Federal Reserve Bank of Minneapolis
Research Department

Sales and the
Real Effects of Monetary Policy[*]

Patrick J. Kehoe and Virgiliu Midrigan

Working Paper 652

April 2007

ABSTRACT

In the data, a sizable fraction of price changes are temporary price reductions referred to as sales. Existing models include no role for sales. Hence, when confronted with data in which a large fraction of price changes are sales related, the models must either exclude sales from the data or leave them in and implicitly treat sales like any other price change. When sales are included, prices change frequently and standard sticky price models with this high frequency of price changes predict small effects from money shocks. If sales are excluded, prices change much less frequently and a standard sticky price model with this low frequency of price changes predict much larger effects of money shocks. This paper adds a motive for sales in a parsimonious extension of existing sticky price models. We show that the model can account for most of the patterns of sales in the data. Using our model as the data generating process, we evaluate the existing approaches and find that neither well approximates the real effects of money in our economy in which sales are explicitly modeled.

At the heart of monetary policy analysis is the question, How large are the real effects of monetary shocks? The answer to this question has important implications for the optimal conduct of monetary policy, the linking of fluctuations in monetary aggregates to those in real activity, and debates regarding the relative potency of fiscal and monetary policy.

Recently there has been a vast surge in work to study this question. The most popular class of models used to quantify the real effects of money assume that goods prices are sticky. In these models, firms typically leave their prices unchanged for a number of periods and, hence, monetary shocks have real effects. A key ingredient in these models that determines the size of the real effects of money 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.

How frequently do prices change in the data? The answer to this question crucially depends on how sales are treated in the data. Bils and Klenow (2004) treat sales just like any other price change and find that prices change frequently: the median consumer good experiences a price change about every 4.3 months. This number suggests that prices change about three times as frequently as was previously believed. For example, in a survey of the empirical literature on the frequency of price changes, Taylor (1999) argues that the average frequency of price changes is about one year. Bils and Klenow (2004) and others have interpreted this high frequency of price changes as casting doubt on the relevance of price stickiness in accounting for business cycle fluctuations.

A recent study by Nakamura and Steinsson (2007) shows that Bils and Klenow's conclusion that prices are fairly flexible is due in large part to the fact that a large number of price changes in the data are sales. When Nakamura and Steinsson exclude sales from the

definition of price changes, they find frequencies close to those cited by Taylor, namely that the median frequency of price changes is about every 8 to 11 months.

In summary, then, if sales are included in the data, prices are fairly flexible; if they are excluded, prices are fairly sticky. To date, there are two approaches to deal with sales. The most popular approach is to exclude sales from the data, write down a model without sales, and then match the frequency of price changes to the data with sales excluded. We refer to this approach as the *take-sales-out* approach.

An alternative approach is to include sales in the data, write down a model without sales, and then match the frequency of price changes to the data with sales included. We refer to this approach as the *leave-sales-in* approach. The leave-sales-in approach obviously generates much smaller effects from monetary shocks than the take-sales-out approach. The leave-sales-in approach implicitly assumes that, in terms of evaluating the real effects of a money shock, a sale is just like any other price change. The take-sales-out approach, in contrast, implicitly assumes that, in terms of evaluating monetary shocks, a sale is similar to no price change. To date there seems to be little guidance from theory as to which approach is preferable.

This paper takes up the issue of how to deal with sales in the price data. In the data we use, sales are defined by a simple AC Nielsen algorithm that looks at the pattern of price changes and classifies price reductions as sales if they are reversed sufficiently quickly and classifies the rest as regular price changes.

Our approach differs from existing ones in that we explicitly include a motive for temporary price reductions in a simple sticky price model and then directly study how large are the real effects of monetary policy shocks. We then treat the model as the data-generating

2

process and apply the two common approaches to data generated from our model. We find neither of the existing approaches provides a good approximation to the key question of interest: the take-sales-out approach leads to much larger effects of monetary shocks than are in the model with sales, and the leave-sales-in approach leads to much smaller effects of monetary shocks than in the model with sales.

We then show that a simple rule of thumb approach of using a model without sales but choosing parameters, not to match the frequency of price changes, but rather to match the fraction of time a price stays at its annual model gives a much better approximation than either of the existing approaches. We end with two proposals to advance the sticky price literature: either explicitly include sales in the model or follow a version of our rule of thumb. We argue that either will represent progress relative to the existing approaches.

Our model is purposefully chosen to be an exceptionally simple and parsimonious extension of the existing sticky price literature: we add one parameter to existing menu cost models, the cost of having a one-period *markdown*. We show that even though the model is simple, it can capture many of the features in the data concerning sales.

In the model a key state variable of each firm is its *regular price* (or reference price) inherited from the previous period. This price is the price it can charge in the current period with no extra costs. If it wants to charge a different price in the current period, it has two options: change its regular price or have a one-period markdown (a "sale"). To change its regular price, the firm pays a fixed 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 one-period markdown, the firm pays a smaller fixed cost which gives it the right to charge a price lower than the existing regular price for the

current period only and keep its regular price unchanged.

The costs of various actions in the future depend on the actions taken in the current period. If in the current period the firm chooses to change its regular price, then in the next period it inherits this new regular price and decides whether to charge this regular price again at no cost, pay a fixed cost and change this regular price, or pay the smaller fixed cost and have a one-period markdown. If in the current period the firm chooses to have a one-period markdown, then in the next period it inherits the unchanged regular price and decides whether to charge this existing regular price at no cost, pay a fixed cost and change its regular price, or pay the smaller fixed cost and again have a one-period markdown. The problem of the firm then proceeds recursively.

In terms of our data analysis, we use data from two sources. We primarily focus on scanner price data from grocery stores. An appealing feature of this data is that it is weekly data and there is independent evidence that pricing decisions are made at a weekly level. Hence, we are comfortable modeling these grocery stores as making weekly decisions on prices. We also do some experiments with the data used by Bils and Klenow (2004) and Steinsson and Nakamura (2007). That data is much more comprehensive than the grocery store data, but it is only collected as point-in-time data at the monthly frequency. Hence, it gives no direct evidence about what happens within a month, and one needs to make a variety of assumptions to come up with the frequency of price changes.

We focus on seven features of prices and sales in our grocery store data. First, prices change frequently. Second, during the year prices spend most of their time at their modal value. Third, prices are much more likely to be below their annual mode than above it. Fourth, most price changes are associated with sales. Fifth, after a sale, the price tends to

4

return to the pre-sale price. Sixth, sales are very transitory. Finally, price changes tend to be clustered in time for a given product.

We show that our simple model does a remarkably good job of generating these features. We then use our model to evaluate existing approaches to dealing with sales. To do so, we treat our model as the data-generating process. We can ask, Which of two alternative practices more closely reproduces the real effects of monetary policy in our model: leaving the sales out of the data or leaving the sales in the data?

We begin by comparing our model to two menu costs models without markdowns. In one version of the model, the *leave-sales-in* version, we set the parameters of the model to match statistics in the data in which sales are included. In the other, the *take-sales-out* version, we set parameters to match statistics in data in which sales are excluded. We find that the leave-sales-in version significantly understates the real effects of monetary policy relative to these effects in the model with markdowns and that the take-sales-out version somewhat overstates the real effects of monetary policy.

We then perform a similar comparison using the more popular Calvo model of pricing, which is the benchmark model in the sticky price literature. These Calvo models are typically viewed as approximations to the underlying menu cost models. We consider a leave-sales-in version and a take-sales-out version of a Calvo model and find the same qualitative results as in our previous comparison: the leave-sales-in version understates the real effects of money, while the take-sales-out version overstates it. At a quantitative level, the approximation error with these models is large. In the leave-sales-in version, the real effects of money, as measured by the standard deviation of consumption, is less than one-fifth of the level in the menu cost model with markdowns. In the take-sales-out version, the real effects of money are about

5

twice the level in the menu cost model with markdowns.

We then propose an alternative procedure to set the parameters of a model without markdowns in order to approximate the real effects of money in a model with markdowns. In this procedure we leave sales in the data, but instead of choosing parameters to match the frequency of price changes (along with other statistics), we choose parameters to match the fraction of prices at the annual mode (along with the same other statistics). We show that for either a menu cost model without markdowns or the Calvo model, this procedure implies real effects of money similar to those in the menu cost model with markdowns.

1. Some Facts about Prices

We begin by documenting seven facts about price changes in the 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 Chicago GSB. The data consists 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. The products available in this database range from non-perishable foodstuffs (some of which are represented by the categories frozen and canned food, cookies, crackers, juices, sodas, beer), to various household supplies (some of which are represented by the categories, detergents, softeners, and bathroom tissue), as well as pharmaceutical and hygienic products. In addition to price data, the database maintained by the Chicago GSB provides information about the number of units of the good sold each period, the average acquisition cost of the goods in each store's inventory, as well as an indicator variable that records sales.

6

Because data are recorded by scanners, price series are frequently interrupted by gaps (periods when information about the price is missing, either because no consumers have purchased that good in that period or because the store has stocked out). To partly address this issue, we study, for each product (identified by its UPC code), the time series of prices of the store at which the product was most frequently available. Given that Dominick's sets prices on a chain-wide basis, price changes across stores are highly correlated, especially those in one of Dominick's three pricing zones (high, low, medium), and little information is lost by restricting our analysis to the price of a single store for each good.[1] Further, we restrict our attention to those goods for which at least 50 weekly price observations are available out of the total maximum of 400 weeks of the sample.

We use an algorithm to identify sales that is a minor extension of the algorithm described in the AC NIELSEN-ERIM database.[2] Here we discuss the algorithm presuming that no data are missing and discuss the general case in the appendix.

Sales are defined relative to an artificial series called a *regular* price series, denoted $\{P_t^R\}$, which is used mainly to define which periods are sales periods. A period is a *sales* period if in that period the original price P_t is lower than the regular price P_t^R. The regular price series is constructed from the original price series according to the following recursive algorithm. For each price cut, defined as a period t in which $P_t < P_{t-1}$, check if the price rises above the lower price within 5 weeks; that is, check if $P_{t+j} \geqslant P_t$ for $j \leqslant 5$. If it does, then let $\bar{j}$ be the first time the price rises above the lower price P_t and replace $P_t, P_{t+1,..}; P_{t+\bar{j}-1}$ with P_{t-1}. If the price never rises above P_t in the next 5 weeks, then leave P_t unchanged. Using the new series, repeat this algorithm 4 more times. The resulting series is the regular price series.

The rationale for repeating the algorithm several times is to allow sales periods to include a whole variety of patterns in which, in the middle of the sale, the price cut is partially reversed or even drops more during a sale. For a first example, suppose the original data are 200, 100, 101, 102, 103, 104, 105, 106. In the first round of the algorithm, the candidate for the regular price series is 200, 200, 101, 102, 103, 104, 105, 106; in the second round it is 200, 200, 200, 102, 103, 104, 105, 106, and so on. The regular price is defined to be the result of this algorithm after 5 such rounds and is given by 200, 200, 200, 200, 200, 200, 200, 106. Comparing the original series to this regular price series, we see that all periods but the first and the last are considered sales periods.

We emphasize that in our definition of sales, we do not require that the price drop by at least some minimum amount or that prices return to at least the pre-sale price. Notice also that we restrict our attention to temporary price cuts and thus exclude clearance sales (price decreases that are never at least partially reversed). As Nakamura and Steinsson (2007) report, clearance sales are uncommon in the processed food industry and constitute a sizable fraction of sales only for apparel products, which are absent in the Dominick's sample. Finally, in addition to using a 5-week duration for sales, we experimented with other durations and report statistics for a 3-week duration as well.

For illustration purposes, in Figure 1 we graph an example of a price series, namely the cost of a six-pack of Diet A&W Cream Soda. The dashed lines are the original transaction prices while the solid line gives the regular price series constructed with the AC Nielsen algorithm. For this series all periods in which the dashed line does not equal the solid line are defined to be sales. The figure makes clear that excluding sales (by using the regular price series rather than the original series) excludes most of the price changes in the original

series. We also see that price changes are large and the price changes tend to spend a lot of time at only a couple of values. As we shall see, these features are common for the prices in our data set.

To that end consider Table 1 in which we report a variety of facts about prices. For each of the 29 product categories we first compute category-level statistics by weighting each good by its sales share 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. (See the data appendix for details.) In discussing the facts about sales, we focus on those computed using the 5-week definition. Those computed using the 3-week definition are broadly similar. Several features of the data stand out.

Fact 1. Prices change frequently.

In Table 1 we see that 33% of prices tend to change every week. (As the second column shows, if sales-related price changes are excluded this number drops to 5.6%.)

Fact 2. Price changes are large and dispersed.

