Econometrica, Vol. 79, No. 4 (July, 2011), 1139–1180

MENU COSTS, MULTIPRODUCT FIRMS, AND AGGREGATE FLUCTUATIONS

BY VIRGILIU MIDRIGAN[1]

Golosov and Lucas recently argued that a menu-cost model, when made consistent with salient features of the microdata, predicts approximate monetary neutrality. I argue here that their model misses, in fact, two important features of the data. First, the distribution of the size of price changes in the data is very dispersed. Second, in the data many price changes are temporary. I study an extension of the simple menu-cost model to a multiproduct setting in which firms face economies of scope in adjusting posted and regular prices. The model, because of its ability to replicate this additional set of microeconomic facts, predicts real effects of monetary policy shocks that are much greater than those in Golosov and Lucas and nearly as large as those in the Calvo model. Although episodes of sales account for roughly 40% of all goods sold in retail stores, the model predicts that these episodes do not contribute much to the flexibility of the aggregate price level.

KEYWORDS: Menu costs, multiproduct firms, sales.

1. INTRODUCTION

A WIDELY HELD VIEW in macroeconomics is that menu costs of price adjustment give rise to aggregate price inertia and thus provide a mechanism through which changes in monetary policy have real effects. This view lies at the heart of new-Keynesian analysis, which often cites menu costs as the microfoundation for the assumptions on price setting it makes.

Golosov and Lucas (2007) recently challenged this view. They argued that when a standard menu-cost model is made consistent with the microprice data,[2] money is nearly neutral. The key feature of the price data they emphasize is that the size of price changes is very large, nearly 10 percent on average. When they choose parameters governing an idiosyncratic productivity shock to produce this feature of the microprice data, the model predicts little stickiness in the aggregate price level. Their paper is viewed as a serious challenge to the key mechanism underlying new-Keynesian business cycle analysis.

This paper revisits the Golosov and Lucas result. I argue that while their analysis captures some features of the microprice data, it does not produce two salient features of that data. The first is that in the data there is a large amount of heterogeneity in the size of price changes, so that relative to their

[1]I am indebted to George Alessandria, Bill Dupor, Paul Evans, Patrick Kehoe, and Mario Miranda for valuable advice and support, as well as Ariel Burstein, Michael Dotsey, Mark Gertler, Narayana Kocherlakota, Ellen McGrattan, John Leahy, Martin Schneider, Takayuki Tsuruga, Ivan Werning, and Harald Uhlig for comments and suggestions. I acknowledge the James M. Kilts Center, GSB, University of Chicago, for making the data used in this study available. Any errors are my own.

[2]Bils and Klenow (2004) and Klenow and Kryvtsov (2008).

 DOI: 10.3982/ECTA6735

model's predictions, the data have both many more small price changes and many more very large price changes. The second feature is that they abstract from sales, even though sales (or more generally, temporary price changes) account for a large majority of price changes.

I add several ingredients to the Golosov and Lucas model to make it consistent with these two additional microfacts as well as the original facts that motivate their study. When I do so, I find that their near-neutrality result is overturned. In particular, my extended model has real effects of money that are of a similar magnitude to those in the popular Calvo model.

To allow the model to produce very large price changes, I allow for a fat-tailed distribution of cost shocks and let the data pin down the shape of the distribution. To allow the model to produce small price changes, I assume economies of scope in price adjustment. In particular, I assume the retailer sells multiple goods and faces a single cost of changing the prices of these goods. Since the idiosyncratic shocks to the different goods are not perfectly correlated, at any point in time some goods' prices may be far from their desired level while others may be close to their desired level. My model formalizes the ideas in Lach and Tsiddon (2007), who provided empirical evidence for this mechanism.

The standard menu-cost model also does not produce the type of temporary price changes observed in the data. To produce such changes, the theory must, of course, provide a motive for retailers to change prices temporarily. A striking feature of temporary price changes is that almost 90% of the time, after a temporary price change, the nominal price returns exactly to the nominal price that prevailed before the change. This feature suggests the need for a nominal friction to account for such behavior.

Existing studies have suggested a number of explanations for temporary price discounts.[3] All of them, however, are about real prices and hence, they cannot explain the striking feature of the data: the nominal price, after a sale, typically returns exactly to the nominal pre-sale price. This latter feature is critical for understanding how nominal prices respond to nominal shocks. Indeed, as emphasized by Kehoe and Midrigan (2008), this feature of the data has an important effect on the model's implications regarding monetary nonneutrality. Even though there is a large amount of high-frequency variation in prices associated with temporary price changes, there is much less low-frequency variation, which is ultimately what matters for how aggregate prices respond to low-frequency variation in monetary policy.

Here instead of building a detailed model of why retailers offer temporary price discounts, I simply assume that these changes are a response to tem-

[3]Search frictions (Butters (1977), Varian (1980), Burdett and Judd (1983)), demand uncertainty Lazear (1986), thick-market externalities (Warner and Barsky (1995)), loss-leader models of advertising (Chevalier, Kashyap, and Rossi (2003)), and intertemporal price discrimination (Sobel (1984)) to name a few.

porary changes in wholesale costs. While this mechanism is simple, it is not without realism if one focuses on retail pricing: as documented by Eichenbaum, Jaimovich, and Rebelo (2011), most retail prices change in response to a change in costs. This also allows me to focus on the nominal rigidity used by Kehoe and Midrigan (2008) that allows the model to reproduce the striking feature of temporary price changes.

To get the model to account for the striking reversion to the pre-sale price, I assume that retailers set two types of prices—posted prices and regular prices—both of which are costly to change. Although not specifically modeled, the idea behind these two costs is that posted prices are set by the retailer's salespeople and regular prices are set by the headquarters. Recent studies find some evidence for such price-setting practices. For example, Zbaracki, Ritson, Levy, Dutta, and Bergen (2004), Zbaracki, Levy, and Bergen (2007) found that pricing is done at two levels: the upper management sets the overalpricing strategy of the store (the regular price), while the sales personnel has the discretion to deviate from the regular price by quoting the consumer a discounted price.

My model has two main differences from the Golosov and Lucas model. First, my model produces more heterogeneity in the size of price changes. Second, I have many temporary price changes, while Golosov and Lucas do not. It turns out that in accounting for why my model produces much larger real effects of money than does the Golosov and Lucas model, the first difference is the most important.

Consider first the role of heterogeneity in price changes. When heterogeneity is small, as it is in Golosov and Lucas, monetary policy has a strong *selection effect* that mitigates the real effects of money. When heterogeneity is large, as in my model, money policy has a very weak selection effect and money has much larger effects. Recall that the real effects of a given money shock are smaller the larger the response of the price level is to this shock. In the Golosov and Lucas model, prices respond strongly to a money shock; in my model, they do not.

The reason that prices respond more strongly to a money shock in the Golosov and Lucas economy is the strong selection effect in which changes in monetary policy alter the mix of adjusting firms toward firms whose idiosyncratic shocks call for larger price increases.[4] The strength of this selection effect depends on the measure of marginal firms whose desired price changes lie in the neighborhood of the adjustment thresholds. When the measure of marginal firms is relatively large, as in Golosov and Lucas, the implied distribution of price changes is bimodal with little size dispersion: most price changes are near

[4]Also related is the Caplin and Spulber (1987) neutrality result. Other contributions in the menu-cost literature are Sheshinski and Weiss (1977), Caplin and Leahy (1991), Caballero and Engel (1993, 2007), Dotsey, King, and Wolman (1999), Danziger (1999), Burstein (2006), and Burstein and Hellwig (2008).

the adjustment thresholds. In contrast, when the measure of marginal firms is small, as it is in my model, most firms are dispersed away from their adjustment thresholds, and hence the distribution of the size of price changes shows large dispersion. Accounting for the dispersion of the size of price changes is thus key to studying the aggregate implications of menu-costs economies.

I show that in the data, the distribution of the size of price changes is, in contrast to the predictions of the Golosov–Lucas model, highly dispersed. Individual goods experience, over their lifetime, both very small and very large price changes. My modified menu-cost model produces these features of the data and because of this has a small selection effect. Indeed, my model's predictions are fairly close to those of a constant-hazard Calvo model in which the selection effect is absent.

Consider next the role of temporary price changes in accounting for my results. Briefly, in my model, retailers frequently adjust their prices for just a couple of periods when they experience idiosyncratic temporary cost shocks. During these periods, the retailers find it optimal to offset the effects of any monetary policy shock by adjusting its price to neutralize such shocks. Hence, all else equal, adding such temporary changes lowers the real effects of money shocks. It turns out that, as Kehoe and Midrigan emphasized, this effect is not large, even though the model accounts for the fact that about 40% of all goods are sold during such periods.

Consider, finally, the robustness of my result. In my model I have focused on economies of scope in price setting as a mechanism to generate small price changes and fat-tailed distributions to generate large price changes. My result seems to be robust to alternative mechanisms to generating such heterogeneity in price changes. For example, a simple alternative way to generate small price changes is to make the menu-cost stochastic, as in Dotsey, King, and Wolman (1999) so that sometimes the menu-cost is so small that retailers change their price even if it is close to the desired price. For this alternative mechanism, I found the real effects of money to be similar to the model with economies of scope. I chose to focus on the model with economies of scope because there is some evidence that economies of scope are indeed a feature of the price-setting technology in retail stores. In addition to the evidence of Lach and Tsiddon (2007), Levy, Bergen, Dutta, and Venable (1997) presented direct evidence of economies of scope by directly describing the technology of price adjustment in a large retail store. Moreover, I present additional evidence for economies of scope in this paper. I show that an individual good is more likely to experience a price change if other goods sold by the retailer show large desired prices. This latter feature of the data is inconsistent with the predictions of a menu-cost model without economies of scope.

Although the main focus of the paper is, due to data availability, on a single grocery store, the results apply more generally to all sectors that comprise the U.S. consumer price index (CPI). To illustrate this point, this paper considers an extension of the analysis to recently available data from Nakamura and Steinsson (2008). I show that the strength of the selection effect in a model that accounts for the large mean and dispersion in the size of price changes in all sectors of the U.S. economy is small as well. Similarly, my results extend to an environment with capital accumulation and interest-elastic money demand.

In addition to the work of Golosov and Lucas (2007), this paper is closely related to recent work by Klenow and Kryvtsov (2008) and Gertler and Leahy (2008). Both of these papers present parameterizations of economies in which menu costs generate aggregate price inertia almost as large as in Calvo-type models. In Klenow and Kryvtsov, the selection effect is weak because of the assumption of time-varying adjustment costs: most price changes occur in periods in which the realization of the adjustment cost is low. The economy studied by Gertler and Leahy is also characterized by a weak selection effect because of the assumption of a Poisson process for idiosyncratic shocks. The economy I study here shares elements of both of these approaches and, importantly, provides empirical evidence for the mechanism that underlies these results.

