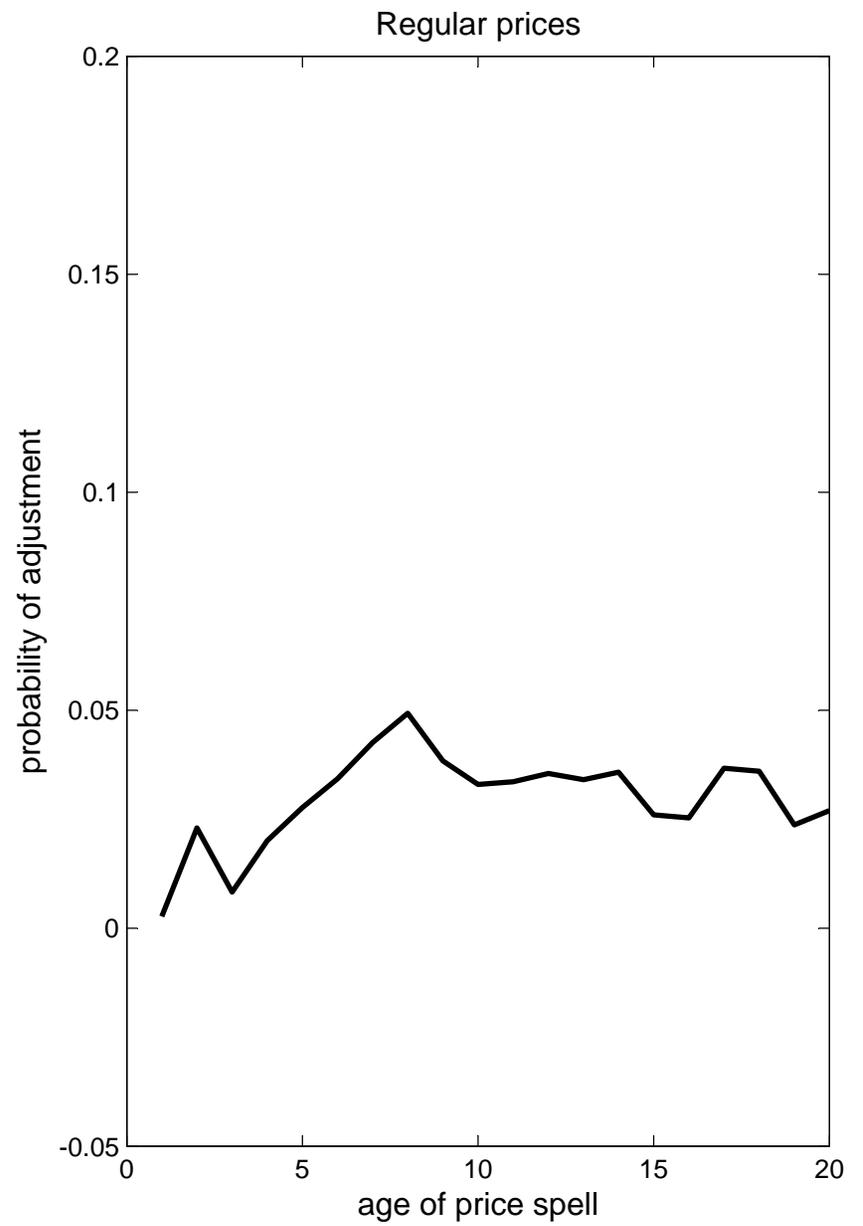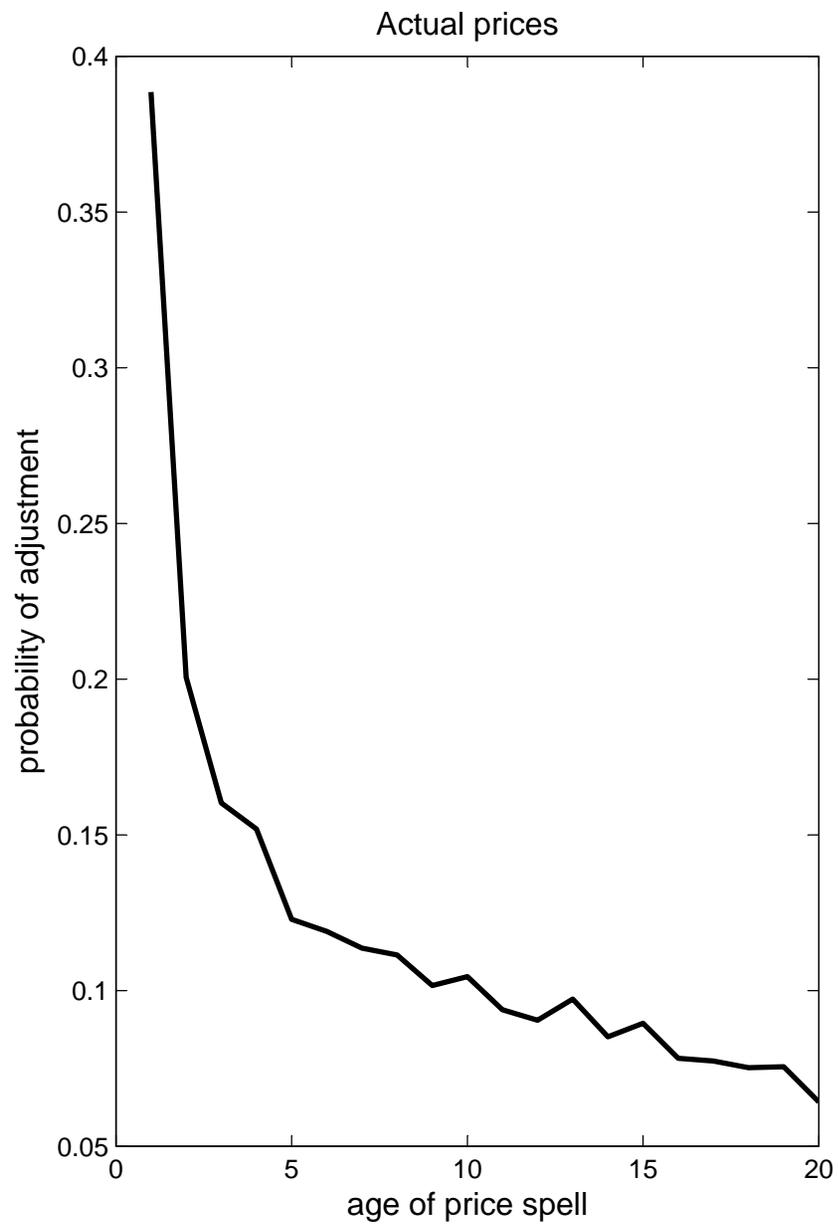