Table 1 shows that mean size of price changes is 17%. The smallest 25% of price changes are less than 4% and the largest 25% of price changes are over 19%. (As the second column shows, if sales-related price changes are excluded the mean size of price changes drops to 10%. Likewise, with sales-related changes excluded price changes are still dispersed, but somewhat less so: with sales-related changes excluded, the smallest 25% of price changes are less than 3% and the largest 25% of price changes are over 12%.)

Fact 3. During a year, prices spend most of their time at their modal value.

Table 1 also shows that, on average during a 50-week period, prices tend to be at their modal value 58% of the time. Prices tend to be at their two most widely used prices 76%

of the time and at one of their three most widely used values 84% of the time. Even in a 100-week interval, prices tend to be at one of their three most widely used values 78% of the time.

Fact 4. *Prices are much more likely to be below their annual mode than above it.*

Table 1 shows that prices are below their annual modal value 30% of the time and above it only 13% of the time. Thus, prices are about 2.5 times as likely to be below the annual mode than above it.

Fact 5. *Most price changes are associated with sales.*

From the lower panel in Table 1 we see that over 83% of price changes occur during a sales period. Thus, if we delete sales periods, then only 5.6% (which is 17% of 33%) of prices change every week, rather than 33% if we do not. Without sales, then, the implied duration of regular prices is about 4 1/2 months, which is much higher than the implied duration of 3 weeks using the prices including the original data with sales included.

Fact 6. *After a sale, the price tends to return to the pre-sale price.*

Table 1 also shows that over 87% of the time, the price after a sale tends to return to the pre-sale price.

Fact 7. *Sales are very transitory.*

Table 1 shows that the probability that a sales ends next week conditional on there being a sale this week is 46%. The implied duration of sales is about 2 weeks.

Fact 8. *Price changes are clustered.*

In Figure 2 we compute the hazard of price changes defined as the probability that prices change at $t+k$ given that the last price change occurred in period t. We computed this hazard as follows. For each of our goods we computed a good-specific hazard. We then took

a weighted average of each of these hazards using the good's share of total sales as the weight. Note that this procedure attempts to account for ex ante heterogeneity in the frequency of price changes across goods. In this figure we see that the hazard at one week is 53%. That is, conditional on a store changing the price of a given good 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 number of weeks since a price change. Notice that Figure 2 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.

2. A Model of Sales

Here we consider a parsimonious extension of a standard menu cost model and show that it can generate temporary price reductions of the kind classified as sales by the AC Nielsen algorithm. More generally, we show that this simple model can account well for the seven facts documented about the pattern of prices.

In our model, as in the standard menu cost model, firms can pay a fixed cost and change their regular price permanently. Our simple innovation is to allow firms the option in any period of paying different and smaller fixed cost and lowering their price for one period but leaving their regular (or reference price unchanged). At an intuitive level, we think of the standard model of requiring that the only way a price can change is that the firm buys a permanent price change. We think of our model as adding an option of renting a price change for one period.

Formally, our model represents a monetary economy populated by a large number of identical, infinitely lived consumers. 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 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. The only aggregate uncertainty is money growth, which is assumed to follow an autoregressive process of the form

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

where $\varepsilon(s^t)$ is a normally distributed i.i.d. random variable with mean 0 and standard deviation σ_μ.

Technology

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, and $a_i(s^t)$ is the good-specific productivity that evolves according to

$$\log a_i(s^t) = (1 - \rho_a)\bar{a} + \rho_a \log a_i(s^{t-1}) + \nu_i(s^t),$$

where $\nu_i(s^t)$ is a random variable.

Consumers

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.

12

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} at $t+1$. One unit of this bond pays one unit of currency 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 $Q(s^{t+1}|s^t) = Q(s^{t+1})/Q(s^t)$.

Consumers have a utility function of the form

$$\sum_t \sum_{s^t} \beta^t \pi(s^t) U(c(s^t), l(s^t)),$$

where $0 < \beta < 1$ is the discount factor, $l(s^t)$ is labor, and $c(s^t)$ is aggregate consumption good, which is a composite of a continuum of goods $c_i(s^t)$ given by

$$(2) \quad c(s^t) = \left(\int_0^1 c_i(s^t)^{\frac{\theta-1}{\theta}} di \right)^{\frac{\theta}{\theta-1}},$$

where θ is the elasticity of substitution across the different goods. The purchases of goods must satisfy a sequence of cash-in-advance constraints

$$(3) \quad P(s^t)c(s^t) \leq M(s^t),$$

where $P(s^t)$ is an aggregate price index given by

$$(4) \quad P(s^t) = \left(\int_0^1 P_i(s^t)^{1-\theta} di \right)^{\frac{1}{1-\theta}},$$

and $P_i(s^t)$ is the price of consumption good i.

Consumers face the following sequence of budget constraints in which the holdings of

real debt $B(s^t)/P(s^t)$ are bounded above and below by some arbitrarily large constants:

(5)
$$M(s^t) + \sum_{s_{t+1}} Q\left(s^{t+1}|s^t\right) B(s^{t+1})$$
$$= R(s^{t-1})W(s^{t-1})l(s^{t-1}) + B(s^t) + \left[M(s^{t-1}) - P(s^{t-1})c(s^{t-1})\right] + T(s^t),$$

where $1/R(s^t) = \sum_{s_{t+1}} Q\left(s^{t+1}|s^t\right)$, $M(s^t)$ is nominal money balances, $W(s^t)$ is the nomi-

nal wage rate, $l(s^t)$ is labor, and $T(s^t)$ is lump-sum transfers. The left side of the budget

constraint is the nominal value of assets held at the end of securities market trading. The

terms on the right side of the budget constraint are the value of nominal debt bought in the

preceding period, the shopper's unspent cash, as well as the transfers of currency and the

returns to last period's labor market activity.

Notice that in (5) we are assuming that firms pay consumers $W(s^{t-1})l(s^{t-1})$ at the end

of period $t-1$ and that the government transfers to consumers $[R(s^{t-1})-1]W(s^{t-1})l(s^{t-1})$ 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

choice 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. In particular,

notice that the first-order conditions for this problem imply that the labor-leisure decision is

undistorted and given by

$$-\frac{U_l(s^t)}{U_c(s^t)} = \frac{W(s^t)}{P(s^t)}.$$

To derive the consumer's demand for individual goods, we solve the consumption problem implicit in the definition of the price index (4) given by

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

subject to (2) and the constraint $\int P_i(s^t) c_i(s^t) = P(s^t) c(s^t)$, which gives the demand function for good i, namely

$$(6) \quad c_i(s^t) = \left(\frac{P_i(s^t)}{P(s^t)} \right)^{-\theta} c(s^t),$$

where $P(s^t)$ is given by (4).

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 reduction in the current period, which, for the sake of brevity we refer to as a sale. Having a sale at s^t entitles a firm for that one period to charge a price below the inherited regular price, that is, to charge $P(s^t) \leq P_R(s^{t-1})$. If the firm wants to continue that sale in the next period, it must again pay ϕ. In the period after the sale ends, the firm inherits the regular price that prevailed before the sale began, namely, $P_R(s^{t-1})$. (Note that when we evaluate our model's predictions relative to data, we will use the AC Nielsen algorithm to

define sales in both the model's data and the actual data. What we referred to above as sales for the sake of brevity would more precisely be termed *a temporary price reduction defined by the paying of the cost* ϕ. In the data generated by our model, such temporary price reductions will be correlated with AC Nielsen–identified sales, but the identification between the two will not be one-to-one.)

In our simple model, the only role of sales is to economize on the costs of changing prices, especially temporary price reductions. In our model, one way of temporarily lowering a price is to have a sale and to pay ϕ in each period of the sale. Another way is to pay the fixed cost to lower the regular price and then shortly thereafter pay a fixed cost again to raise it back up. As we will see in our quantitative model, the parameters will be such that firms often choose to temporarily lower the price by having a sale.

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

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

where we have used the demand function (6). The present discounted value of profits of the firm, expressed in units of date 0 dollars, is given by

$$(7) \quad \sum_t \sum_{s^t} Q(s^t) \left[R_i(P_i(s^t); s^t) - W(s^t)(\kappa \delta_{R,i}(s^t) + \phi \delta_{S,i}(s^t)) \right].$$

16

In equation (7) the term

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

is the labor cost of changing prices. The variable $\delta_{R,i}(s^t)$ is an indicator variable that equals one when the firm changes its regular price and zero otherwise, which $\delta_{S,i}(s^t)$ is an indicator variable that equals one when a firm has a sale and is zero otherwise. The constraints are that $P_i(s^t) = P_R(s^{t-1})$ if there is neither a sale nor a regular price change in that $\delta_{R,i}(s^t) = \delta_{S,i}(s^t) = 0$; $P_i(s^t) \leq P_R(s^{t-1})$ if there is a sale $\delta_{S,i}(s^t) = 1$; $P_i(s^t) = P_R(s^t)$ if there is a regular price change $\delta_{R,i}(s^t) = 1$, and that $\delta_{S,i}(s^t)\delta_{R,i}(s^t) = 0$ so that firms never have a regular price change and a sale in the same period.

Equilibrium

Consider now the market-clearing conditions and the definition of equilibrium. The market-clearing condition on labor,

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

requires that the labor used in production as well as the labor costs of making regular price changes and having sales 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 $l(s^t)$, allocations and prices for firms $\{P_i(s^t), y_i(s^t)\}_i$,; aggregate prices $W(s^t)$, $P(s^t)$, and $Q(s^{t+1}|s^t)$ that 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; (iv) the money supply processes and transfers satisfy the specifications above.

Computing the Equilibrium

It will be convenient to write the problem recursively. At the beginning of s^t, after the realization of the current money and productivity shocks, the state of an individual firm i is characterized by its regular price in the last period, $P_{Ri}(s^{t-1})$ and its idiosyncratic productivity level $a_i(s^t)$. It is convenient to normalize all of the nominal prices and wages by the current money supply. We let $p_{R-1,i}(s^t) = P_{Ri}(s^{t-1})/M(s^t)$, $w(s^t) = W(s^t)/M(s^t)$ and use similar notation for other prices. With this normalization we can write the state of an individual firm i at s^t as $(p_{R-1,i}(s^t), a(s^t))$.

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 on i, we write the state variables of a firm as $x = (p_{R,-1}, a)$ and the aggregate state variables as $S = (\mu, \lambda)$. Let

$$(8) \quad R(p_i, S) = \left(p_i - \frac{w(S)}{a} \right) \left(\frac{p_i}{p(S)} \right)^{-\theta} c(S),$$

where $w(S), p(S)$, and $c(S)$ are all known functions of the aggregate state. The function is the static gross profit function, normalized by the current money stock M. Let $\lambda' = \Lambda(\lambda, S)$ denote the transition law on the measure over the firm's state variables.

The value of a firm that does not adjust its price and instead uses its existing regular

18

price is

$$V^N(p_{R,-1}, a; \mu, \lambda) = \frac{U_c(S)}{p(S)} R(p_{R,-1}, S) + \beta E \left[V(p_{R,-1}, a'; \mu', \lambda') | a, \mu \right],$$

where $\lambda' = \Lambda(\lambda, S)$ and $U_c(S)$ and $p(S)$ are known functions of the aggregate state. (Notice that here we have expressed the value function in units of utility at date t. If we multiply these value functions at s^t by the constant $\beta^t \pi(s^t | s^0) P(s^0) / U_c(s^0)$, we convert them to units of dollars at s^0.)

The value of a firm that has a markdown is

$$(9) \quad V^M(p_{R,-1}, a; \mu, \lambda) = \max_{p_M \le p_{R,-1}} \frac{U_c(S)}{p(S)} \left[R(p_M, S) - \phi w(S) \right] + \beta E \left[V(p_{R,-1}, a'; \mu', \lambda') | a, \mu \right],$$

while that of a firm that adjusts its regular price is

$$(10) \quad V^R(p_{R,-1}, a; \mu, \lambda) = \max_{p_R} \frac{U_c(S)}{p(S)} \left[R(p_R, S) - \kappa w(S) \right] + \beta E \left[V(p_R, a'; \mu', \lambda') | a, \mu \right].$$

An intuitive way to think about the difference between a markdown and a regular price change is as follows. A markdown 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, and 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 (9) makes it clear that, conditional on having a markdown, the optimal pricing decision for p_M is static and that optimal price sets the marginal gross profit $R_p(p, S) =$

0. Note from the form of R in (8) that the optimal markdown is

$$(11) \quad p_M = \frac{\theta}{\theta - 1} \frac{1}{a} w(S).$$

Note that this price is a simple markup over the nominal unit cost of production and that, given this unit cost, does not depend on other features of the company such as the stochastic processes for shocks, the discount factor, the utility function, and so on. Indeed, 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 and depends on all of these properties (and, in particular, p_R will not typically equal p_M).