2. DATA

Here I describe the data and the algorithm I use to distinguish between regular and temporary prices. I then report several salient features of the data that motivate the subsequent analysis.

A. *Description of the Data Set*

I use a data set of scanner price data, maintained by the Kilts Center for Marketing at the University of Chicago Graduate School of Business (GSB).[5] The data set is a by-product of a randomized pricing experiment[6] conducted by the Dominick's Finer Foods[7] retail chain in cooperation with the Chicago GSB. Nine years (1989–1997) of weekly store-level data on the prices of more than 9000 products for 86 stores in the Chicago area are available. The products available in this data base range from nonperishable food products to various household supplies, as well as pharmaceutical and hygienic products. Dominick's sets prices on a weekly basis (it changes prices, if at all, only once a week),[8] so I report all statistics at the weekly frequency. Data are available on the actual transaction prices faced by consumers, as well as quantity sold and

[5]The data are available online in the Supplemental Material.

[6]Hoch, Dreze, and Purk (1994) discussed Dominick's experiment in detail.

[7]An earlier version of this paper also looked at data collected by AC Nielsen that also is available from the same source. The AC Nielsen data are constructed from an underlying panel of the purchasing history of households and subject to more measurement and time-aggregation problems. For this reason, I chose to eliminate this additional data set from the analysis in this paper.

[8]Its pricing cycle is weekly, so price change, if at all, occurs only once during the week. See Dutta, Bergen, and Levy (2002, p. 1854).

the wholesale cost of the good.[9] Dominick's sets prices on a chainwide basis, so prices across stores are highly correlated.[10] I choose thus to work with the price of one single store, number 122, the store with the largest number of price observations that was part of the control zone and thus not subject to the pricing experiment.

B. *Algorithm to Identify Regular Prices*

To identify regular prices, I make use of an algorithm discussed in Kehoe and Midrigan (2008). The algorithm is based on the idea that a price is *regular* if the store charges it frequently in a window of time adjacent to that observation. The regular price is thus equal to the modal price in any given window surrounding a particular week provided the modal price is used sufficiently often. The algorithm is somewhat involved and I relegate a formal description to Appendix 1 in the Supplemental Material (Midrigan (2011)).

Figure 1, reproduced from Kehoe and Midrigan (2008), presents several time series of the original price, p_t, as well as the regular price, p_t^R, constructed using the algorithm. The algorithm does quite a good job of identifying the much more transitory price changes associated with temporary discounts from the more persistent price changes associated with a change in the regular price. Throughout this paper, I use the term "temporary prices" to denote prices that are not equal to the regular price in any given period. Although the algorithm I use is symmetric (it treats temporary price increases identically to temporary price decreases), most temporary price changes are associated with sales.

To summarize, the algorithm I use is a filter that distinguishes high-frequency price variation from low-frequency price movements, much like the Hodrick–Prescott (HP) filter is used to identify trend from cycle. As with the Hodrick–Prescott filter, the algorithm is characterized by a parameter (the size of the window around the current observation used to compute the modal price: I use a 10-week window) that determines the relative frequency of the two types of price changes. Importantly, the results I report are not too sensitive to the exact length of the window (results for a 6- and 20-week window are virtually identical). Moreover, as is standard in business cycle research, I apply an identical filter to both my model and the data. Finally, I will report below a number of additional statistics that are independent of the filtering method I use so as to gauge the sensitivity of my results to alternative decompositions of high-/low-frequency price changes.

[9]The cost variable is an "average acquisition cost," a weighted average of the cost of goods available in inventory. See Peltzman (2000) for a description of the cost data.

[10]See Peltzman (2000) for a discussion of Dominick's pricing practices. He noted that the retail price at Dominick's stores is set by the chain's head office and the correlation between price changes within a particular pricing zone are in the neighborhood of 0.8 to 0.9.

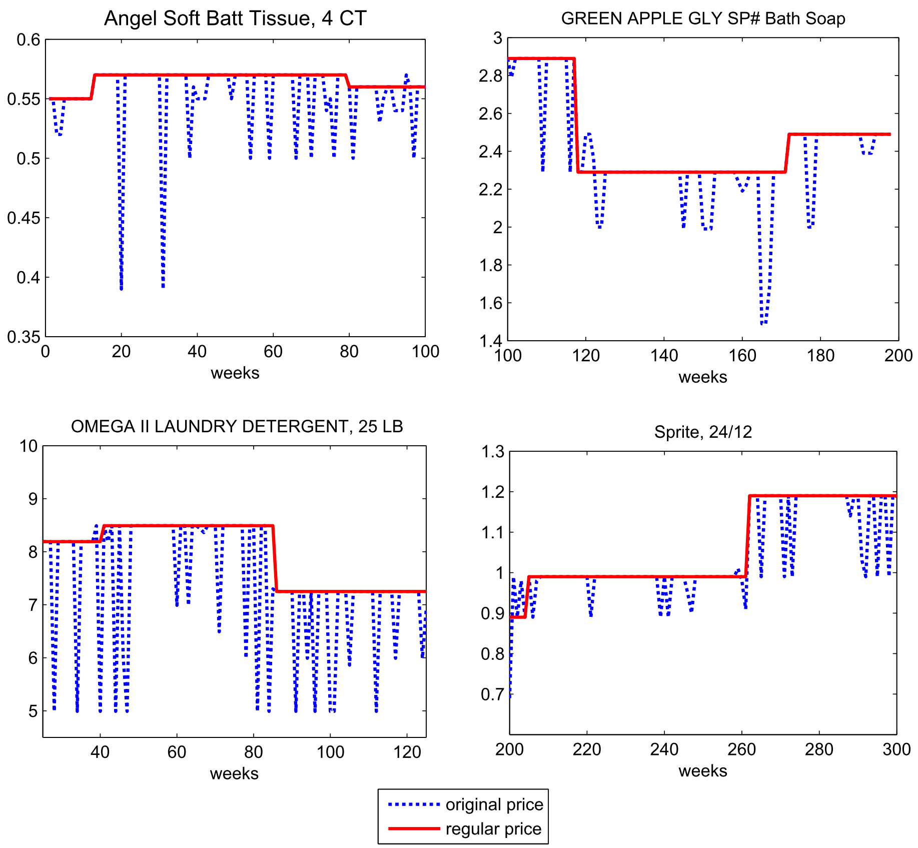

FIGURE 1.—Example of the algorithm.

C. *Motivating Facts*

Golosov and Lucas argued that the standard Calvo model vastly overstates the real effects of monetary shocks relative to a model that is consistent with the microdata on price changes. The Golosov and Lucas model, however, is inconsistent with three important features of the microdata. First, their model predicts little heterogeneity in the (absolute) size of price changes, while in the data there is a large amount of heterogeneity. Second, their model abstracts from temporary price changes, even though these changes account for the vast bulk of price changes in the data and periods with temporary discounts account for a sizeable fraction of goods sold. Third, I document a striking feature of temporary price changes: after a temporary price change, the price tends to revert to the exact nominal price (to the penny) in effect before the change. Here I briefly document these features of the data and use them to motivate my model.

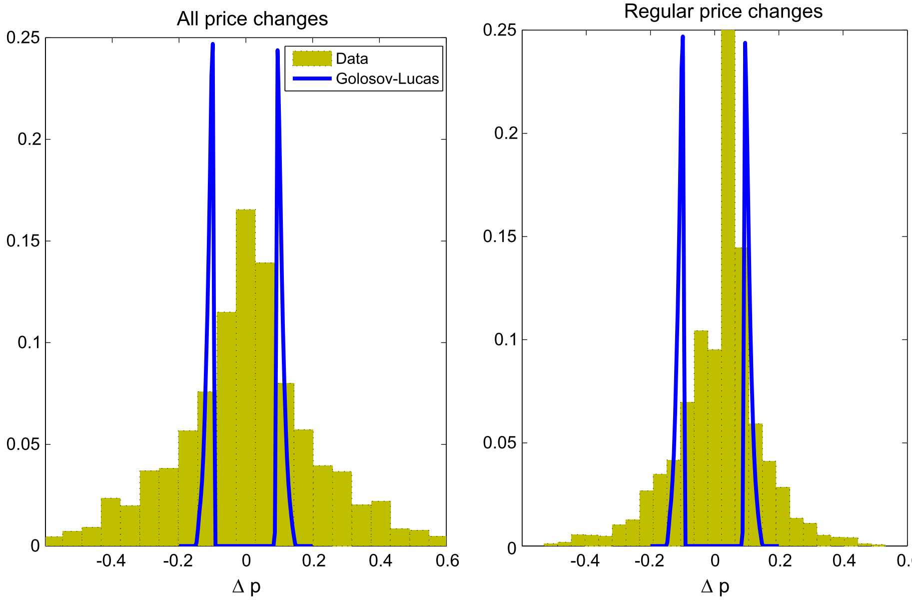

FIGURE 2.—Distribution of nonzero price changes: Dominick's versus Golosov–Lucas model.

To document the heterogeneity of price changes, I present, in Figure 2, the distribution of (nonzero) price changes in the data and contrast it with the distribution implied by the Golosov and Lucas model (described in the next section). The two panels give the histogram for all price changes and regular price changes. Superimposed on these histograms is a solid line that gives the density implied by the Golosov and Lucas model. Clearly, the Golosov and Lucas model is unable to produce heterogeneity in the absolute size of price changes seen in the data. Specifically, it has both too few small price changes and too few large price changes relative to the data. For example, the standard deviation of the absolute size of price changes in the Golosov–Lucas model is 1.2%, while the corresponding number for regular prices in the data is 8.2%.

To document the importance of temporary price changes, I note two facts. First, most price changes in the data are temporary price changes. Second, a sizeable fraction of goods are sold during periods in which a temporary price is chosen. The first fact is evident in Figure 1 for the selected series. For the data set as a whole, 96.5% of price changes are temporary. For the second fact, 40% of a store's goods are sold during periods with a temporary price cut. Thus, in two senses, temporary price changes are an important feature of the data.

To document the striking reversion of prices following temporary price changes, I note that 86% of the time, the nominal price following a temporary price change returns to the exact level it had before the change. Because of this feature of the data, although the frequency of all price changes is large (34% of prices change every week), the frequency of regular price changes is much smaller (2.9% per week).

3. A MENU-COST ECONOMY

I extend the standard menu-cost model by adding several ingredients that allow the model to reproduce the features of data discussed above. I then calibrate the model and contrast its implications for the real effects of money with those of Golosov and Lucas.