Figure 4: Hazard of price changes, benchmark model

Actual prices

Regular prices

Data
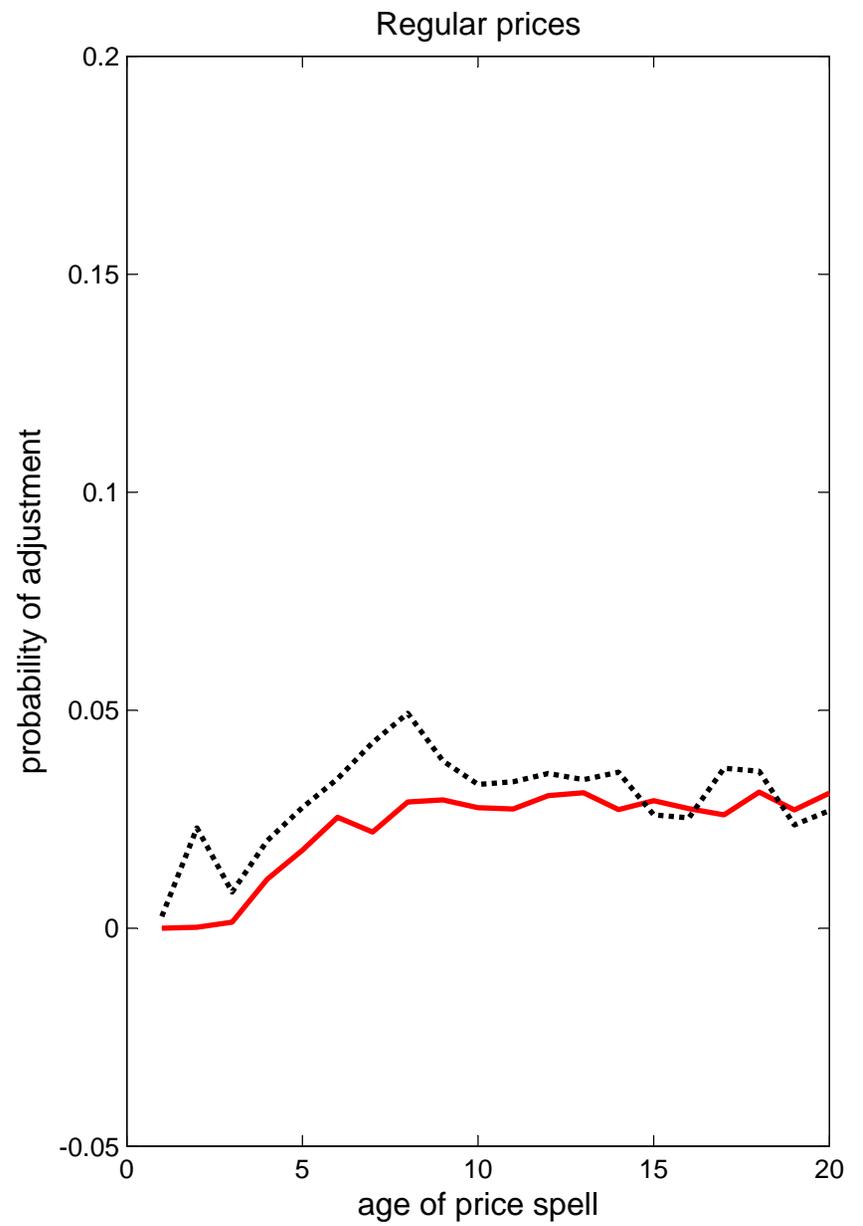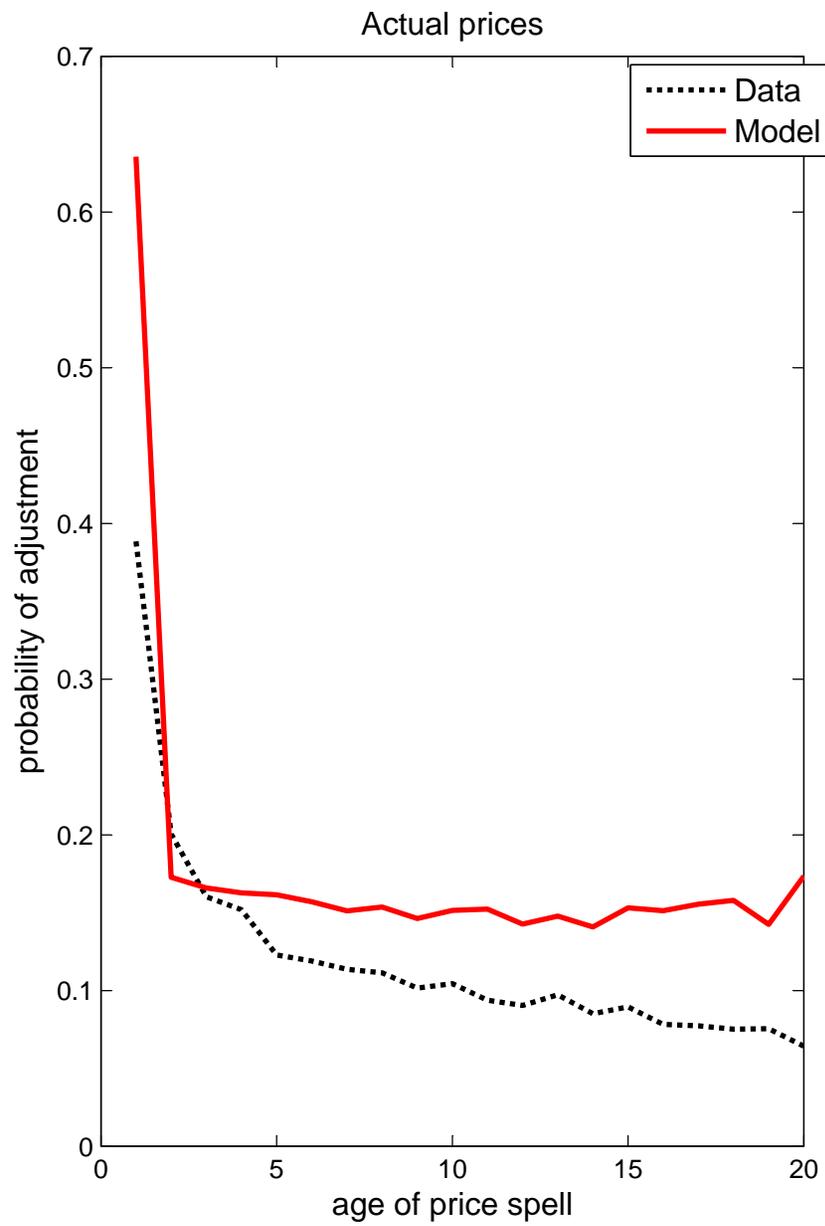Model

probability of adjustment