As (11) makes clear, conditional on having a markdown, the inherited regular price $p_{R,-1}$ is irrelevant so that we can write $p_M(a, S)$. Likewise, as inspection of (10) makes clear, conditional on having regular price change, the inherited regular price $p_{R,-1}$ is also irrelevant so that we can write $p_R(a, S)$.

3. Parameterization and Results

In this section we describe how we choose functional forms and benchmark parameter values. We then examine if our parsimonious model can account for the key features of prices that we have documented.

A. Parameterization

We consider a utility function of the form

(12) $\quad U(c, l) = \log(c) - \psi l.$

This specification follows Hansen (1985) by assuming indivisible labor decisions implemented with lotteries. We set the length of the period to one week, and therefore choose a discount factor $\beta = (.96)^{1/52}$. We choose ψ to ensure that in the absence of aggregate shocks, households supply one-third of their time to the labor markets. We choose $\theta = 3$, a number in the range of estimates of demand elasticities available in the retail industry. (See, for example, Nevo (2001), Barsky et al. (2000), and Chevalier, Kashyap, and Rossi (2003).)

Our model is weekly so that the process for money growth (1) in our numerical experiments is weekly as well. Given that the highest frequency at which M1 data are available is monthly, we pin down the serial correlation ρ_μ and volatility σ_μ^2 of weekly money growth in the model by requiring the model to generate a monthly growth rate of money that has the same serial correlation and volatility as the US M1 monthly growth rate during 1989–1997, the years for which the micro-price data used to calibrate the model are available.

The rest of the parameters are calibrated: κ, the cost the firm incurs when changing its regular price, ϕ, the cost of having a temporary sale, the persistence of the productivity process, ρ_a, as well as the distribution of innovations to the firm's productivity process, $\varepsilon_i(s^t)$. Midrigan (2006) shows when $\varepsilon_i(s^t)$ is normally distributed, the model generates counterfactually low dispersion in the size of price changes and 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 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 $[-\bar{\nu}, \bar{\nu}]$. This is the approach we take in our numerical experiments:

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

where $\nu_i(s^t)$ is distributed uniformly on the interval $[\underline{\nu}, \bar{\nu}]$.

One reason to pay particular attention to the distribution of productivity shocks is that this distribution plays an important role in determining the real effects of money. For example, Golosov and Lucas (2006) show that money is approximately neutral when productivity shocks are normally distributed. But, as Midrigan (2006) shows with a fat-tailed distribution of productivity shocks, money has much larger real effects. The reason is that as the kurtosis of the distribution of productivity shocks increases, the measure of marginal firms whose decision to adjust prices changes depending on the size of the money shocks is reduced, and changes in the identity of adjusting firms are thus muted.

B. Results

We show that our parsimonious model accounts well for the eight facts about prices and sales. We then give some intuition for how the model works.

Accounting for Our Eight Facts

We ask, Can the other five parameters, $\{\kappa, \phi, \rho_a, \lambda, \bar{\nu}\}$, be jointly chosen to mimic well the patterns of prices and sales in the data. In setting these parameters we target 13 moments

in the data indicated in Table 2. These moments include 2 on the frequency of price changes (including and excluding sales), 6 on the size and dispersion of price changes (including and excluding sales). These moments also include the fractions of prices at the annual mode, the fractions of annual prices below the model, the fraction of sales-related price change, the fraction of sales that return exactly to the pre-sale price, as well as the probability that a sale ends next week, conditional on there being a sale today.

In Table 2 we see that our parsimonious model does a remarkably good job at reproducing these facts. The frequency of price changes is high: .33 in both the data and the model with all prices (and much lower both in the data, .056, and the model, .062, when they are excluded.) The mean size of all price change is high in both the data (.17) and the model (.13) and the dispersion is high in both as well. The proportion of price changes that are the annual mode is also high: .58 in both the data and the model. When prices are not at their annual model, they tend to spend more time below the annual model than above it. Specifically, in the data prices spend 30% of their time below the annual mode and in the model about 27%. Most price changes are sale-related: 83% are in the data and 80% are in the model. Most sales tend to return to their pre-sale price: 87% do in the data and 89% do in the model. We also see that sales are very transitory: in the data the fraction of weeks with sales that are followed by weeks without sales is 46% in the data and 61% in the model.

Probably the most evident discrepancy between the model and the data is that the model generates somewhat smaller and less dispersed price changes than in the data, a result driven by our assumption that the process for idiosyncratic productivity shocks is symmetric and of the trade-off between generating a sufficiently large number of sales in the model (which requires making firms reluctant to temporarily increase prices) and generating sufficiently

sizable price changes. Moreover, the model lacks a mechanism to generate the large number of small price changes in the data, and this further reduces its ability to generate dispersed but sufficiently large price changes. Midrigan (2006) argues that a multi-product extension of a menu-cost model in which the cost of changing prices is shared among several products is capable of accounting for the large number of small price changes observed in the data. We chose not to pursue this route here in order to keep the model as simple as possible and highlight the special nature of sale-related price changes.

We also investigate our model's implications for some other moments. Consider, then, the eight 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 few weeks after a price change. Figure 3 shows that our model generates a sharply declining hazard in the first few weeks, as in the data. After the first few weeks, the hazard in the data continues to decline somewhat, while in our model the hazard is essentially flat. In Table 2 we also consider various statistics about what fraction of time over a given period that prices tend to remain at one of their most common prices over that period. We see that for most of these, the model produces numbers similar to those in the data.

Table 3 lists the parameter values that we choose to mimic the moments in the data problem. The menu cost of changing regular prices is fairly high—2.8% of a firm's steady-state labor bill. In contrast, the cost of temporarily marking down the price is 1.2% of a firm's steady-state labor bill, or about 40% of the cost of changing the regular price. Productivity shocks arrive with probability $\lambda = 0.28$ and have an upper bound of $\bar{\nu} = 0.23$. Moreover, the productivity process is highly transitory, $\rho_a = 0.3$. Given that sales last on average only 2 weeks, marginal cost shocks must be highly transitory for firms to be willing to return to

their regular price sufficiently fast.

The Workings of the Model

Consider the firm's optimal policy rules. These rules are a function of the individual states, the normalized regular price $p_{R-1} = P_{R,-1}/M$ and the current productivity shock a, as well as the aggregate state, the money growth rate. (Recall that in our approximation to the policy rules, we treat the distribution of λ as a constant.) We illustrate the optimal policy rules for the firm that has the option of a temporary markdown in Figure 4. In particular, Figure 4 shows the regular price the firm chooses conditional on adjusting it, as well as the regions of the state space in which the firm optimally chooses to have a regular price change, a markdown, or do nothing.

Recall that the firm-specific state variables is its normalized regular price $p_{R-1} = P_{R,-1}/M$ and its productivity level a. Figure 4 illustrates features of the optimal decision rule. The axes are the (logs of the) individual state variables. Movements in the horizontal direction come from movements in the exogenous productivity level. Movements in the vertical direction come from three sources. The first two come from either changing the regular price or having a markdown. The last can be roughly thought of as coming from a change in the money supply M, for a fixed level of the regular price $P_{R,-1}$. (To be precise, one way a firm could move horizontally on the graph from periods $t-1$ to t is for the money supply to change and for that firm not to change its regular price so that its normalized price p_{R-1} changes.) An increase in the money supply corresponds to a decrease in $p_{R-1} = P_{R,-1}/M$. (Throughout, we suppress explicit dependence on the aggregate state variables, S.) The shaded regions show where the firm finds it optimal to change its regular price (R), have a markdown (M),

and leave its price unchanged (N). The function $p_M(a)$, the solution to (9), is the price the firm would choose conditional on having a markdown. The function $p_R(a)$, the solution to (10), is the price the firm would choose conditional on it changing its regular price.

As we have noted above when discussing (9) and (10), the price $p_M(a)$ is a constant markup over marginal cost given by (11) and it does not equal the price $p_R(a)$. The markdown $p_M(a)$ in log space falls one-for-one with a for all values of a because the log of marginal costs fall one-for-one with a. In contrast, $p_R(a)$ is very flat in a when a is near its mean (of $\log a = 0$) but is steeper as a differs more from its mean. The reason for this difference in slopes reflects the dynamic considerations underlying the choice of p_R. For small deviations of a from its mean, the likelihood that the firm will adjust again next period is low. Given that a is rapidly mean-reverting, the firm chooses not to respond to small changes in a, as its current choice of p_R will last for a number of future periods. Hence, p_R is flat for small deviations in a. In contrast, when a is further away from its mean, the firm finds it optimal to respond more aggressively to this marginal cost disturbance today.

To understand why, consider two strategies for the firm. First, it could respond aggressively today, get higher static profits today, and then, with high probability, pay a second menu cost next period to adjust its regular price back toward its original level when a reverts closer to the mean. Alternatively, it could respond less aggressively today and then tomorrow plan to leave its price unchanged with high probability. In the second strategy it loses more in static profits in both periods but saves on paying the second fixed cost. When deviations in a are large, the first strategy is more profitable than the second, as the losses in profits outweigh the fixed cost and it responds aggressively.

The figure also illustrates the regions of the state space for which the firm finds optimal

to have a markdown, a regular price change, or sell at its previous period's regular price. Consider first the lower-left shaded region in which prices are much lower than $p_R(a)$. In this region, marked with an R, the firm would like to raise its price so its relevant choices are to leave its price unchanged or to change its regular price—as having a markdown is dominated by leaving its price unchanged. In this shaded region the firm's current regular price $p_{R,-1}$ is sufficiently low relative to $p_R(a)$ that the firm's benefit from increasing its regular price change outweighs the fixed cost. Roughly speaking, when the firm's current price $p_{R,-1}$ is below $p_R(a)$, the relevant "gap" in determining whether or not a firm adjusts is the distance between the current regular price $p_{R,-1}$ and $p_R(a)$.

When the firm's current price $p_{R,-1}$ is above both $p_R(a)$ and $p_M(a)$, all three of its options are relevant. When the current price is sufficiently high, the firm either has a temporary markdown or it changes its regular price. As the figure shows, when a is sufficiently far away from its steady-state level (the regions denoted with M), the firm finds it optimal to have a markdown. Since the stochastic process for a is not very persistent, when a is far from its mean it expects a large change in a in the next few periods as a reverts to its mean.

For intuition's sake, consider two strategies for a firm in such a circumstance. First, it could change its regular price today and then keep changing its regular price as the a reverts to its mean and stop changing it when it gets sufficiently close. Second, it could have a markdown today and each subsequent period until a gets close enough to its mean. For our parameter values in the regions marked M, the second strategy beats the first: it substitutes paying the smaller cost to have several markdowns for the larger cost of having several regular price changes.

Finally, consider the shaded region marked R above the $p_R(a)$ and $p_M(a)$ lines. In

that region, the productivity shock a is close to its mean and expected to stay there. Hence, the current price $p_{R,-1}$ is expected to be suboptimally high for a long period of time. In such a circumstance, having a markdown is not very attractive because the firm would expect to conduct on every period in the near future, and instead the firm finds it optimal to pay for a permanent price change to $p_R(a)$.

We can summarize the intuition on the choices by the firm using our earlier analogy on renting or buying a price. When the productivity level a is far from its mean and the current price $p_{R,-1}$ is too high, the firm prefers to rent a lower price today and expects to rent a series of them in the future as it waits for a to revert to its mean. In this region, it does not pay for the firm to invest in the high cost of buying a regular price because the value of that investment is expected to quickly deteriorate. When the productivity level a is far from its mean and the current price $p_{R,-1}$ is too low, the only price it can rent is one that is even lower and that is not optimal, so instead it buys a regular price change.

To be concrete, consider first the transition path of a firm that starts at steady state, the point labelled 0 in Figure 6A, and experiences a rise in productivity (of 20%) that moves its state at the beginning of period 1 to the point 1. This firm chooses a markdown in that period and lowers its price to be consistent with the point $1'$. Absent any new shocks, at the beginning of the next period productivity has fallen (to 6%) and since the firm's regular price is unchanged it starts that period at the point 2. It chooses to have a second markdown and hence charges a price consistent with the point labelled $2'$. Absent future shocks, the firm's productivity continues to revert to its mean and the firm stays in the inaction region and moves to points 3, 4, and so on.

Consider next the transition path of a firm that, starting from the steady state, expe-

riences a fall in productivity (of 20%) that moves it to point 1 at the beginning of period 1. In that period the firm chooses to raise its regular price and move to point 1′. By the beginning of period 2, productivity has somewhat reverted to the mean (so that it is now only 6% below the steady state). The firm chooses to have a markdown and charges a price consistent with point 2′ in that period. Absent future shocks, the firm's productivity continues to revert to its mean and the firm stays in the inaction region and moves to points 3, 4, and so on as it drifts back to the steady state.