As I have documented, the standard menu-cost model produces too few small price changes as well as too few very large price changes. To allow the model to produce very large price changes, I follow Gertler and Leahy (2008) and assume a Poisson process for idiosyncratic productivity shocks. The frequency with which retailers receive these shocks is chosen so that the model matches a number of moments of the distribution of price changes in the data. To allow the model to produce small price changes, I assume economies of scope in price adjustment. In particular, I assume the retailer sells multiple goods and faces a single cost of changing the prices of these goods. I present evidence for both of these departures from the standard model in an empirical section (Section 3.C).

To account for the pattern of temporary and regular price changes in the data, I make several additional assumptions. Retailers choose two prices for each good: a regular price, p_t^R, as well as a price to quote to the consumer (the posted price), p_t. The latter is the price at which the consumer purchases the good and thus determines the store's revenue. The regular price affects the retailer's profits because every time the retailer posts a price that differs from the regular price, it must incur a fixed cost, κ. As a result, the posted price will deviate from the regular price infrequently, only when the benefit from doing so exceeds the fixed cost. I assume that changing prices is costly. Changing the regular price entails a fixed cost, ϕ^R. Similarly, changing the posted price requires a fixed cost, ϕ.

These assumptions are motivated by evidence on the pricing practices of firms in the data. Zbaracki et al. (2004), Zbaracki, Levy, and Bergen (2007) provided evidence that prices are set at two levels: (a) the managerial level (e.g., by the headquarters) that sets list (regular) prices and (b) the sales force that is responsible for posted prices (including discounts, rebates, etc.). These authors found that managerial costs of changing list prices are substantially greater than the physical costs of changing posted prices (the costs of printing new price lists, etc.). Moreover, they also found that salespeople must coordinate departures of their prices from the regular (list) prices with the upper-level managers, activities that presumably involve a cost. See, for example, the following quote from a manager they interviewed[11]:

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

A. *Setup*

The economy is populated by a representative consumer, a unit measure of monopolistically competitive retailers, indexed by z, and a continuum of manufacturers. Each retailer sells N products, indexed by $i = 1, \dots, N$. Retailers purchase goods from manufacturers. I discuss the problem of the representative consumer, that of the retailer, and that of manufacturers, and then define an equilibrium for this economy.

Consumers

Consumers' preferences are defined over leisure and a continuum of imperfectly substitutable goods purchased from retailers. Consumers sell part of their time endowment to the labor market, own shares in all firms, and trade state-contingent Arrow securities. Their problem is

$$\max_{\{c_t(i,z)\}, l_t, \mathbf{B}_{t+1}} E_0 \sum_{t=0}^{\infty} \beta^t U(c_t, l_t)$$

subject to

$$\int_0^1 \sum_{i=1}^{N} p_t(i, z) c_t(i, z)\, dz + \mathbf{Q}_t \cdot \mathbf{B}_{t+1} \le W_t l_t + \Pi_t + B_t,$$

where

$$c_t = \left(\frac{1}{N} \sum_{i=1}^{N} c_t(i)^{(\theta-1)/\theta} \right)^{\theta/(\theta-1)},$$

$$c_t(i) = \left(\int_0^1 [a_t(i, z) c_t(i, z)]^{(\gamma-1)/\gamma}\, dz \right)^{\gamma/(\gamma-1)}.$$

Here z is an index over retailers, i is an index over goods, c_t is an aggregator over the different goods, and $c_t(i)$ is an aggregator over the consumption

[11] Zbaracki et al. (2004, p. 524).

of good i purchased from the different retailers. I have in mind an economy where a good i sold by a retailer z is fairly substitutable with a good i sold by a retailer z' (Saltines at Dominick's versus Saltines at Cub Foods) so that γ is relatively high. In contrast, the different goods z are fairly imperfectly substitutable, so that θ is relatively low. I refer to an (i, z) good–retailer combination as *variety*.

I refer to $a_t(i, z)$ as the *quality* of the good. On one hand, a higher a_t increases the marginal utility from consuming that good. On the other hand, a higher a_t good is also more costly to sell, as I describe below. Changes in the quality of the good will thus make it optimal for the retailer to change its prices and will provide one source of price variation in the model, akin to the idiosyncratic productivity shocks assumed by Golosov and Lucas (2007). The assumption that idiosyncratic shocks affect both the cost at which a good is sold and the consumer's preferences for the good is made for purely technical reasons. It allows me to reduce the dimensionality of the state space and thus the computational burden.[12]

The consumer's budget constraint says that expenditure on goods and purchases of state-contingent securities must not exceed the consumer's labor income, $W_t l_t$, profits from ownership of firms, Π_t, and the returns to last period's purchases of state-contingent bonds, B_t. Here $\mathbf{B}_{t+1}$ is a vector of state-contingent securities purchased at date t and $\mathbf{Q}_t$ is a vector of prices of these securities.

The formulation above implies that demand for any individual variety is

$$c_t(i, z) = a_t(i, z)^{\gamma-1}\left(\frac{p_t(i, z)}{P_t(i)}\right)^{-\gamma}\left(\frac{P_t(i)}{P_t}\right)^{-\theta} c_t,$$

where P_t is the minumum expenditure necessary to deliver one unit of the final consumption good, c_t, and $P_t(i)$ is the price index for good i:

$$P_t(i) = \left(\int_0^1 a_t(i, z)^{\gamma-1} P_t(i, z)^{1-\gamma}\, dz\right)^{1/(1-\gamma)},$$

$$P_t = \left(\frac{1}{N}\sum_{i=1}^{N} P_t(i)^{1-\theta}\right)^{1/(1-\theta)}.$$

[12]A frequently employed alternative assumption that serves a similar purpose is that the menu-cost is proportional to $a_t^{\gamma-1}$ (e.g., Gertler and Leahy (2008)).

Manufacturers

Each variety (i, z) is produced by a perfectly competitive sector. Firms in this sector hire labor $l_t(i, z)$ and produce output according to

$$y_t(i, z) = \frac{l_t(i, z)}{e_t(i, z)},$$

where $e_t(i, z)$ is the inverse of that sector's productivity. Manufacturers sell output to retailers at a price $\omega_t(i, z)$. Perfect competition implies that manufacturer's profits are equal to 0:

$$\omega_t(i, z) y_t(i, z) - W_t l_t(i, z) = 0$$

so that

$$\omega_t(i, z) = e_t(i, z) W_t.$$

I assume that $e_t(i, z) \in \{\bar{e}, 1\}$ and is independent across retailers, but common across varieties of goods sold by a given retailer, $e_t(i, z) = e_t(z)$. A retailer will thus synchronize its temporary price changes, as is the case in the data I describe below. This variable evolves over time according to the Markov transition probability

$$\Pr(e_t(z) = \bar{e} | e_{t-1}(z) = 1) = \alpha,$$
$$\Pr(e_t(z) = \bar{e} | e_{t-1}(z) = \bar{e}) = \rho.$$

Retailers

Retailers purchase goods from manufacturers at a unit price $\omega_t(i, z)$ and sell these goods to the consumers at a price $p_t(i, z)$. The cost to the retailer of selling a good of quality $a_t(i, z)$ is $a_t(i, z)\omega_t(i, z)$. I assume that shocks to a_t are correlated across the two varieties produced by a given retailer and that the log of a_t follows a random walk:

$$\log a_t(i, z) = \log a_{t-1}(i, z) + \varepsilon_t^a(i, z).$$

The retailer's (nominal) profits from sales of good i are, therefore,

$$\Pi_t(i, z) = [p_t(i, z) - a_t(i, z)\omega_t(i, z)] a_t(i, z)^{\gamma-1} \times \left(\frac{p_t(i, z)}{P_t(i)}\right)^{-\gamma} \left(\frac{P_t(i)}{P_t}\right)^{-\theta} c_t.$$

Notice that absent price adjustment frictions, the retailer's optimal price is a constant markup over the unit cost:

$$p_t(i, z) = \frac{\gamma}{\gamma - 1} a_t(i, z)\omega_t(i, z).$$

Letting $\mu_t(i,z) = \frac{p_t(i,z)}{a_t(i,z)\omega_t(i,z)}$ denote the actual markup the retailer charges, we can write the profits from selling good i as

$$\Pi_t(i,z) = \omega_t(i,z)^{1-\gamma}[\mu_t(i,z) - 1]\mu_t(i,z)^{\gamma-1}P_t(i)^{\gamma}\left(\frac{P_t(i)}{P_t}\right)^{-\theta} c_t.$$

This expression shows that the retailer's profits from selling any individual good (conditional on its markup) are independent of the quality of the good, $a_t(i,z)$. The only effect of shocks to $a_t(i,z)$ is to alter the retailer's markup in the presence of price adjustment costs and thus provide a source of low-frequency price variation. It is this feature of the process for a_t, together with the assumption that it follows a random walk, that allows me to reduce the dimensionality of the state space.

Price Adjustment Technology

I assume that retailers set two prices: a regular price p_t^R and a posted (transactions) price p_t. I assume that the adjustment costs apply to all goods the retailer sells. Whenever at least one of its posted prices $p_t(i)$ deviates from the regular price $p_t^R(i)$, the retailer incurs a fixed cost κ. This cost is independent of the number of prices that deviate. Moreover, changing the regular price is also costly and entails a fixed cost ϕ^R. This cost is again incurred when at least one regular price is changed and is independent of the total number of regular prices that the retailer adjusts. Finally, changing posted prices is costly as well: the retailer pays a fixed cost ϕ every time at least one of its posted prices changes. Given these assumptions, the problem of any retailer z is

$$\max_{p_t(i,z),\, p_t^R(i,z)} E_0 \sum_t q_t \begin{pmatrix} \sum_{i=1}^{N}[p_t(i,z) - a_t(i,z)\omega_t(i,z)]c_t(i,z) - \\ \phi^R \times [\text{any } p_t^R(i,z) \neq p_{t-1}^R(i,z)] - \\ \phi \times [\text{any } p_t(i,z) \neq p_{t-1}(i,z)] - \\ \kappa[\text{any } p_t^R(i,z) \neq p_t(i,z)] - \end{pmatrix},$$

where $q_t = \frac{U_{c,t}/P_t}{U_{c,0}/P_0}$ is the date 0 price of a date t Arrow security. The first term in this expression denotes the profits the retailer makes and the last three terms reflect the cost associated with exercising the three different options.