age of price spell
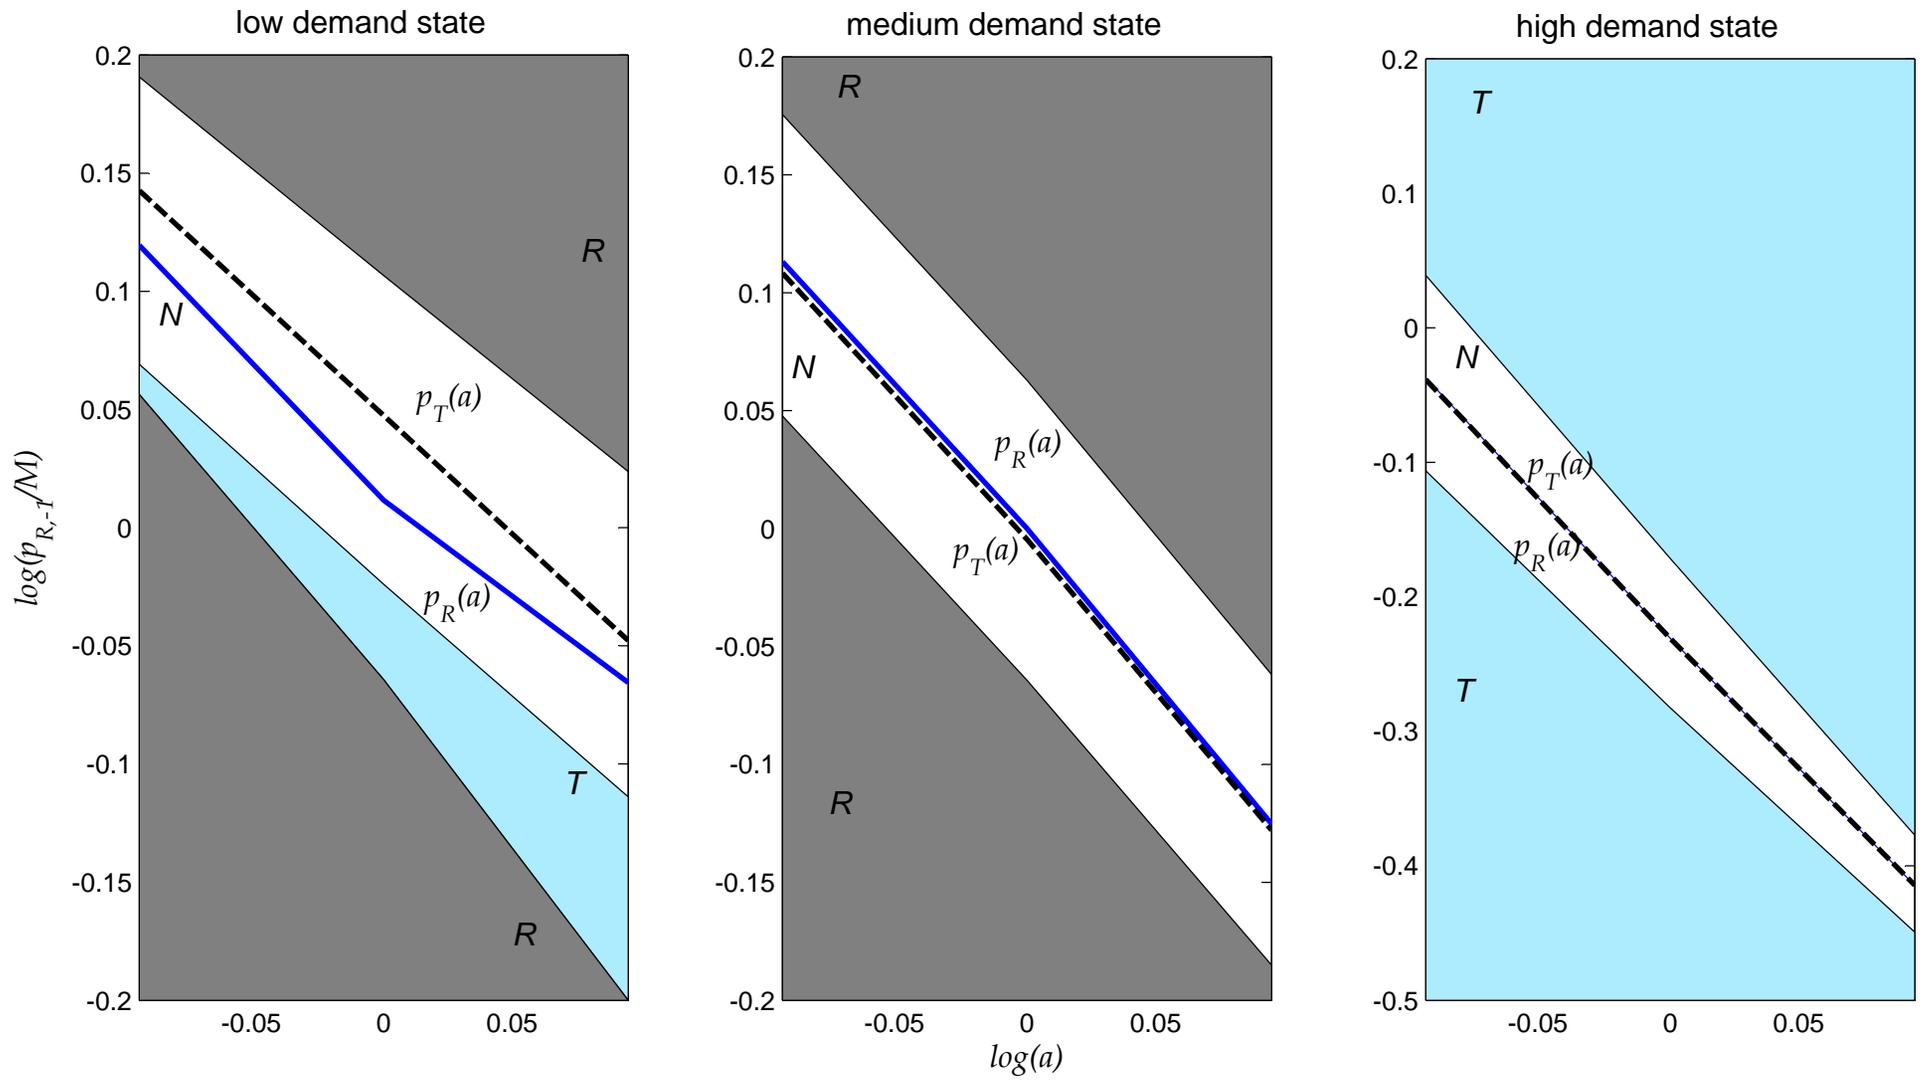
Figure 5: Optimal policy rules, benchmark model

Figure 6: Real effects of money in Calvo alternatives vs. Benchmark