Notice that two markdowns in the first transition path in Figure 6A would be classified as sales by the AC Nielsen filter but that the markdown in the second transition path in Figure 6B would not. As we shall see the first transition is much more likely than the second. More generally, most of the markdowns in our model are classified as sales by the AC Nielsen filter, and hence we can roughly associate markdowns in the model with sales in the data when giving intuition.

We now turn to giving some intuition for how our model can generate our facts about prices and sales. To help us do so, consider Figure 6 in which we illustrate the nature of the ergodic distribution. To do so, we partition the state space into a large number of equally sized cells and then shade each cell according to the ergodic probability that a realization of a state lies in that cell. The dark cells are most frequently visited cells that together account for 50% of the realizations. Likewise, the dark cells plus the medium dark cells account for 75%; the dark, medium dark, and the medium cells account for 95%; and all the shaded regions together account for 99%.

Consider now how our model can generate our facts. Let us start with Facts 1, 2, 3, and 4. Clearly our model generates frequent price changes because there is a substantial mass

outside of the inaction region N. It also generates large price changes because the average distance between most of the mass in the M and R regions to $p_M(a)$ and $p_R(a)$ is large. Next, prices spend most of their time within a year at the modal value because the mass inside the inaction region and the markdown region is large. Price changes are more likely to be below their annual mode than above it because there is more mass in the M regions (in which the firm has a temporary markdown below its annual mode) than in the R region below the inaction region (in which the firm typically raises its price above the annual mode).

Consider now how our model can generate the last four facts. In the model, most price changes are associated with sales because the mass in the M regions is much larger than the mass in the R regions. In the model, prices tend to return to the pre-sale price after sales because the transition path with the same features as those in Figure 6A is fairly typical. Sales are transitory in the model again because the process for productivity is transitory and firms that start in the M region with high productivity shocks, similar to point α or β in Figure 6A, tend to have a small number of markdowns in a row. The transitory nature of productivity shocks also helps explain why price changes are often clustered. If a firm starts with either high productivity shocks, as in Figure 6A, or low productivity shocks, as in Figure 6B, it tends to have a couple of changes in a row as the productivity reverts to the mean.

Next, consider next the impulse response of economic aggregates to an innovation in money growth. In Figure 7A we plot the response of the money supply and the response of prices. Since money growth is serially correlated, an innovation in money growth leads the level of the money supply to slowly increase to a new level. Since the relation

$$P_t C_t = M_t$$

holds in the model, the less this increase in money shows up in prices, the more it shows up in real consumption. From Figure 7A we see that as the money supply increases, the aggregate price level increases less and the gap between money and prices starts small, then widens, then narrows, giving the real effects on consumption in Figure 7B. Finally, in Figure 7C we also see that although the money shock decreases the fraction of firms with markdowns, the effects are rather minor.

Our summary measure of the real effects of money is the standard deviation of consumption, which from Table 5 is .44%. In this model if money has no real effect, this standard deviation would be zero, and the larger are the real effects, the larger is the standard deviation. Since we have made no attempt to incorporate into the model the sorts of "real rigidities" popular in the literature, we find this number useful as setting the benchmark against which to compare the relative sizes of the real effects of money in alternative versions of the model rather than interesting in its own right.

4. Experiments

The key question asked in the sticky price models is, How large are the real effects of monetary policy? We contrast the answer to this question in our model with markdowns to that of simpler models without them. Our analysis serves a dual role. First, it helps in giving intuition for how our model works. In particular, it helps illustrate why the effects of monetary policy differ after price changes that are the result of markdowns relative to price changes that are the result of regular price changes. Second, if we treat our model as the data-generating process, we can ask which of two alternative practices more closely reproduces the real effects of monetary policy in our model: leaving the sales out of the data

31

or leaving the sales in the data.

We begin by comparing our model to two menu cost models without markdowns. In one version of the model, the *leave-sales-in* version, we set the parameters of the model to match statistics in the data in which sales are included. In the other, the *take-sales-out* version, we set parameters to match statistics in the data in which sales are excluded. We find that the leave-sales-in version significantly understates the real effects of monetary policy relative to these effects in the model with markdowns and that the take-sales-out version somewhat overstates the real effects of monetary policy.

We then perform a similar comparison using the more popular Calvo model of pricing, which is the benchmark model in the sticky price literature. These Calvo models are viewed as approximations to the underlying menu cost models. We also consider a leave-sales-in version and a take-sales-out version of a Calvo model. We find the same qualitative results as in our previous comparison: the leave-sales-in version understates the real effects of money, while the take-sales-out version overstates it. At a quantitative level, the approximation error with these models is large. In the leave-sales-in version, the real effects of money, as measured by the standard deviation of consumption, is less than one-fifth the level in the menu cost model with markdowns. In the take-sales-out version, the real effects of money are about twice the level in the menu cost model with markdowns.

We then propose an alternative procedure to set the parameters of a model without markdowns in order to approximate the real effects of money in a model with markdowns. In this procedure, we leave sales in the data but instead of choosing parameters to match the frequency of price changes (along with other statistics), we choose parameters to match the fraction of prices at the annual mode (along with the same other statistics). We show that

32

for either a menu cost model without markdowns or the Calvo model, this procedure implies real effects of money similar to those in the menu cost model with markdowns.

A. Comparison to Fixed Cost Models without Markdowns

Here we conduct the following experiment. We generate data from our model with markdowns and then fit two models without markdowns to this data. These models are special cases of our model with markdowns in which we disallow temporary markdowns. We then ask, How well do the two models without markdowns do in reproducing the real effects of money in the generated data? We measure the real effects of money by the standard deviation of consumption.

We then consider fit a third model without markdowns to the generated data: one in which we give up on matching the frequency of price changes and instead match the fraction of prices that are at the annual mode.

The Leave-Sales-In and the Take-Sales-Out Versions

In the *leave-sales-in* version, we choose parameters to match four statistics of price changes in this generated data in which we make no attempt to take out "sales". The parameters are κ, the fixed cost of price adjustment, the arrival probability of productivity shocks, λ, and the upper bound on productivity shocks, $\bar{\nu}$. The statistics, listed in Table 2, are the fraction of price changes per week (.31) and three moments of the distribution of the absolute value of price changes: its mean, the 25th percentile, and the 75th percentile (.13, .08, and .17). The resulting parameters are listed in Table 3.

In the *take-sales-out* version, we remove sales from the generated data by applying the same AC Nielsen–type filter to the generated data that we have used in our analysis of

Dominick's data. We then choose the same parameters, κ, λ, and $\bar{\nu}$, to match the same four statistics computed for the filtered data. For the statistics see Table 2 under the column Model with markdowns, exclude sales and for the resulting parameters see Table 3.

We can shed some light on the difference in parameters in the two versions as follows. From Table 2 we see that when we leave sales in the data generated by our markdown model, price changes are much more frequent (about once every three weeks) than when we exclude them (about once every 16 weeks). A model without markdowns can generate such frequent price changes in two ways: small menu costs and more volatile shocks. Relative to the model calibrated to the generated data with sales removed, the model calibrated to the data with sales included has both lower fixed costs and more volatile shocks. Specifically, in the take-sales-out version, the cost of a regular price change is 1.25% of the steady-state labor bill and the productivity shocks arrive with 3.8% probability. In the leave-sales-in version, the costs are lower, only .30%, and the shocks arrive more frequently, with 18.5% probability.

Consider now the regions of adjustment and inaction and the policy rules in the two models. Figure 8 plots these regions together with the policy rules for firms resetting the regular price $p_R(a)$, where again we suppress dependence on the aggregate state S. Figure 9 plots the analogous regions and policy for the take-sales-out version. In both of these figures, the policies for price adjustment are much simpler than in the model with markdowns. In both, if the regular price is sufficiently high, it is optimal to pay the fixed cost and adjust it down, while if it is sufficiently low, it is optimal to pay the fixed cost and adjust it up. Comparing the two figures, note that, as one might expect, the inaction region in the take-sales-out version is somewhat wider than in the leave-sales-in version.

More informative are Figures 10 and 11, which illustrate the nature of the ergodic

distribution in the two versions. As before, the dark cells account for 50% of the realizations; the dark cell plus the medium dark cells account for 75%; the dark, medium dark, and the medium cells account for 95%; and all the shaded regions together account for 99%. In Figure 10 the mass in the adjustment regions marked is 34%, while in Figure 11 this mass is only 5.7%. Most of the reason for this difference in mass is in the variance of shocks. Recall from Table 2 that in the leave-sales-in version, shocks arrive nearly five times as frequently as in the take-sales-out version. The rest of the difference is that the inaction region in the leave-sales-in version is wider than that of the take-sales-out version, so that for the same shocks a firm sometimes adjusts in the leave-sales-in version when it would not in the take-sales-out version.

Let us turn now to the main question of this section. How do the real effects of monetary policy differ in the models without markdowns relative to the model with markdowns? We measure the real effects of monetary policy by the volatility of consumption. In Table 4 we see that leaving sales in the data leads us to underestimate the effects of monetary policy, while taking sales out of the data leads us to overestimate the effects of monetary policy. The real effects in the version that leaves sales in are 22% (.10/.45) of what they are in the model with markdowns, while the real effects in the model that takes sales out is 131% (.59/.45) of what they are in the model with markdowns.

To get some intuition for why the benchmark model generates larger real effects than the leave-sales-in version even though prices in both change equally frequently, consider the following heuristic example. Consider a firm that starts at time t in the deterministic steady state and then at time $t + 1$ experiences a 10% increase in productivity and a 5% decrease in the money supply. For simplicity, suppose that from period $t + 2$ on, productivity returns

35

to the steady state and that the money supply stays permanently lower by 5%. If there were no costs of changing prices, this firm at $t + 1$ would lower its prices 15% and then at period $t + 2$ would increase its prices by 10% and leave them unchanged thereafter. Such a path of prices would lead to there being no real effects from money in any period.

Now consider the behavior of a firm that faces this same sequence of shocks in the benchmark model. We heuristically describe this behavior with points marked 0, 1, 2, and 3 on Figure 12 and with these same points on the graphs in lower panels of Figure 13. (Note that we are associating a 5% drop in money with a 5% shift in the state variable $p_{R,-1} = \frac{P_{R,-1}}{M}$.) In the benchmark model, this firm in period 1 has a markdown and lowers its prices by 15%, thus offsetting fully the real effects of money in that period. From period 2 on the firm, faced with just the permanently lower level of money, finds it unprofitable to pay the fixed cost to change its regular price. Hence, the firm is stuck forever with its old regular price, which is 5% higher than the frictionless price, and this sequence of shocks leads to real effects from period 2 on.

Next, consider the behavior of a firm that faces this same sequence of shocks in the leave-sales-in version in Figure 15 and see also the dashed line in Figure 16). In period 1 this firm lowers its regular price by approximately 15%, approximately offsetting the money shock in that period. By period 2 the productivity shock has disappeared so that the firm is at the beginning of the period the firm is at point 2. It then raises its regular price approximately 10% and leaves it there from then on. This firm's behavior then well approximates that of a firm in a frictionless environment and, at least for this firm, there are no real effects from money.

Notice that in this simple example, even though the firm in the benchmark model and

the firm in the leave-sales-in version both had two price changes, the real effects were quite different. The reason is the nature of the price changes. In the benchmark model the changes were a temporary markdown and then a reversal of that markdown back to the original price. The fact that the second price change restored the original price meant that it restored the real effects of money to be what they would have been from periods 2 on to what they would have been if it had not changed prices at all. In the leave-sales-in version, the changes were both regular price changes. Both of these changes offset the change in the money supply, thus leading to no real effects at all from this firm.

Finally, consider a heuristic description of a similar experiment in the take-sales-out version. In this version the productivity shocks arrive about one-fifth as frequently as in either the benchmark model or the leave-sales-in version. In Figures 16 and 17 we mimic this by supposing that no productivity shock occurs while money again decreases 5% permanently. Here the firm never changes its regular price so the money has a real effect in every period. Here money has a larger real effect in the benchmark model because in that model the temporary markdown offset the real effects in period 1.

We can summarize the key message of our heuristic example as follows. In assessing the real effects of money, the nature of price changes matters as well as their frequency. In particular, temporary markdowns are a very different type of price change than a change in a regular price: markdowns, unlike regular price changes, typically return to their original price and hence do not let firms permanently respond to changes in money.

A Version Matching the Fraction at the Annual Mode

Finally, we consider an alternative procedure for setting the parameters in models in which markdowns are not explicitly modeled. We choose to match the fraction of prices at the annual mode. In Table 2 we report on the statistics we attempt to match again treating the benchmark model with markdowns as the data-generating process. The statistics are the mean size of price changes, the 25 and the 75 percentile of price changes and, most importantly, the fraction of prices at the annual mode. In Table 3 we report on the parameter values we choose and in Table 4 the results. From Table 4 we see that, again using the standard deviation of consumption as the measure of the real effects of money, we see that this procedure works much better than either of the two approaches that we have discussed.