Retailer Decision Rules

Kehoe and Midrigan (2008) discussed in detail the optimal decision rules in a (single-product) economy described by the assumptions I made above. Here I briefly summarize the workings of the model under the assumption that the manufacturer's cost shocks, $e_t(i,z)$, are transitory: $e_t(i,z) = 1$ most of the time, infrequently enters the low-value state $e_t(i,z) = \bar{e}$, and leaves the low-value state with high probability.

Assuming that κ, the cost of deviating from the regular price, is relatively low, the retailer will find it optimal to use this option to respond to transitory changes in its wholesale cost, $\omega_t(i, z)$. Doing so entails paying a cost $\kappa + \phi$ in the period in which it initiates a price change, a cost κ in every subsequent period in which the posted price is below the regular price, and then a final cost ϕ to return its posted price to the old regular price. As long as ϕ^R is high relative to ϕ and κ, temporarily deviating from the regular price will be less expensive than undertaking two regular price changes (one down and another one up) to respond to this transitory cost shock.

In contrast, a series of a_t shocks change the retailer's desired price permanently: the retailer will choose to respond to such shocks by undertaking a regular price change at a cost ϕ^R. The alternative would be to deviate from the regular price, but such an option is too expensive given that the retailer would have to pay a cost κ in every period in which the posted price deviates from the regular price.

Figure 3 illustrates this discussion by plotting a time series of simulated decision rules. The frequent transitory cost shocks (reflected in a lower desired price, the dotted line) lead the retailer to temporarily deviate from the regular

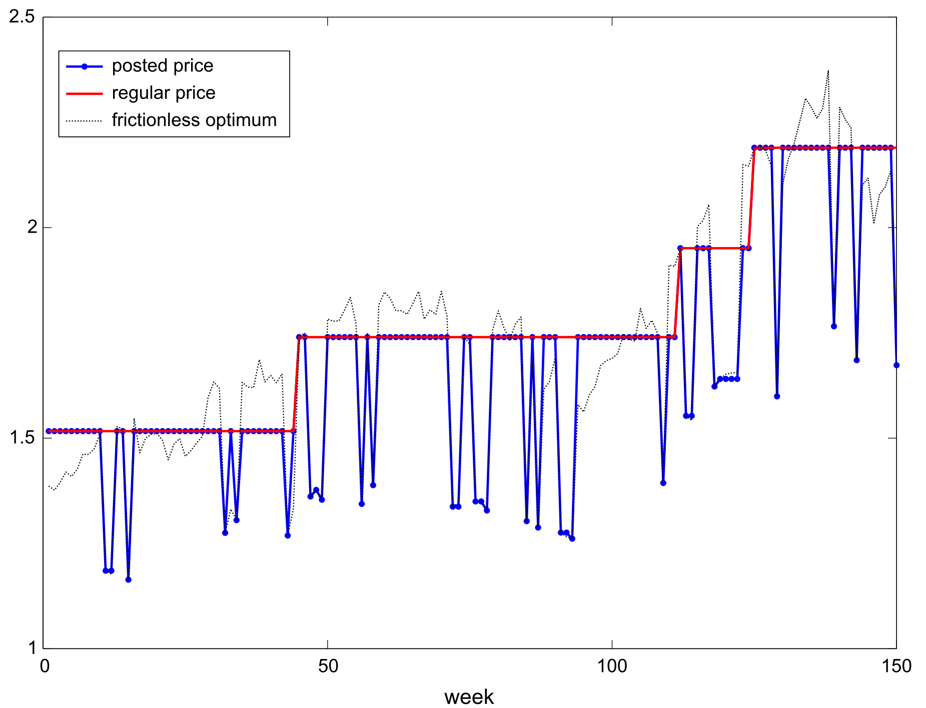

FIGURE 3.—Prices in an economy with temporary changes.

price by changing the posted price only. In contrast, when the wholesale cost is in the high state and the retailer's regular price is too far from the desired level (because of a change in a_t), the retailer finds it optimal to undertake a regular price change. Because of discounting, it never pays off to change the regular price without using it. Hence in such periods, the posted price changes as well.

Equilibrium

I assume that nominal spending must be equal to the money supply:

$$\int_0^1 \sum_{i=1}^{N} p_t(i, z)c_t(i, z)\, dz = P_t c_t = M_t.$$

The money growth rate $g_t = \frac{M_t}{M_{t-1}}$ evolves over time according to

$$\log g_t = \rho_m \log g_{t-1} + \varepsilon_t^m,$$

where ε_t^m is an independent and identically distributed (i.i.d.) $N(0, \sigma_m^2)$ disturbance.

An equilibrium is a collection of prices and allocations $p_t(i, z)$, $P_t(i)$, P_t, W_t, $c_t(i, z)$, c_t, and l_t such that, taking prices as given, allocations and prices solve the consumer, retailer, and manufacturer's problems and the goods, labor, and bond markets clear.

B. *Computing the Equilibrium*

I assume preferences of the form $U(c, l) = \log(c) - l$. This specification follows Hansen (1985) and Rogerson (1988) by assuming indivisible labor decisions implemented with lotteries. These assumptions ensure that the nominal wage is proportional to nominal spending, $W = Pc$, and thus proportional to the money supply, $W = M$. This closely follows Golosov and Lucas (2007) and ensures that shocks to the money supply translate one-for-one into changes in the nominal marginal cost and hence the retailer's desired price. I consider a richer specification (in which I allow for interest-elastic money demand and capital accumulation) below.

The quantitative analysis I report below assumes, for simplicity, that the cost of changing posted prices, ϕ, is equal to zero. This assumption considerably simplifies the notation below, but has little effect on my results. Appendix 3 in the Supplemental Material studies the economy with a nonzero cost of changing posted prices, $\phi > 0$, and shows that calibration of the model to the data indeed requires a very small cost of changing posted prices.[13]

[13]The only feature of the data that a model with $\phi = 0$ misses is the fact that prices are sometimes sticky even during a sale. However, in the data, conditional on the retailer undertaking a

The assumption that $\phi = 0$ implies that the retailer's state is fully characterized by the markup of its existing regular prices, $p_{t-1}^R(i)$, over the unit cost of selling the good, $a_t(i)\omega_t(i)$, as well as the cost, to the manufacturer, of producing the good, $e_t(i)$. Let

$$\mu_{t-1}^R(i) = \frac{p_{t-1}^R(i)}{a_t(i)\omega_t(i)}.$$

The aggregate state of this economy is characterized by the growth rate of money, g, and the distribution of retailers' markups, $\boldsymbol{\mu}_{-1}^R = \{\mu_{i,-1}^R\}$, and efficiency levels, e. Let Λ denote this distribution.

Let $V^R(\boldsymbol{\mu}_{-1}^R, e; g, \Lambda)$, $V^T(\boldsymbol{\mu}_{-1}^R, e; g, \Lambda)$, and $V^N(\mu_{-1}^R, e; g, \Lambda)$ denote a firm's value of (i) adjusting its regular and posted price (and selling at $p_t(i) = p_t^R(i) \neq p_{t-1}^R(i)$ in that period), (ii) undertaking a temporary price change and leaving its regular price unchanged, $p_t(i) \neq p_t^R(i) = p_{t-1}^R(i)$, and (iii) selling at the old regular price, $p_t(i) = p_t^R(i) = p_{t-1}^R(i)$. Letting $V = \max(V^R, V^T, V^N)$ denote the envelope of these three options, the following system of functional equations characterizes the retailer's problem (I impose here the log-utility specification above and substitute several equilibrium conditions to simplify the notation):

$$\begin{aligned}
V^R(\boldsymbol{\mu}_{-1}^R, e; g, \Lambda) = \max_{\mu_i^R}\Bigg(&\sum_{i=1}^N e_i^{1-\gamma}(\mu_i^R - 1)\mu_i^{R,-\gamma}\hat{P}^{\gamma-1} - \phi^R \\
&+ \beta E V(\boldsymbol{\mu}_{-1}^{R\prime}, e'; g', \boldsymbol{\Lambda}')\Bigg), \\
V^T(\boldsymbol{\mu}_{-1}^R, e; g, \Lambda) = \max_{\mu_i}\Bigg(&\sum_{i=1}^N e_i^{1-\gamma}(\mu_i - 1)\mu_i^{-\gamma}\hat{P}^{\gamma-1} - \kappa \\
&+ \beta E V(\boldsymbol{\mu}_{-1}^{R\prime}, e'; g', \boldsymbol{\Lambda}')\Bigg), \\
V^N(\boldsymbol{\mu}_{-1}^R, e; g, \Lambda) = \Bigg(&\sum_{i=1}^N e_i^{1-\gamma}(\mu_{i,-1}^R - 1)\mu_{i,-1}^{R,-\gamma}\hat{P}^{\gamma-1} - \kappa \\
&+ \beta E V(\boldsymbol{\mu}_{-1}^{R\prime}, e'; g', \boldsymbol{\Lambda}')\Bigg),
\end{aligned}$$

sale two periods in a row, the sale price changes from one period to another 67% of the time. Thus the model requires a very low ϕ to account for the fact that sales prices are, in fact, quite flexible. (This feature of the data has also been documented by Nakamura and Steinsson (2009).)

where $\hat{P} = \frac{P}{M}$ denotes the aggregate price level detrended by the money supply.

The law of motion for the aggregate state is $\Lambda' = \Gamma(g, \Lambda)$. As for idiosyncratic states, $\mu_{i,-1}^{R\prime} = \mu_i^R \frac{e}{g' \exp(\varepsilon_{i,a}')e'}$ if the retailer adjusts its price and $=$ $\mu_{i,-1}^R \frac{e}{g' \exp(\varepsilon_{i,a}')e'}$ otherwise. Here the retailer's markup is eroded by the three types of shocks: aggregate shocks to the money supply, g, permanent shocks to the quality of the two goods, a, and transitory shocks to its wholesale cost, e.

Solving for the equilibrium in this economy requires characterizing the objects $V^R(\cdot)$, $V^T(\cdot)$, and $V^N(\cdot)$, as well as $\Gamma(\cdot)$, the law of motion for the distribution of markups. To solve this system of functional equations, I follow an approach developed by Krusell and Smith (1998)[14] that restricts the aggregate state space to g, the growth rate of the money stock, and a single moment of the distribution Λ. In particular, I use $\hat{P}_{-1}$, the last period's aggregate price level (detrended by the money stock). The latter is a sufficient state variable in a (log-linearized) version of the Calvo pricing model. I show below that this state variable alone characterizes the evolution of aggregate variables in this menu-cost economy quite well.