B. Comparison to Calvo Models

Here we perform a similar experiment to those above with two versions of a Calvo model. Specifically, we again generate data from our benchmark model and then fit two Calvo models to this data: a leave-sales-in version in which the frequency of price changes is set to match the generated data with sales left in, and a take-sales-out version in which this frequency is set to match the generated data after the AC Nielsen filter is applied to it. We then ask, How well do these two Calvo models do in reproducing the real effects of money in the generated data?

The Calvo models are very similar to the benchmark model described above except that it has time-dependent sticky prices and no temporary markdowns. The consumers in 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 Calvo (1983). Specifically, in a

given period with probability λ a firm is allowed to adjust, and with probability $1 - \lambda$ the firm is not allowed to adjust. We refer to λ as the *frequency of price adjustment*.

The problem of a firm that is allowed to adjust at s^t is

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

where

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

Since there are no costs to price adjustment, the resource constraint is simply

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

The parameters of technology, preferences, and stochastic processes are set to be equal to those in our benchmark model. The additional parameter that needs to be set is λ. We consider two parameterizations corresponding to the two alternative practices discussed above.

In the leave-sales-in version, we set $\lambda = .34$ so that firms get an opportunity to change their prices 34% of the time, and in the take-sales-out version we set $\lambda = .057$ so that firms get an opportunity to change their prices 5.7% of the time. As the appendix shows, the implications of a log-linear approximation to this model do not depend on moments of the distribution of the idiosyncratic shock process other than the mean. In particular, the implications depend neither on the arrival rate nor on the variance of the idiosyncratic productivity shocks.

In Table 4 we see that, as with our earlier model, leaving sales in the data leads us to underestimate the effects of monetary policy, while taking sales out of the data leads us to overestimate the effects of monetary policy. The real effects in the Calvo version that leaves sales in are less than one-fifth (.08/.44) of what they are in the benchmark model, while the real effects in the version that takes sales out are almost double (.82/.44) of what they are in the benchmark model.

Finally, we see from Table 4, if instead of matching the frequency of price changes, we choose λ to match the fraction of prices at the annual model we do much better than the other two procedures.

5. Conclusion

In the data a sizable fraction of price changes are sales-related. Existing sticky price models abstract from explicitly modeling these changes. We have added one parameter to the standard sticky price model and have shown that the parsimoniously-extended model can well account for many of the patterns of prices and sales in the data.

We have evaluated existing approaches by explicitly including a motive for temporary price reductions in a simple sticky price model and then directly study how large are the real effects of monetary policy shocks. We then treat the model as the data-generating process and apply the two common approaches to data generated from our model. We find neither of the existing approaches provides a good approximation to the key question of interest: the take-sales-out approach leads to much larger effects of monetary shocks than are in the model with sales, and the leave-sales-in approach leads to much smaller effects of monetary shocks than in the model with sales.

We then show that a simple rule of thumb approach of using a model without sales but choosing parameters, not to match the frequency of price changes, but rather to match the fraction of time a price stays at its annual model gives a much better approximation than either of the existing approaches. We end with two proposals to advance the sticky price literature: either explicitly include sales in the model or follow a version of our rule of thumb. We argue that either will represent progress relative to the existing approaches.

Notes

[1]See Hoch and Purk (1994), Peltzman (2000), and Chevalier et. al (2003) for a discussion of Dominick's experiment, pricing strategy and of the data.

[2]We prefer to use this algorithm to identify temporary price cuts in the data rather than flags in the data base that supposedly indicate sales because the Dominick's sales flag is inconsistently coded and fails to record all sales.

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

Burstein, Ariel. 2002. Inflation and output dynamics with state-dependent pricing decisions. Mimeo.

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. 2006. Menu costs and Phillips curves. *Journal of Political Economy.*

Hoch, Stephen, Xavier Dreze and Mary E Purk, 1994, "EDLP, Hi-Lo, and Margin Arithmetic," *Journal of Marketing*, 58: 16-27

Klenow, Peter J., and Oleksiy Kryvtsov. 2005. State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? Working Paper 11043, National Bureau of Economic Research.

Midrigan, Virgiliu. 2006. Menu costs, multiproduct firms, and aggregate fluctuations. Mimeo, Federal Reserve Bank of Minneapolis.

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

Peltzman, Sam, 2000, "Prices Rise Faster Than They Fall," *Journal of Political Econ-*

omy, 108(3): 466-502

6. Appendix: Sales Filter

Step 1: return above or to the same level.

For each price cut (defined as a period in which $p_t < p_{t-1}$ & p_t, p_{t-1} are non-missing), check if $p_{t+j} \geqslant p_{t-1}$ for $j \leqslant 3$ (5). Let $\bar{j}$ be the minimum j for which this condition is satisfied (if at all). Replace $p_t, p_{t+1,..}, p_{t+\bar{j}-1}$ etc. with p_{t-1} for those js between p_t and $p_{t+\bar{j}}$ for which the original price is available. The figure below plots the output (new vs. old price series). Notice in the first sale that we impose no restrictions on what the old price does prior to the period it returns to a level equal to or above the pre-sale price. It can fall gradually, stay fixed, gradually increase, etc. Also notice (slightly prior to week 260) 2 price cuts in the data, of which only one (the second) is called a "sale" at this point. The first one is excluded because the price rises to a level below the pre-sale price.

There is no need to iterate this filter more than once, as it leaves no "sales" in the newly created regular price series by construction (i.e., after running the sales filter once, there are no other price cuts left in the new "regular" price series that will revert to something $\geqslant$ the original in $\leqslant 3$ weeks). The next step needs to be iterated more than once, however, if gradual price increases, or price decreases following the original price cut, are to be allowed for.

Step 2: return below the original level.

For each price cut (defined as a period in which $p_t < p_{t-1}$ & p_t, p_{t-1} are non-missing), check if $p_{t+j} \geqslant p_t$ for $j \leqslant 3$ (5). Let $\bar{j}$ be the minimum j for which this condition is satisfied (if at all). Replace $p_t, p_{t+1,..}, p_{t+\bar{j}-1}$ etc. with p_{t-1} for those js between p_t and $p_{t+\bar{j}}$ for which there are no gaps in the original price series.

45

We apply this second procedure 3 (5) times in order to filter out sales periods associated with price changes following the original price cut.

Table 1: Frequency and size of price changes

	All price changes	Exclude sale-related price changes
Frequency of price changes	0.33	0.056
Mean size of price changes	0.17	0.10
25 percentile size of price changes	0.04	0.03
75 percentile size of price changes	0.19	0.12
Fraction of annual prices at mode	0.58	
Fraction of 50-week prices at top 2 prices	0.76	
Fraction of 50-week prices at top 3 prices	0.84	
Fraction of 100-week prices at mode	0.50	
Fraction of 100-week prices at top 2 prices	0.69	
Fraction of 100-week prices at top 3 prices	0.78	
Fraction of annual prices below mode	0.30	
Fraction of sale-related price changes	0.83	
Proportion of returns to pre-sale p	0.87	
Probability a sale ends	0.46	

Table 2: Calibration targets

Moments	Data		Model with markdowns		Model w/o markdowns		
	Include sales	Exclude sales	Include sales	Exclude sales	including sales	excluding sales	prices at annual
Used for calibration							
Frequency of price changes	0.33	0.056	0.31	0.062	0.32	0.061	0.087
Mean size of price changes	0.17	0.10	0.13	0.14	0.13	0.14	0.13
25 percentile size of price changes	0.04	0.03	0.08	0.06	0.07	0.05	0.06
75 percentile size of price changes	0.19	0.12	0.17	0.20	0.17	0.21	0.18
Fraction of annual prices at mode	0.58		0.58		0.30	0.66	0.59
Fraction of annual prices below mode	0.30		0.27				
Fraction of sale-related price changes	0.83		0.80				
Proportion of returns to pre-sale p	0.87		0.89				
Probability a sale ends	0.46		0.61				
Other moments							
Fraction of 50-week prices at top 2 prices	0.76		0.75				
Fraction of 50-week prices at top 3 prices	0.84		0.81				
Fraction of 100-week prices at mode	0.50		0.40				
Fraction of 100-week prices at top 2 prices	0.69		0.63				
Fraction of 100-week prices at top 3 prices	0.78		0.74				

Table 3: Parameter values

	Markdowns	No markdowns		
		Match data including sales	Match data excluding sales	Match fraction of prices at annual mode
Cost of changing regular price, % of SS labor bill	2.80	0.30	1.25	0.85
Cost of temporary markdown, relative to menu cost	0.18	x	x	x
Arrival rate of technology shock	0.28	0.185	0.038	0.052
Upper bound on technology shock	0.23	0.26	0.27	0.25
Persistence of technology	0.30	0.30	0.30	0.30
Arrival rate of preference shock	x	x	x	x
Upper bound on measure of high-elasticity customers	x	x	x	x
Substitution elasticity of type-1 consumers	3	3.00	3.00	3.00
Discount factor	$0.96^{1/52}$	$0.96^{1/52}$	$0.96^{1/52}$	$0.96^{1/52}$
Autocorr. of growth rate of money supply	0.90	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}	8.31×10^{-4}

Table 4: Results

	Markdowns	No markdowns		
Model with menu-costs		Match data including sales	Match data excluding sales	Match fraction of prices at annual mode
Std. dev. of HP-detrended monthly consumption	0.45	0.10	0.59	0.44
Autocorr. of HP-detrended monthly consumption	0.88	0.65	0.86	0.79
Pass-through of monthly ΔM into ΔP	0.83	0.96	0.78	0.80
Model with Calvo pricing				
Std. dev. of HP-detrended monthly consumption		0.08	0.82	0.57
Autocorr. of HP-detrended monthly consumption		0.77	0.93	0.92
Pass-through of monthly ΔM into ΔP		0.97	0.63	0.74

Figure 1: Example of price series

A&W Cream Soda Diet, 6/12 oz, 4 bottles per case

"Regular" price

Transaction Price

week

Figure 2: Hazard of price adjustment

length of spells, weeks
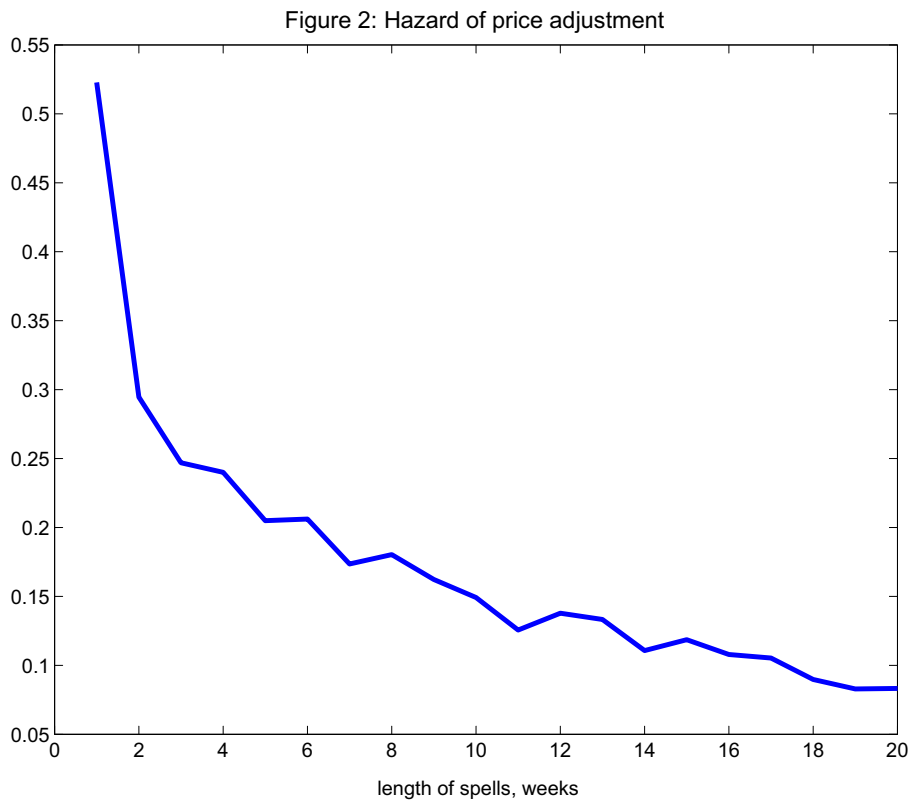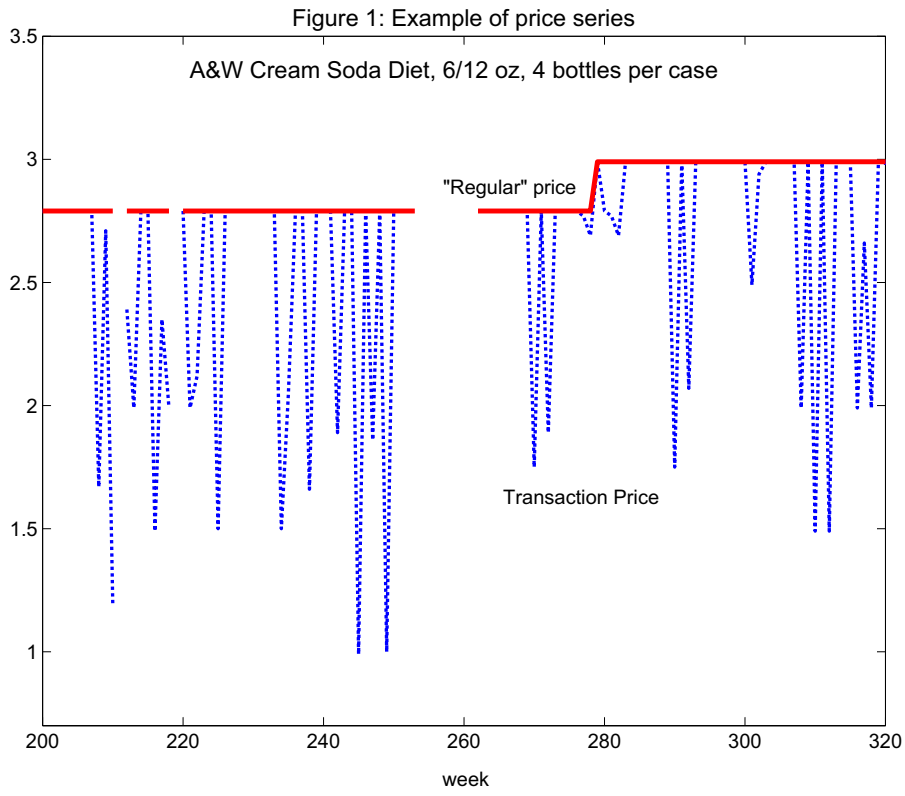
Figure 3: Hazard of price adjustment
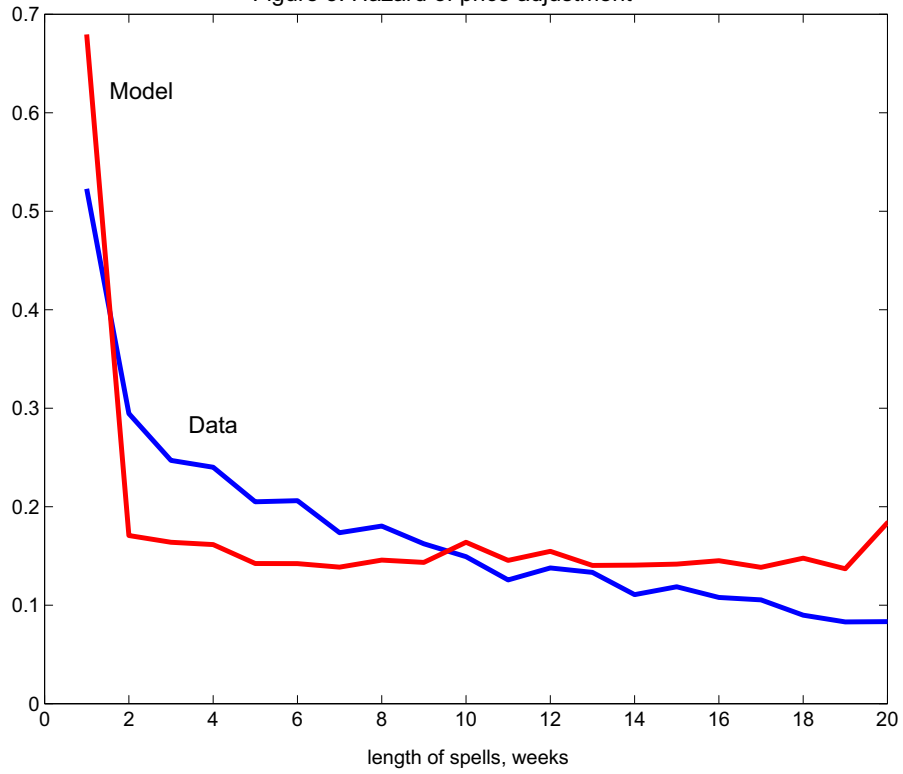
Figure 4: Optimal policy rules: model with markdowns
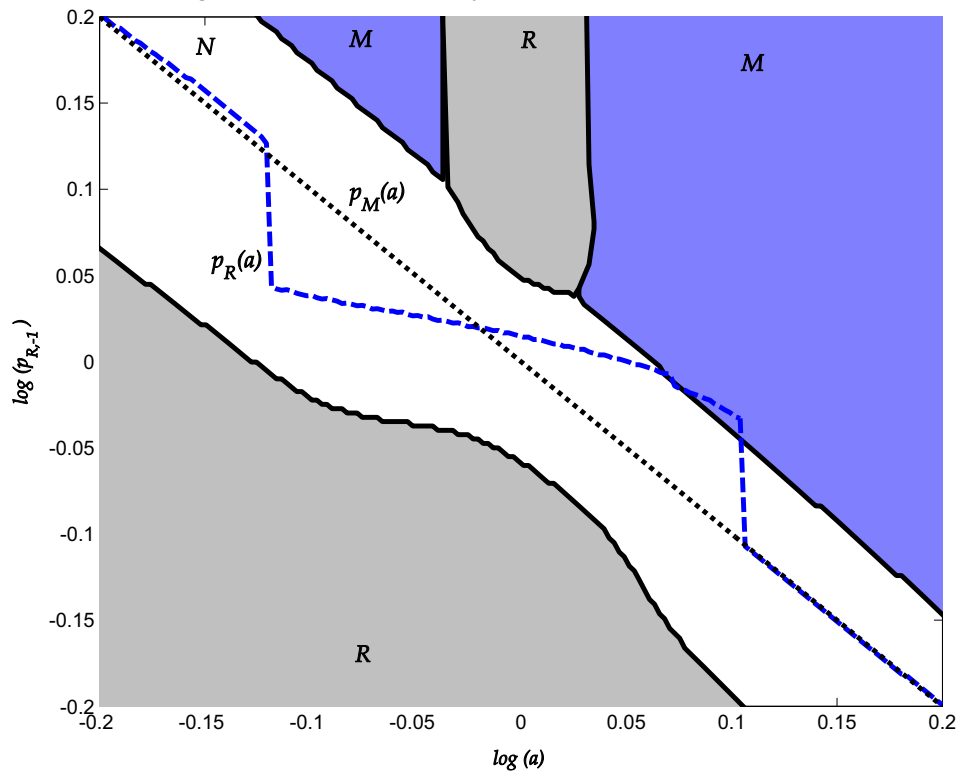
Figure 5A: Transition path for firm with high productivity
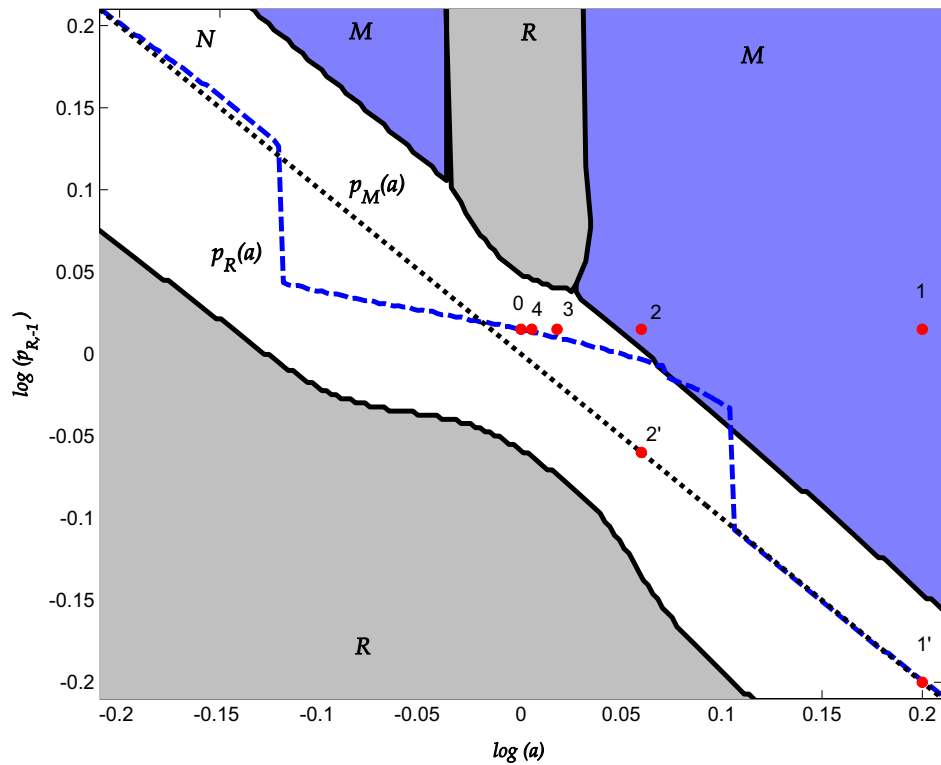
Figure 5B: Transition path for firm with low productivity
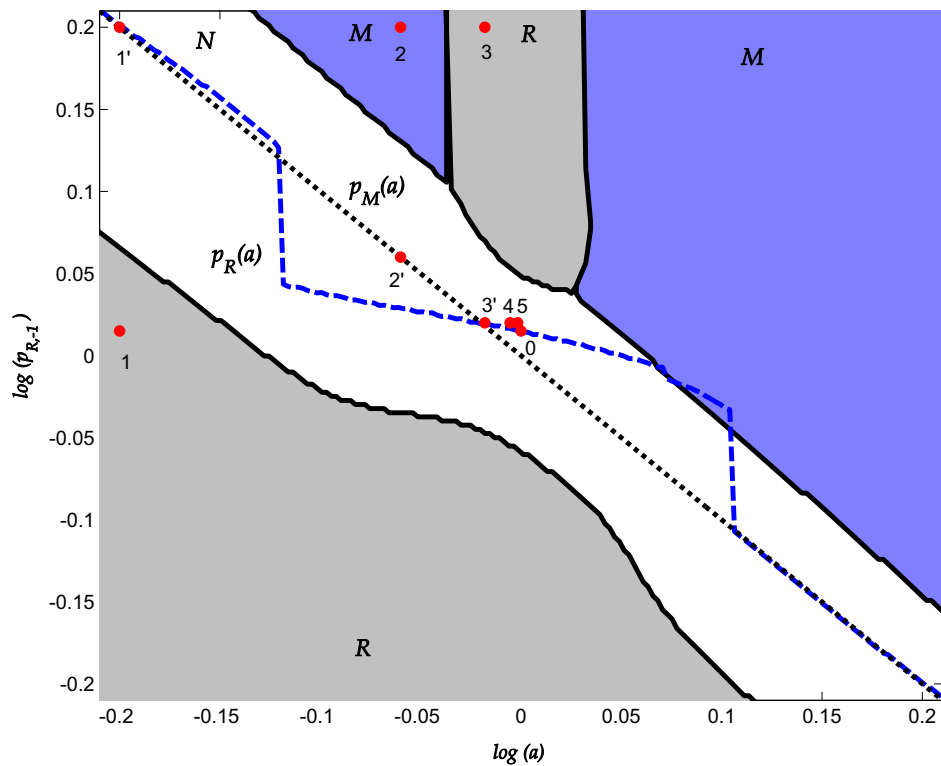
Figure 6: Ergodic distribution, model with markdowns
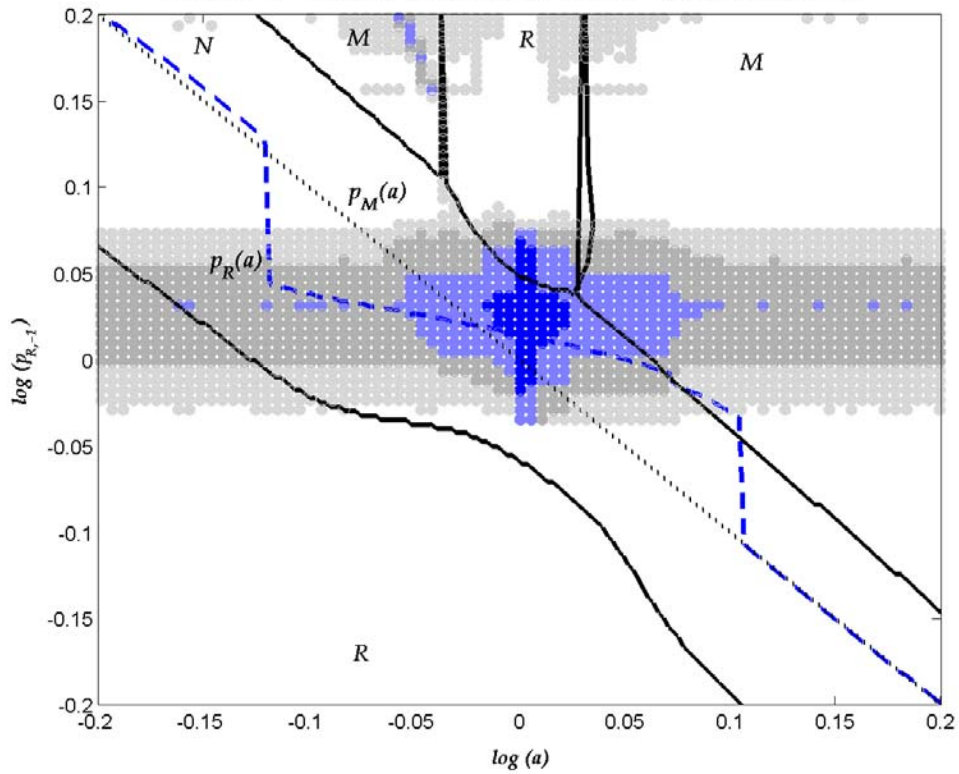
Figure 7: Impulse Response to Money Shock