To implement the Krusell and Smith (1998) approach, I assume that the aggregate price level is a log-linear function of the two aggregate state variables

$$\log \hat{P}_t = \varsigma_0 + \varsigma_1 \log g_t + \varsigma_2 \log \hat{P}_{t-1}.$$

Given a guess for the coefficients in this expression, I solve the retailer's problem using projection methods and a combination of cubic and linear spline interpolants.[15] I then simulate retailer decision rules and use the simulated data to reestimate the coefficients in the postulated law of motion. These updated coefficients are used to recompute firm decision rules, simulate new data, and update the coefficients. Once these coefficients converge, the distance between actual (in simulations) and predicted (by the coefficients in the aggregate functions) aggregate time series is insignificant (the out-of-sample forecasts have an R^2 in excess of 99.9%). Moreover, adding higher-order moments of Λ adds little precision.[16] Notice that the accuracy of this first-moment approach is not inconsistent with my argument that higher-order moments of the distribution of desired price changes matter for the economy's aggregate implications. The precision of the simple pricing rules reflects the fact that monetary disturbances induce little time series variation in these higher-order moments in the model.[17] This is because the distribution of desired price

[14]See also Klenow and Willis (2006) and Khan and Thomas (2007) for applications of this approach to models with nonconvexities.

[15]See Miranda and Fackler (2002) for a detailed description of these methods as well as a toolkit of routines that greatly facilitate their implementation. The solution method is described in some more detail in the Supplemental Material.

[16]These coefficients and the R^2 are reported in Table III.

[17]For example, the kurtosis of μ_{-1}^R in simulations ranges from a 10th percentile of 3.28 to a 90th percentile of 3.32.

changes is mainly pinned down by the distribution of exogenous idiosyncratic shocks.

The existence and continuity of the value functions can be established using standard theorems (Stokey and Lucas (1989)). Although these value functions are not concave, it can be shown that they satisfy K-concavity, a property introduced by Scarf (1959). In turn, K-concavity guarantees uniqueness of the optimal price functions. Aguirregabiria (1999) proved K-concavity in the context of a model similar to the one presented here in which the two control variables are each subject to a fixed cost of adjustment. Sheshinski and Weiss (1992) studied a special case of the economy presented here and also proved uniqueness of the optimal decision rules.

C. *Calibration and Parameterization*

I assume that the period is a week (the frequency at which Dominick's sets prices). I set $\beta = 0.96^{1/52}$ and $\gamma = 3$, a typical estimate of elasticity of substitution in grocery stores.[18] As I show below, this value of γ allows the model to match well the response of quantities to temporary price cuts in the data. In particular, the model will be shown to account for the fraction of goods sold in periods with sales. I set $\theta = 1$, consistent with the idea that different goods sold by a particular retailer are relatively less substitutable. However, the exact choice of θ plays little role since there are no aggregate good-specific shocks in the model. Finally, I assume that retailers sell $N = 2$ goods each.

I allow for persistence in the growth rate of money given that most applied work assumes persistence in monetary policy.[19] I calibrate the coefficients in the money growth rule by first projecting the growth rate of (monthly) $M1$ on current and 24 lagged measures of monetary policy shocks.[20] I then fit an (autoregression) AR(1) process for the fitted values in this regression and obtain an autoregressive coefficient of 0.61 and standard deviation of residuals of $\sigma_m = 0.0018$. I then adjust these parameters to reflect the weekly frequency in my model economy.

The rest of the parameters are calibrated to allow the model to match the microprice facts documented in the earlier section. I follow Gertler and Leahy

[18]Nevo (1997), Barsky, Bergen, Dutta, and Levy (2000), and Chevalier, Kashyap, and Rossi (2003).

[19]Christiano, Eichenbaum, and Evans (2005) reported that the growth rate of money increases persistently in response to an (identified) exogenous monetary policy shock and postulated a process for μ_t that is well approximated by an AR(1) with a (quarterly) persistence coefficient of 0.5. Alternatively, Smets and Wouters (2007) assumed an interest rate rule and also found evidence of inertia: their estimate of the coefficient on lagged interest rates in the interest rate rule is 0.8.

[20]The results reported below use a new measure of shocks due to Romer and Romer (2004) that is available for 1969–1996. I also use the measure used by Christiano, Eichenbaum, and Evans (2005) and find very similar results. I thank Oleksiy Kryvtsov for sharing the Christiano, Eichenbaum, and Evans (2005) data with me.

(2008) and assume that the permanent shocks to the good's quality, a_t, arrive infrequently, according to a Poisson process[21]:

$$\tilde{\varepsilon}_{it} = \begin{cases} 0 & \text{with probability } = 1 - p^a, \\ N(0, \sigma_a^2) & \text{with probability } = p^a. \end{cases}$$

Conditional on their arrival, the shocks are drawn from a Gaussian distribution with standard deviation σ_a. My specification of the process for shocks nests that of Golosov and Lucas (2007) ($p^\alpha = 1$), but is more flexible and allows me to account for the distribution of the size of price changes in the data.

I assume that the permanent (quality) shocks are correlated across the goods sold by a given retailer. In particular, the actual innovations to a good's quality, ε_{it}, depend on the underlying draws as $\varepsilon_{it} = \tilde{\varepsilon}_{it} + \chi \,\text{mean}(\tilde{\varepsilon}_{it})$, where χ is a parameter that governs the correlation of productivity shocks across the two goods.

The eight parameters that I calibrate are σ_a—the standard deviation of permanent shocks, χ—the parameter governing the correlation of quality shocks across the two varieties sold by a retailer, α—the parameter governing the frequency of shocks to the manufacturer's cost, ρ—the persistence of the low manufacturer cost state, $\bar{e}$—the manufacturer's cost in the low state (recall that the high state is normalized to 1), ϕ^R—the fixed cost of changing the regular prices, κ—the cost of deviating from the regular price, and p^a—the frequency with which retailers experience a shock to the good's quality.

I choose these parameters so as to match the salient properties of the micro-price data, specifically, moments of the distribution of the size of price changes, as well as statistics that capture the relative importance of high- versus low-frequency price variation. I discuss these moments below. Specifically, in addition to the moment emphasized by Golosov and Lucas, the large mean size of price changes, I ask the model to capture (i) the large heterogeneity in the size of price changes, (ii) the frequency of posted and regular price changes, and (iii) the fact that after a sale, prices typically revert to the preexisting regular price.

Data Moments

The moments I use to calibrate the model are listed in the data column of Table I. Some of these moments were computed using the algorithm I discussed in the previous section. I emphasize that I apply an identical algorithm to the *posted* (transactions) price in the model and in the data so as to report statistics about regular prices. For example, when I report the frequency of "regular price changes," this reflects the frequency of changes in the regular price

[21]This assumption is a simpler alternative to the mixture of betas I have used in earlier work that produces very similar results.

TABLE I

MOMENTS IN MODEL AND DATA[a]

Moments	Data	Benchmark	No Temp. Changes	Golosov–Lucas
Used in calibration				
1. Fraction of prices at annual mode	0.58	0.58	0.58	0.58
2. Frequency of price changes	0.34	0.32	0.053	0.046
3. Frequency regular price changes	0.029	0.025	0.045	0.044
4. Fraction of price changes that are temporary	0.97	0.96	0.29	0.06
5. Probability a temporary price spell ends	0.47	0.45	0.41	0.30
6. Probability temp. price returns to old regular	0.86	0.89	0.00	0.00
7. Fraction of periods with temporary prices	0.25	0.22	0.02	0.00
8. Fraction of periods with sales	0.22	0.22	0.01	0.00
9. Fraction of goods sold when sales	0.37	0.40	0.01	0.00
10. Mean size of price changes	0.20	0.20	0.11	0.11
11. Mean size of regular price changes	0.11	0.10	0.11	0.11
12. 10th percentile size regular price changes	0.03	0.03	0.03	0.10
13. 25th percentile size regular price changes	0.05	0.05	0.05	0.10
14. 50th percentile size regular price changes	0.09	0.09	0.09	0.11
15. 75th percentile size regular price changes	0.13	0.14	0.15	0.11
16. 90th percentile size regular price changes	0.21	0.20	0.22	0.12
17. Mean abs(fup − fdn) within store	0.89	0.91	0.94	–
Additional moments				
18. Frequency changes annual mode	0.61	0.73	0.89	0.89
19. Fraction prices at quarterly mode	0.70	0.73	0.86	0.87
20. Frequency changes quarterly mode	0.32	0.29	0.47	0.47
21. Std. dev. size of price changes	0.18	0.15	0.08	0.01
22. Kurtosis price changes	3.15	1.65	2.97	1.06
23. Fraction changes $< 1/2$ mean	0.36	0.36	0.28	0.00
24. Fraction changes $< 1/4$ mean	0.19	0.34	0.07	0.00
25. Std. dev. size of regular price changes	0.08	0.07	0.08	0.01
26. Kurtosis regular price changes	4.02	2.53	3.13	1.07
27. Fraction regular price changes $< 1/2$ mean	0.25	0.28	0.28	0.00
28. Fraction regular changes $< 1/4$ mean	0.08	0.07	0.07	0.00
29. Fraction of price changes during a sale (within)	0.67	0.94	0.08	0.01
30. Fraction of changes in the sale price (between)	0.82	1	1	1

[a]Moments not included in the criterion function in the calibration are underlined.

identified by the algorithm, not the frequency with which retailers exercise the option of the regular price change in theory.[22]

Table I (rows 2–4) shows that 33% of goods experience a price change in any given week. Most of these price changes reflect temporary price discounts.

[22]The two are very similar, since the algorithm does a very good job identifying changes in the theoretical regular price.

The fraction of price changes that are temporary (in periods in which the actual price is not equal to the regular price) is 0.965. As a result, only 2.9% of goods experience a regular price change in any given week.

Also notice (rows 5–9) that the fraction of periods in which the store charges a temporary price is equal to 0.25, and most (0.22 of the weeks) of these are periods with sales (the posted price is below the regular price). Sales are also periods in which the retailer sells a disproportionately larger amount of its goods: 37% of all goods are sold during sales.

Temporary price changes are very transitory: conditional on the price being temporary at date t, it returns to a regular price the next week 47% of the time. When temporary changes do return, they do so to the old regular price 86% of the time.

Rows 1 and 18–20 report additional statistics that capture the pattern of low-frequency price variation that are independent of the algorithm I use. The annual mode for any particular good accounts for 58% of all prices the retailer charges during the year.[23] The annual mode also shows quite a bit of stickiness and changes from one year to another only 61% of the time. As Kehoe and Midrigan (2008) showed standard menu-cost models cannot account, simultaneously, for the high frequency of price changes in the data as well as these alternative measures of low-frequency nominal stickiness. For example, the standard model, when calibrated to match a frequency of price changes of 34% per week, predicts a fraction of prices at annual mode of only 23%, much less than the 58% in the data. The model I study here will be shown to fit all of these facts very well.

I finally focus on the size of price changes (rows 10–16, 21–28), and again distinguish between changes in the posted price and changes in the regular price. A key feature of the data is that there is large heterogeneity in the size of price changes. One concern is that this heterogeneity reflects permanent good-level heterogeneity. To address these concerns, I report, following Klenow and Kryvtsov (2008), moments of the "standardized" distribution of price changes. In particular, I scale each price change by each good's mean size of price changes, where a "good" represents a given manufacturer $\times$ product category. By construction, these data are free of good-specific heterogeneity.

As in Klenow and Kryvtsov (2008), the mean size of price changes is very large: the retailer in question adjusts posted prices by 20% and regular prices by 11% on average.[24] The standard deviation of price changes is, however, large as well: 18% for posted prices and 8% for regular prices.

Given my focus on the distribution of the size of price changes, rows 12–16 report several percentiles of this distribution. I focus on regular prices because,

[23] Hosken and Reiffen (2004), Kehoe and Midrigan (2008) and more recently Eichenbaum, Jaimovich, and Rebelo (2011) have also used modal statistics to characterize microprice data.

[24] Because sales make up most of posted price changes, the statistics for sale prices are very similar to those for posted prices, and are not reported here.

as shown by Kehoe and Midrigan (2008), it is the pattern of low-frequency price variation that mostly matters for the aggregate predictions of a menu-cost economy. Notice that 10% of all regular prices are less than 3% in absolute value, while 25% of all regular prices are less than 5% in absolute value. A lot of price changes are thus very small. Similarly, many price changes are also very large: the 75th percentile of this distribution is equal to 13%, while the 90th percentile is equal to 21%.

Price changes within the store are strongly synchronized, especially for goods in the same product category, as I document in a subsequent section. Moreover, price changes also tend to be of the same sign. Table I (row 17) shows that the average (absolute value) of the difference between the fraction of regular price increases (fup) and decreases (fdn) in any given week is 0.89. This suggests that idiosyncratic shocks to the different goods are strongly correlated and I use this statistic to pin down the size of this correlation.

Finally, rows 29 and 30 report that sales prices are highly flexible in the data. Conditional on the good being on sale and not returning to a regular price next period, the probability that the sale price changes is equal to 67%. Such price changes are large (18% on average) and dispersed (the standard deviation is 16%), just like all other price changes. Similarly, there is very little price rigidity in the sale price across different sales episodes: 82% of the time the store charges a price during a sale that is different than the price it has charged during the previous sale. The model is able to capture this flexibility, since I have assumed that the cost of changing posted prices, ϕ, is equal to 0.

Criterion Function

The criterion function I use to pin down the model's parameters is the sum of the squared deviations of the moments in the model from those in the data, those listed in rows 1–17 in Table I. I target the frequency and average size of posted and regular price changes, the fraction of periods with temporary price changes, the fraction of goods sold during sales, and the fraction of prices at annual mode. In addition, I include the percentiles of the distribution of the size of regular price changes and require the model to account for the fact that price changes are typically of the same sign within a store.

Benchmark Model

Table II reports the parameter values used in the model (second column, benchmark). Quality (a_t) shocks arrive infrequently, with $p^a = 0.030$. The volatility of these shocks, σ_a, is equal to 0.08 and thus fairly large. Shocks to the quality of a good sold by a given retailer are strongly correlated, $\chi = 0.84$, implying a correlation of shocks equal to 0.53. Manufacturers' costs change more frequently. The likelihood that the cost transits from a high to a low state is $\alpha = 0.126$. The low-cost state is much less persistent: the probability of transiting back to the high-cost state is $\rho = 0.524$. Thus, retailers' wholesale costs

TABLE II

PARAMETER VALUES

	Benchmark	No Temp. Changes	Golosov–Lucas	BLS Calibration
Calibrated				
σ_a	0.080	0.100	0.023	0.112
ϕ^R, relative to SS revenue	0.022	0.018	0.038	0.012
χ	0.845	0.999	–	0.466
p^a	0.030	0.040	–	0.079
α	0.126	–	–	
ρ	0.524	–	–	
e	0.741	–	–	
κ, relative to SS revenue	0.012	–	–	
Assigned				
γ	3	3	3	3
β (annual)	0.96	0.96	0.96	0.96
ρ_m	0.884	0.884	0.884	0.610
σ_m (%)	0.032	0.032	0.032	0.180

experience frequent, but temporary, price declines, leading retailers to offer temporary price discounts.

The menu costs (expressed as a fraction of steady-state (SS) revenue) are $\phi^R = 0.022$ and $\kappa = 0.012$. These are the costs the retailer pays every time it makes a regular price change or deviates from the regular price. Given that these events are relatively infrequent, adjustment costs are a much smaller share of the overall (across all periods) revenue of the retailer: the total resources used up by price changes are equal to 0.34% of revenue. For comparison, Levy et al. (1997) reported menu costs as large as 0.7% of revenue, while Zbaracki et al. (2004) reported price adjustment costs as large as 1.2% of revenue. The menu costs I estimate are thus by no means large relative to direct estimates in earlier work.

The second column of Table I reports the moments in the model. The model does well at accounting for both high- and low-frequency price variation in the data. As in the data, posted prices change frequently (32% of the weeks), while regular prices much less so (2.5% of the weeks). Temporary price spells are transitory and last roughly 2 weeks. Most periods of temporary prices are periods with discounts (a good is on sale 22% of the time), and periods of sales account for a disproportionate amount of goods sold (40%). Finally, as in the data, temporary price changes revert to the old regular price most (86%) of the time. As a result, the frequent price changes in the model do not impart much low-frequency price variation: the retailer's prices are equal to the annual mode 58% of the time, as in the data.

The model also does a good job at reproducing the distribution of the size of regular price changes in the data. In particular, many price changes are small in absolute value: the 10th percentile is equal to 0.03 and the 25th percentile is equal to 0.05. Similarly, many price changes are large in absolute value: the 75th percentile is equal to 0.14, while the 90th percentile is equal to 0.20. All these numbers are very close to their counterparts in the data. Finally, the model accounts well for the mean size of regular and posted price changes, the key statistics emphasized by Golosov and Lucas.

The fit is also good vis-à-vis the other moments that were not directly used in the criterion function: the fraction of prices at the quarterly mode and the frequency with which the quarterly mode changes, as well as the standard deviation and fraction of small posted and regular price changes. The model understates the kurtosis of price changes in the data (2.6 vs. 4 for regular prices), but since the kurtosis is very sensitive to a few outliers in the tails, I use the much more robust percentiles of the distribution to calibrate the model.[25]

To summarize, the model reproduces well salient features of the microprice data. This itself is not surprising, since I have explicitly targeted these facts when choosing parameters. Recall, however, that the question I ask in this paper is, "What are the aggregate implications of a model that accounts for the microlevel facts?" Since I have shown that my model is indeed consistent with the microdata, I can now study its aggregate implications. Before I do so, I discuss a number of additional experiments I conduct to isolate the role of the assumptions I have made above.

Model Without Temporary Changes

I report results from an economy in which the technology for changing prices is the same as in the standard menu-cost model. The retailer incurs a cost, ϕ, every time it changes its menu of posted prices. Permanent shocks to quality are the only source of idiosyncratic uncertainty. The difference between this economy and the benchmark economy is the absence of a motive for temporary price changes.

I target the same set of moments as above, with the exception of those that characterize the low- versus high-frequency price variation (underlined in Table I). I choose the size of the menu-cost, ϕ, so that the model matches the fraction of prices at the annual mode. Targeting this statistic, Kehoe and Midrigan argued, is a more accurate method of comparing economies with and without temporary price changes. The alternative—targeting the frequency of posted price changes—would underpredict the real effects of money in the benchmark model, since most of these changes are temporary and do not allow retailers to respond to low-frequency changes in monetary policy. Similarly, targeting the frequency of regular price changes would overpredict the real effects of money, since temporary prices do respond to monetary policy shocks, albeit for only a short period. In contrast, the fraction of prices at the annual model is a robust

[25]I thank an anonymous referee for the suggestion to no longer focus on this feature of the data.

measure of low-frequency price variation in the model and more appropriate for measuring the response of prices to monetary shocks.

Once again, the model does a good job capturing the heterogeneity in the size of price changes in the data. In contrast, it accounts less well for the pattern of temporal price variation. The frequency of posted price changes is only 0.053, much lower than in the data, while that of regular price changes (as identified by the algorithm) is equal to 0.045, higher than in the data.

Golosov–Lucas Model

I also report results from a version of the model without temporary price changes and in which the distribution of permanent shocks is Gaussian as opposed to Poisson, and there are no economies of scope. I choose the fixed cost to match the fraction of prices at the annual mode. The only additional parameter that remains to be chosen is the standard deviation of idiosyncratic shocks, which I choose, as did Golosov and Lucas, so as to match the average size of regular price changes.

This economy fails to account for both the pattern of temporal price variation and the heterogeneity in the size of price changes. Price changes are now much more concentrated near the adjustment thresholds: the interquartile range is now only 2% (10% in the data). I show below that the model's failure along this dimension explains the much lower real effects from monetary shocks that it predicts.

Calvo Model

I finally compare the predictions of the menu-cost economies to those of the Calvo model, as did Golosov and Lucas. I choose the adjustment hazard $1 - \lambda = 0.055$ so as to match the fraction of prices at the annual mode.[26]

4. AGGREGATE IMPLICATIONS

I next turn to the models' aggregate implications. My measures of the real effects of money are the volatility and persistence of HP-filtered consumption. I construct a measure of monthly consumption, as in the data, for comparability with other work.

A. *Business Cycle Statistics*

Table III shows that the standard deviation of consumption is 0.29% in the benchmark setup, 0.31% in the economy without temporary price changes, and as low as 0.07% in the Golosov and Lucas economy. In contrast, the Calvo

[26]The alternative, of choosing λ so as to match the frequency of price changes in the menu-cost economy, produces similar results, since the frequency of price changes in the economy without temporary price changes is equal to 0.053. I adopt this alternative approach in a robustness section below.