7A
Price index, Money

7B
Consumption

7C
Fraction with markdowns
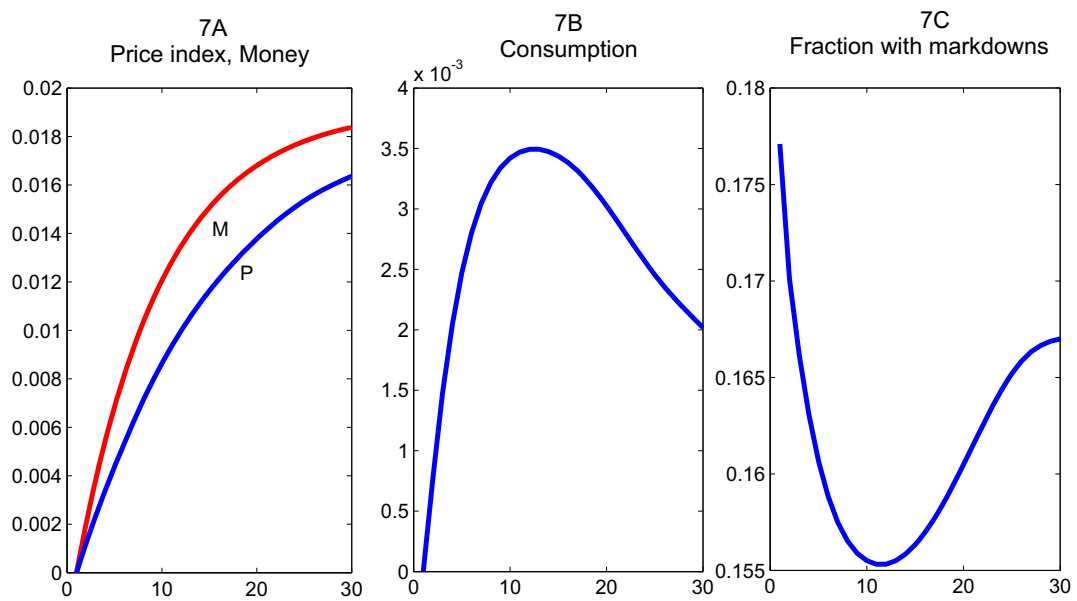
Figure 8: Optimal policy rules: no markdowns, leave in sales
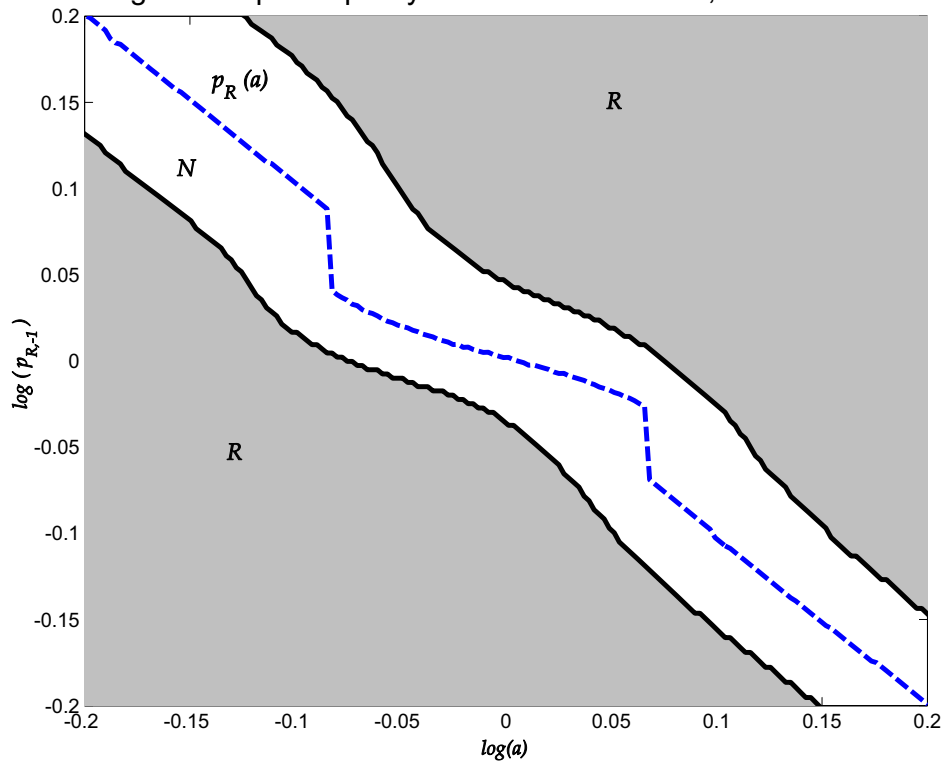
Figure 9: Optimal policy rules: no markdowns, leave sales out
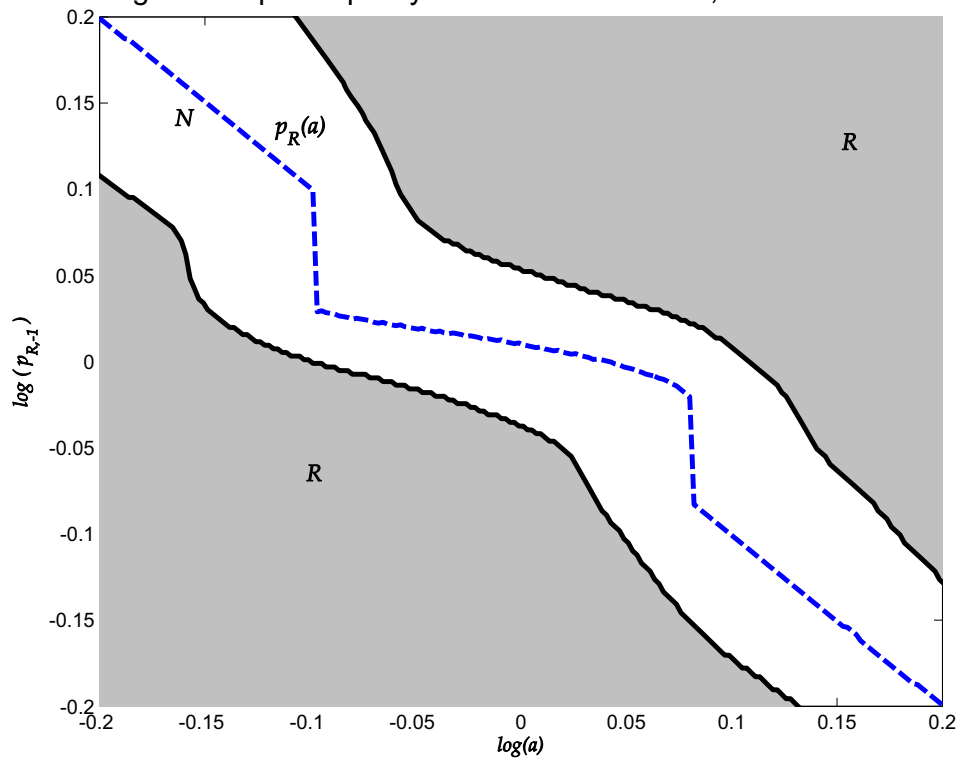
Figure 10: Ergodic distribution: no markdowns, leave in sales
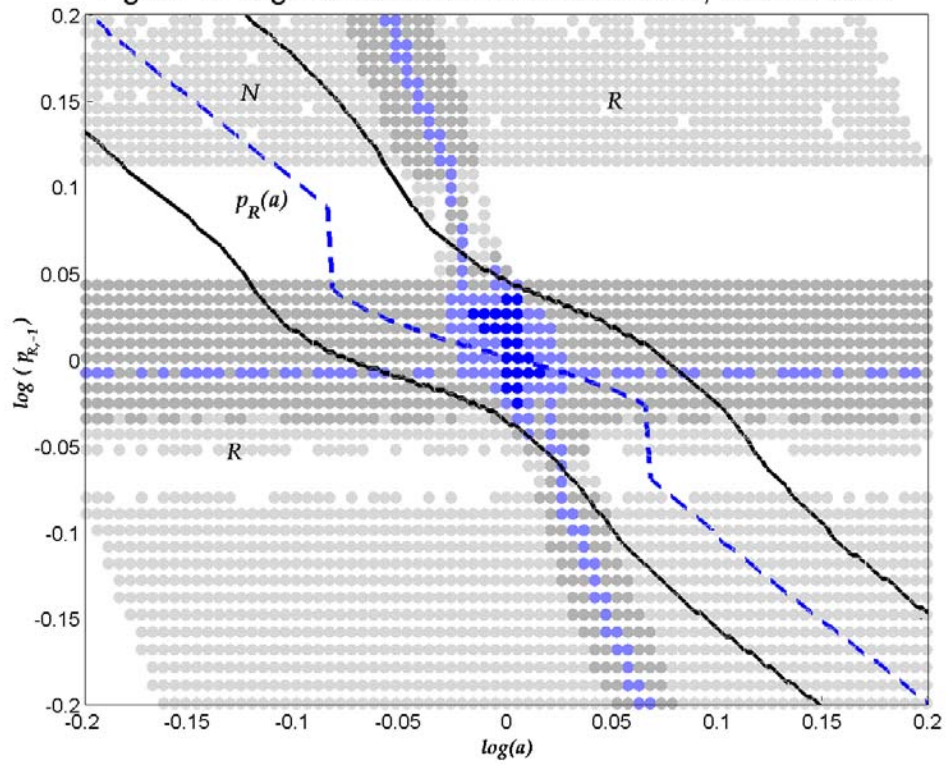
Figure 11: Ergodic distribution, no markdowns, leave out sales
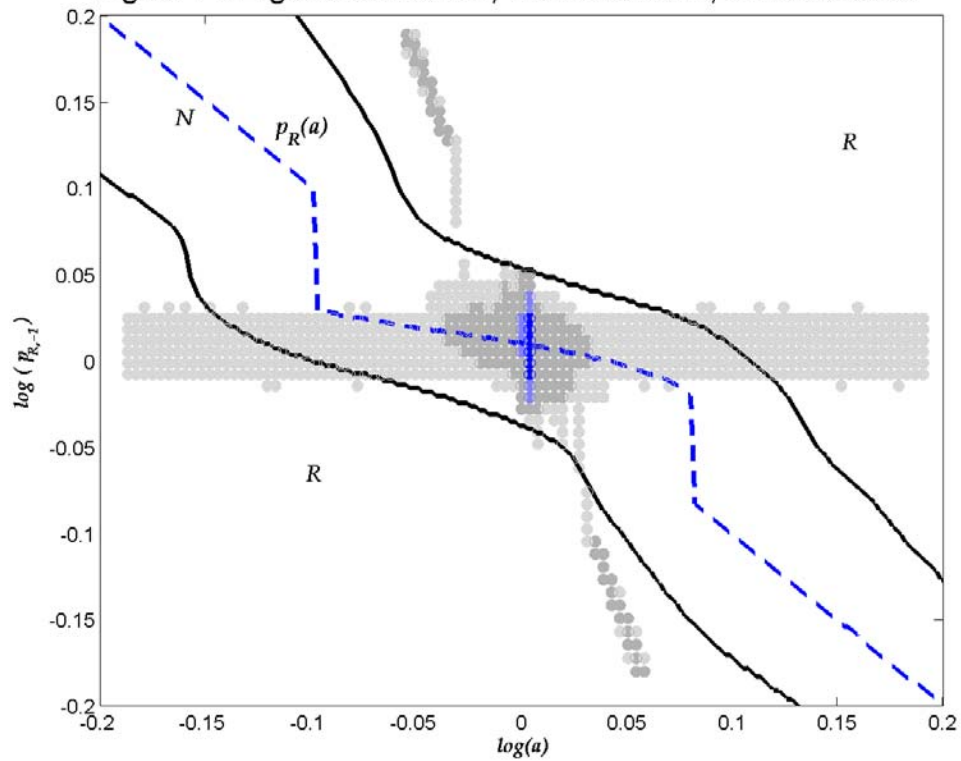
Figure 12: response to productivity and money shock, model w/ markdowns
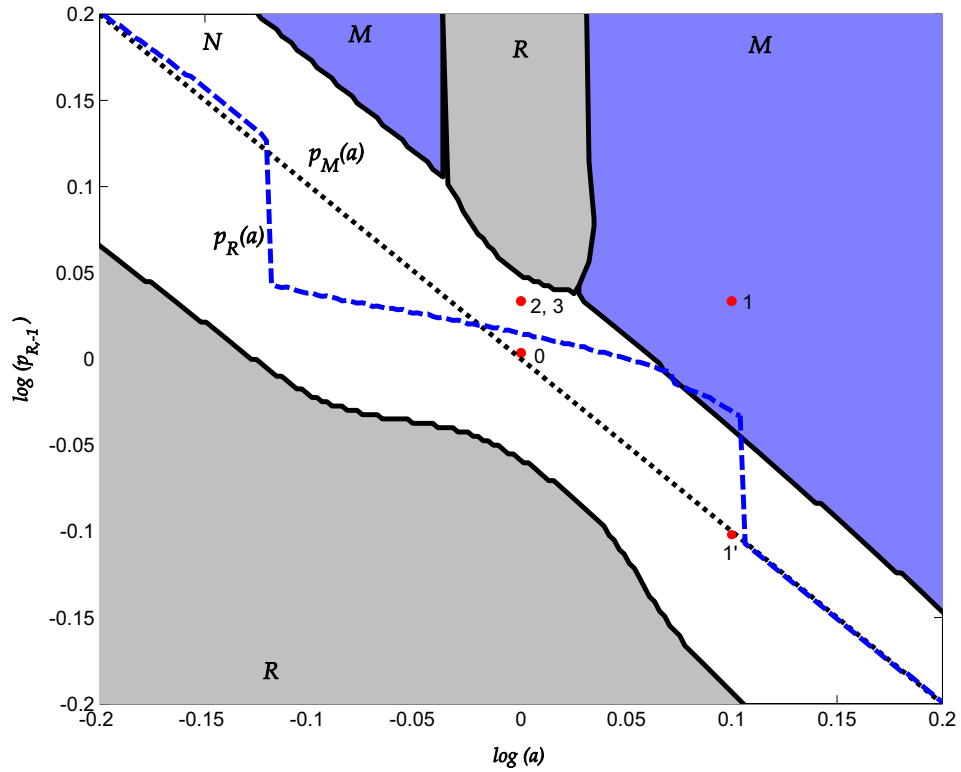
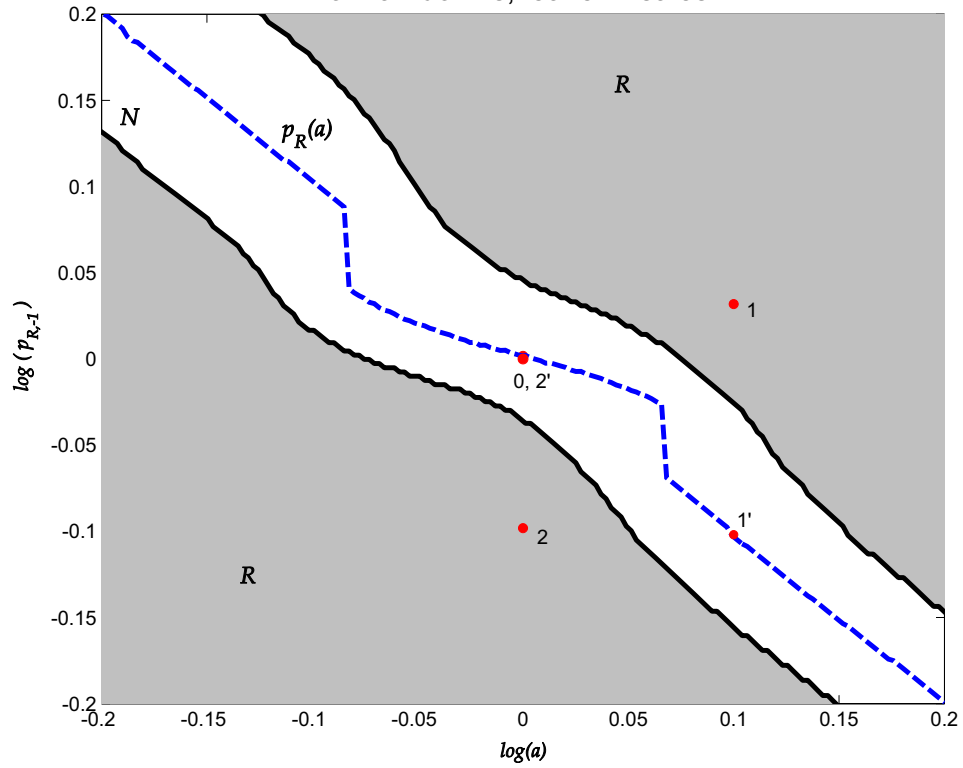Figure 13: response to productivity and money shock:
model with markdowns
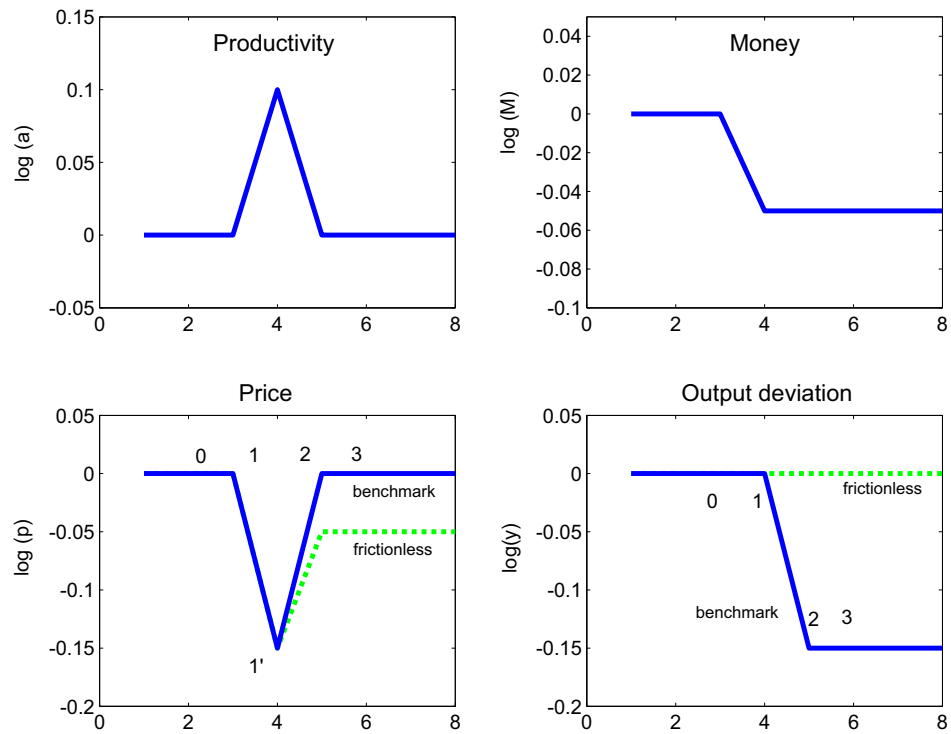
Figure 14: response to productivity and money shock:
no markdowns, leave-in-sales