TABLE III

AGGREGATE IMPLICATIONS

	Calvo	Benchmark	No Temp. Changes	Golosov–Lucas
Business cycle statistics[a]				
$\sigma(C)$ (%)	0.35	0.29	0.31	0.07
Serial correlation C	0.93	0.93	0.92	0.84
Inflation dynamics				
var(π) intensive margin	1.00	0.90	0.96	0.99
correlation (π, fractionally adjusted)	–	0.65	0.40	0.16
Law of motion P[b]				
$P(t-1)$	0.946	0.947	0.932	0.816
$g(t)$	−0.668	−0.544	−0.652	−0.270
R^2	1	0.9997	0.9998	0.9808

[a]Business cycle statistics are computed for data sampled at monthly frequency. An HP (14400) filter is applied to the data.
[b]Laws of motion are computed at the actual frequency (weekly) in the model.

model predicts a standard deviation of 0.35%. Thus, my benchmark economy that accounts for the salient features of the microprice data produces real effects of money that are 80% as large as in the Calvo model. In contrast, the Golosov and Lucas economy produces real effects of money that are 20% as large as in the Calvo model.

Table III also reports the serial correlation of HP-filtered consumption, another measure of the real effects of money. Once again, the menu-cost models that do account for the heterogeneity in the size of price changes produce persistence that is as high as in the Calvo model (0.93). In contrast, the Golosov and Lucas model features somewhat less persistent business cycles (0.84).

All of these economies produce little synchronization of price changes in response to monetary shocks: the fraction of inflation variance accounted for by the intensive margin (the mean size of price changes) is in excess of 90% (91% as reported by Klenow and Kryvtsov (2008) for U.S. data). This suggests that the extensive margin (the fraction of adjusting firms) accounts for little of the inflation variability, despite the fact that it is positively correlated with inflation (the correlation ranges from 0.16 in the Golosov–Lucas model to 0.65 in the benchmark model vs. 0.25 in the U.S. data). In both of these economies, idiosyncratic shocks are large and adjustment decisions are driven mostly by idiosyncratic rather than aggregate shocks. Notice also that these features of the data cannot be reproduced by a model in which (a) there are no price adjustment costs and (b) all shocks, including aggregate shocks, arrive infrequently. In such a model, prices would indeed change infrequently, only in periods with shocks. However, all prices would change in periods with money shocks and the model would predict, counterfactually, perfect synchronization of price changes at the aggregate level.

I next explain the intuition behind the aggregate implications of the economies I study. Two separate forces account for the somewhat smaller variability of aggregate consumption in my benchmark economy vis-à-vis the Calvo setup. First, the menu-cost model features a selection effect due to the endogenous timing of price changes. Second, a large number of goods are sold during periods of sales in which prices are flexible and respond one-for-one to monetary shocks. To see the latter effect, notice that the price level for any individual good is

$$P_t(i) = \left(\int_0^1 \omega_t(i, z)^{1-\gamma} \mu_t(i, z)^{1-\gamma} \, dz \right)^{1/(1-\gamma)} M_t.$$

If, in response to a monetary shock, the retailer does not reset its price, its markup, μ_t, decreases, and hence the price index does not fully respond to the change in M_t. The weight, in the price index, on each individual price setter, depends on the amount sold by retailers, and hence is inversely proportional to its cost, as capturer by the $\omega_t(i, z)^{1-\gamma}$ term. Because the low-cost retailers undertake temporary price discounts and thus react to monetary shocks, the aggregate price level becomes more flexible than in the absence of such temporary price changes. In the limit, if the only retailers that sell goods are those that sell on sale, the price level would respond one-for-one to monetary shocks.

The experiments I report above allow me to disentangle the relative importance of these two effects. The difference between the benchmark economy and the economy without temporary price changes is the effect of sales. Clearly, this effect is fairly small (0.29% vs. 0.31%), as pointed out by Kehoe and Midrigan (2008), who addressed this issue in much detail. The reason for the small difference is that both economies are characterized by similar low-frequency variability of prices. Even though prices change frequently in the benchmark economy, they usually return to the pre-sale price and respond to monetary policy shocks only temporarily. Hence, the stickiness of the regular price (at which almost 67% of all goods are sold) generates a nontrivial amount of stickiness in the aggregate price level.

The difference between the real effects of money in the economies without temporary price changes and the Calvo model is solely accounted for by the endogenous timing of price changes in menu-cost economies. Clearly, accounting for the heterogeneity in the size of price changes has an important effect on the real effects of money the model predicts. I explain why this is the case below.

B. *Selection Effect*

I next focus on the key result of this paper: that the real effects of money are much greater in an economy with heterogeneity in the size of price changes. The key difference between my economy and that of Golosov and Lucas is what Golosov and Lucas (2007) refer to as the *selection effect* and Caballero and

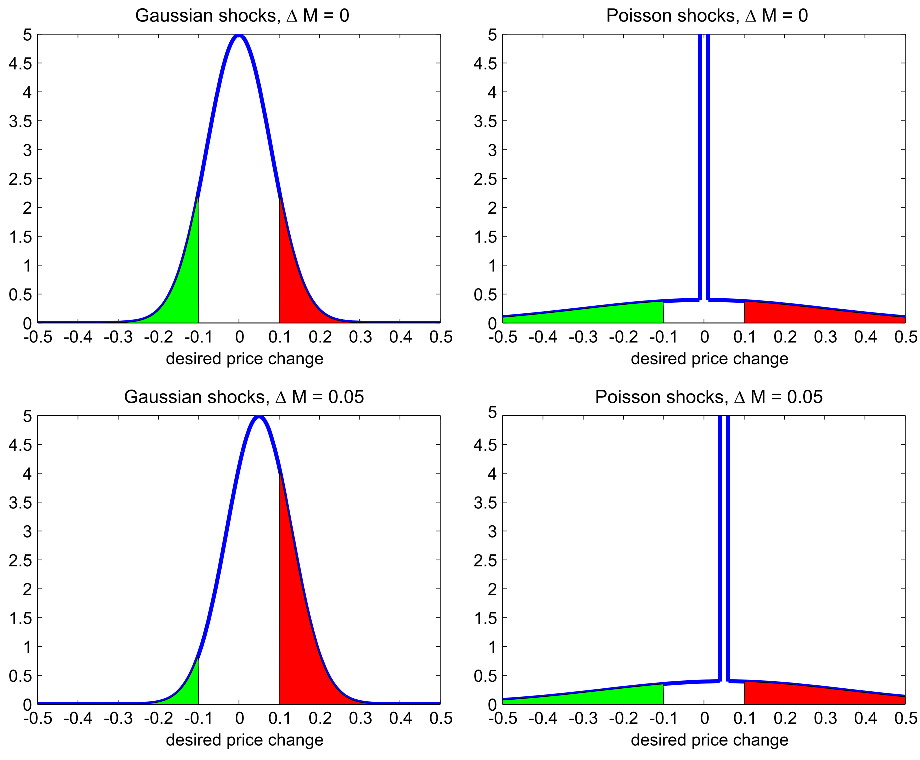

FIGURE 4.—Selection effect.

Engel (2007) refer to as the *extensive margin effect.* The endogenous timing of price changes implies that the mix of adjusters varies with the aggregate shock: in times of monetary expansion, adjusters are mostly firms whose idiosyncratic state is such that they need to raise prices. The strength of this effect critically depends on the shape of the distribution of desired price changes.

I illustrate this using a heuristic example in Figure 4. The upper panels show the distribution of desired price changes in the absence of a monetary shock, in an economy with Gaussian (left panel) and Poisson shocks (right panel).[27] In both cases I assume that the adjustment thresholds are equal to ± 0.1: the shaded areas thus reflect the distribution of actual price changes. Notice how the distribution of the absolute size of actual price changes is much more dispersed in the Poisson economy and much more concentrated in the Gaussian economy.

[27]This is an extreme example. In the benchmark economy, the distribution of desired price changes has no mass point at 0 because of the correlation in cost shocks across goods, as well as the fact that past shocks (aggregate and idiosyncratic to which the retailer has not yet responded) spread the measure of desired price changes away (but close) to 0.

Absent a money shock, the distribution of desired price changes is symmetric: half of the retailers raise and lower their prices. After a monetary shock, the distribution of desired price changes shifts to the right: retailers would like, on average, to increase their prices. Notice in the lower panel of the figure that in the Gaussian economy, this shift in the distribution has a large impact on the distribution of retailers that change their prices: many more retailers are now increasing prices. This sizeable effect on the distribution arises from the fact with Gaussian shocks, the measure of retailers in the neighborhood of the adjustment thresholds is large. Hence changes in the money supply have a large effect on the identity of adjusting firms and thus on the aggregate price level.

This selection effect is much weaker in the Poisson economy. Because distribution of price changes is much more dispersed, the measure of retailers in the neighborhood of adjustment thresholds is much smaller. Hence, the identity of retailers that adjust prices is virtually unchanged by the monetary shock: almost as many retailers find it optimal to lower their prices as before because the idiosyncratic shocks are large and the monetary shock does little to offset them. In the Dominick's data, the distribution of price changes does not change much from periods of high to low inflation, as shown in Figure 5,[28] thus providing additional support for the Poisson economy.

This simple example abstracts from two features in the benchmark economy I study: (i) economies of scope generate many price changes that are very small and (ii) the adjustment thresholds are closer to 0 because a smaller menu cost is required to account for the frequency of price changes. Economies of scope flatten the adjustment hazard and thus weaken the strength of the selection effect even further. In the limit, if the number of goods sold by a retailer is sufficiently large, a good's adjustment hazard is independent of its own desired price change; hence the economy resembles the constant-hazard Calvo setup much more. Similarly, smaller adjustment thresholds weaken the selection effect because most of the price increases caused by the monetary shock are small.

All of these ideas can be summarized by recognizing that the response of the price level to a monetary shock in this economy can be approximated[29] (in the limit, as $\Delta m \to 0$) by

$$\frac{\Delta p}{\Delta m} = \int_x f(x)h(x)\,dx + \int_x x f(x)h'(x)\,dx,$$

where $f(x)$ is the distribution of desired price changes and $h(x)$ is the adjustment hazard. The first term represents the fraction that change their prices—

[28] I thank an anonymous referee for suggesting this test of the theory. Inflation is measured by the change in the CPI for food and beverages. High (low) corresponds to periods with inflation above (below) the mean. The figure reports the distribution of regular price changes; that for all price changes is very similar.

[29] See Caballero and Engel (2007) for a derivation. Also see Burstein and Hellwig (2008) for a similar decomposition.