Figure 15: response to productivity and money shock:
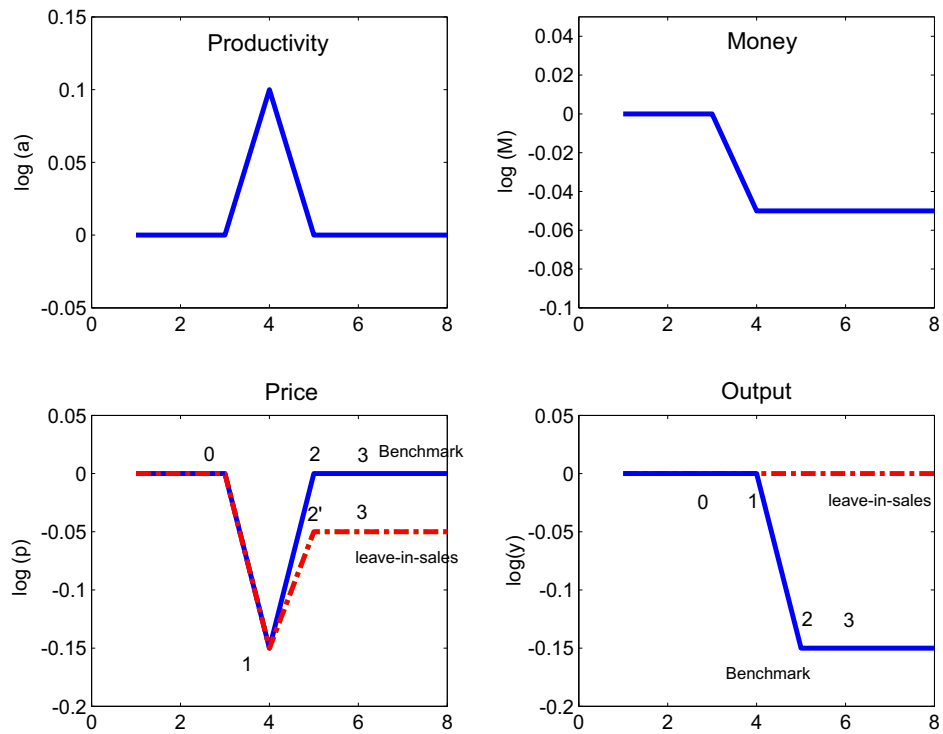no markdowns, leave-in-sales

Figure 16: response to productivity and money shock:
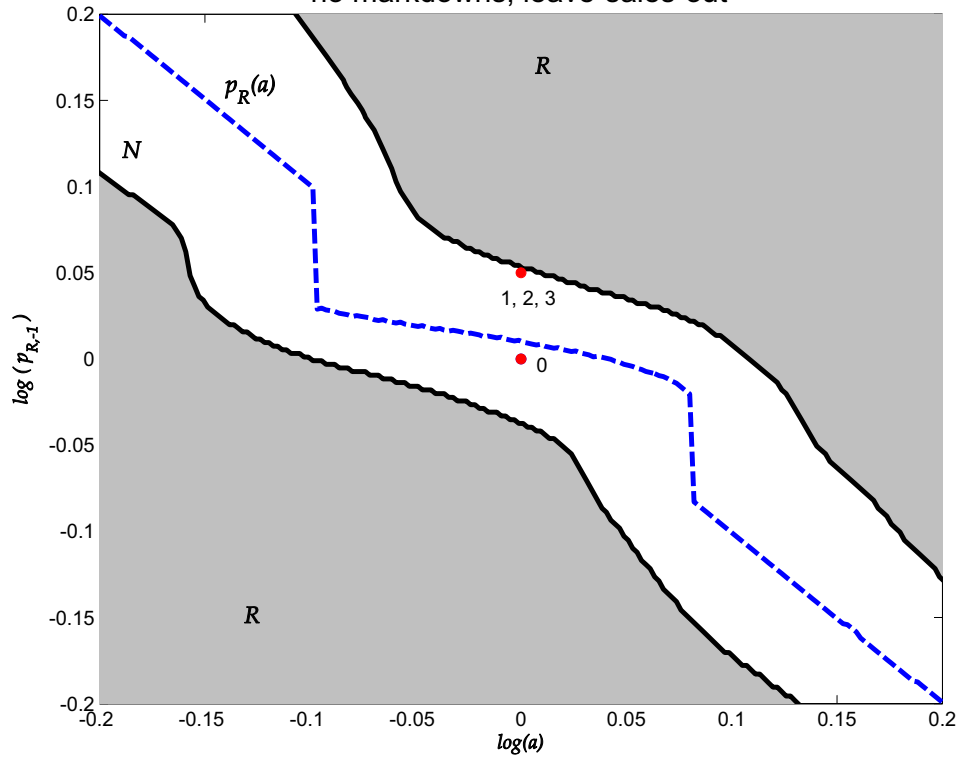no markdowns, leave-sales-out

Figure 17: response to productivity and money shock:
no markdowns, leave-sales-out