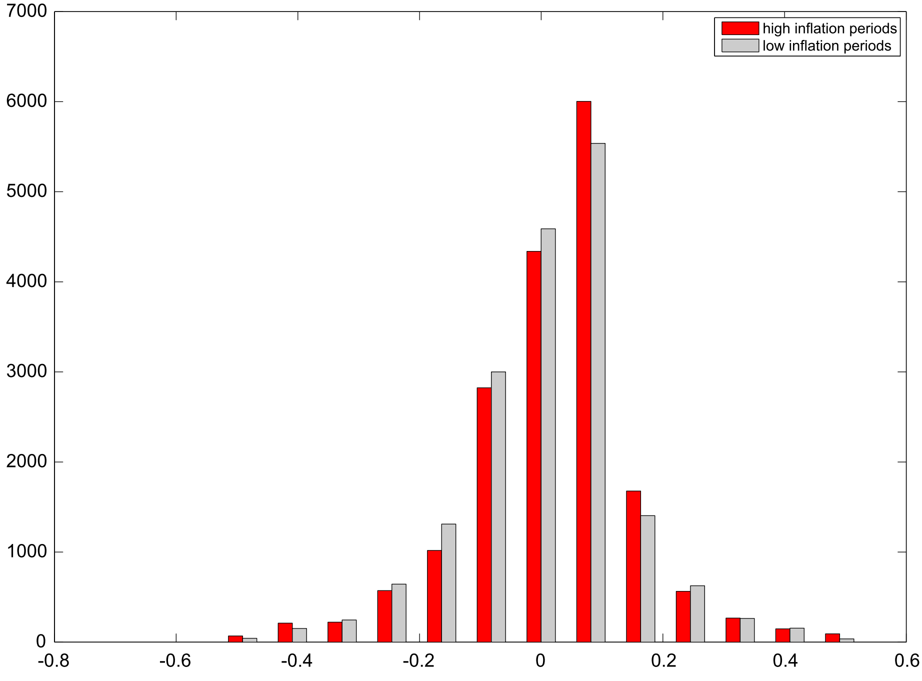

FIGURE 5.—Distribution of (regular) price changes in high and low inflation periods.

present in Calvo and menu-cost economies alike. The second term captures the selection effect: variation in the adjustment hazard, $h(x)$. This effect is larger when there is more mass, $f(x)$, in the region of increasing hazard $h'(x)$ (more marginal firms) and when the marginal firms need larger price changes, x.

C. *Additional Experiments*

I next perform two experiments to gauge the sensitivity of my results. In particular, I study (i) a variation of my model without temporary price changes calibrated to the economy-wide data from the Bureau of Labor Statistics (BLS),[30] and (ii) an extension of the model with capital, interest-elastic money demand, and a roundabout input–output structure, again calibrated to the aggregate data.

BLS Data

Table IV reports several moments of the distribution of the size of price changes in the BLS data. These statistics were computed by Nakamura and

[30]See Kehoe and Midrigan (2008), who studied the BLS data through the lens of the economy with posted and regular prices.

TABLE IV
MOMENTS IN MODEL AND BLS DATA

Moments	Data	Model With Scope Economies	Model With Random Menu Cost
Used in calibration			
1. Fraction of prices at annual mode	0.75	0.75	0.75
2. Mean size of regular price changes	0.11	0.09	0.09
4. 25th percentile size regular price changes	0.03	0.03	0.03
5. 50th percentile size regular price changes	0.07	0.07	0.09
6. 75th percentile size regular price changes	0.13	0.13	0.16
Additional moments			
1. Frequency of price changes	0.22	0.10	0.10
2. Frequency regular price changes	0.07	0.09	0.07
3. Frequency with which annual mode changes	0.64	0.70	0.70

Steinsson (2008) using the same algorithm I used above for the Dominick's data.[31] As in the Dominick's data, price changes (I focus on regular prices) are large (11% on average) and dispersed: 25% of price changes are less than 3% in absolute value and another 25% are greater than 13% in absolute value. The frequency of posted price changes (22% per month) is greater than that of regular price changes (7% per month) and roughly 75% of prices are at the annual mode.

I study two versions of a menu-cost economy using these data. The first has economies of scope and Poisson shocks, as above. In a second variation, I assume away the economies of scope, but introduce random menu costs, so that the menu-cost for any given good is equal to either 0, with probability p^{ϕ}, or ϕ, with probability $1 - p^{\phi}$. This is an alternative way to generate small price changes in the model. I choose parameters in each of these models so as to match the moments listed in rows 1–6 in Table IV and again target the fraction of prices at the annual mode.

Table V reports the business cycle statistics in these economies. As earlier, the model with economies of scope predicts a volatility of consumption that is 85% as large as in the Calvo model with the same frequency of price changes. The economy with random costs also predicts real effects of money that are nearly as large. Thus, the conclusions I draw are not sensitive to the exact mechanism I use to generate small price changes in the model or to the fact that I have focused on the Dominick's data.

[31]The full set of statistics they computed is available in the online supplement titled "More Facts About Prices" (Nakamura and Steinsson (2008, Tables 24–28)). I only report aggregate statistics (computed as a weighted median of ELI-level statistics) here.

TABLE V

AGGREGATE IMPLICATIONS, BLS CALIBRATION

	Benchmark			K and Inelastic M^d		Add Intermediates	
	Calvo	Economies of Scope	Random Menu Cost	Calvo	Economies of Scope	Calvo	Economies of Scope
Businss cycle statistics							
$\sigma(C)$ (%)	0.45	0.39	0.42	0.46	0.33	0.58	0.55
Serial correlation C	0.93	0.92	0.93	0.85	0.80	0.89	0.88
$\sigma(Y)$ (%)				0.75	0.55	0.75	0.70
Serial correlation Y				0.85	0.80	0.89	0.88
Law of motion P							
$P(t-1)$	0.897	0.851	0.859	0.868	0.811	0.935	0.924
$g(t)$	−0.777	−0.715	−0.775	−0.695	−0.482	−0.842	−0.792
$K(t-1)$	–	–		−0.045	−0.056	−0.020	−0.026
R^2	1	0.999	0.999	1	0.996	1	1.000

Richer Business Cycle Dynamics

I next extend the analysis to allow for capital accumulation, interest-elastic money demand, and use of intermediate inputs in production.[32] I describe next the additional assumptions I make.

Consumers: The consumer owns all the capital stock in this economy and rents it to manufacturers. The consumer chooses how much to invest in the capital stock and faces the budget constraint

$$P_t[c_t + i_t + \xi i_t^2] + \mathbf{Q}_t \cdot \mathbf{B}_{t+1} \leq W_t l_t + \Pi_t + B_t + P_t r_t^k k_t,$$
$$i_t = k_{t+1} - (1-\delta)k_t,$$

where i_t is investment, ξi_t^2 is a quadratic capital adjustment cost, and r_t^k is the (real) rental rate of capital. The resource constraint for final goods now reads

$$c_t + i_t + \xi i_t^2 + n_t = y_t = \left(\frac{1}{N}\sum_{i=1}^{N} c_t(i)^{(\theta-1)/\theta}\right)^{\theta/(\theta-1)},$$

where $c_t(i)$ is defined as earlier. Here the final good has four uses: consumption, investment, adjustment costs, and use as an intermediate input in production by manufacturing firms, n_t.[33] GDP in this economy is thus equal to $y_t - n_t$.

[32]Basu (1995) and Nakamura and Steinsson (2009).

[33]The interpretation here, as is standard in the literature, is that the Dixit–Stiglitz aggregator represents the production function for a final good produced by firms in a perfectly competitive sector that buy intermediate goods i, z from retailers.

I follow Dotsey, King, and Wolman (1999) and assume a particular form for the money demand without explicitly modeling the source of the transactions demand for money:

$$\log \frac{M_t}{P_t} = \log c_t - \eta R_t,$$

where R_t is the gross nominal risk-free rate.

Manufacturers: I modify the production function to

$$y_t(i,z) = c_y[l_t(i,z)^{\alpha} k_t(i,z)^{1-\alpha}]^{1-\nu}[n_t(i,z)]^{\nu},$$

where c_y is a constant of normalization, $k_t(i,z)$ is the amount of capital hired by manufacturers in sector (i,z), and $n_t(i,z)$ is the amount of the intermediate good they purchase. Perfect competition implies that, since manufacturers buy the final good at price P_t, the price at which manufacturers sell the good to retailers is

$$\omega_t(i,z) = [W_t^{\alpha}(P_t r_t^k)^{1-\alpha}]^{1-\nu} P_t^{\nu}.$$

Retailers: The problem of the retailer is similar to that discussed earlier. I augment the aggregate state space to include the stock of capital, k. I approximate the law of motion for capital, real wages, and rental price of capital, since these affect the retailer's profits. I solve the problem using a Krusell–Smith algorithm as above, now allowing aggregate variables to depend on the capital stock as well.[34]

Results: As in the exercise above, I assume a monthly frequency in the model, the frequency in the BLS data. I choose the size of the capital adjustment costs, ξ, to reproduce a volatility of investment that is three times greater than output, as in the data. The share of capital in value added, α, is equal to 0.33. I set the semielasticity of money demand equal to where $\eta = 1.9$, an estimate from Ireland (2009). I report results from two experiments: one without intermediate inputs and the other from an experiment in which ν is equal to 0.66, a value that implies a materials share of slightly below 50%, as in the data.

Notice in Table V that my earlier results about the relative variability of output in the menu-cost models are robust to allowing for richer business cycle dynamics. The standard deviation of consumption (GDP) is 72% (73%) as large in the menu-cost model as in the Calvo model when I assume away intermediate inputs. When I do allow for intermediate inputs, the selection effect becomes slightly weaker, so the gap between menu costs and the Calvo model is bridged even further. The volatility of consumption is 0.55% in the menu-cost model (0.58% in the Calvo model), while that of GDP is equal to 0.70% (0.75%) in the Calvo model.

[34]Once again, this approximate law of motion is very accurate. See Table V.

The economy with capital and interest-elastic money demand alone suffers from the persistence problem discussed by Chari, Kehoe, and McGrattan (2002): the serial correlation of output is 0.80 versus 0.92 in the economy without capital and interest-elastic money. This problem is shared, however, by Calvo and menu-cost models alike. The key result of my paper—that the real effects of money in a menu-cost model with heterogeneity in the size of price changes are similar to those in the Calvo model—is invariant to the persistence problem.

5. EMPIRICS

I have assumed economies of scope in price setting as well as infrequent cost shocks so as to reproduce the dispersion in the size of price changes. I next present empirical support for these assumptions using the Dominick's data.

A. *Economies of Scope*

My results are robust to alternative mechanisms of generating the distribution of price changes in the data, as shown above. Nevertheless, I provide next some empirical support for economies of scope in price setting.[35] My theory assumes economies of scope at the retail level (among different goods sold by the retailer). Hence in the data I attempt to establish evidence for economies of scope across different goods sold by Dominick's.

Synchronization

One prediction of my model is that the retailer synchronizes its price changes, both regular and sale related, since the costs of such actions, ϕ^R and κ, are shared across goods.

Table VI provides evidence that price changes within a store are synchronized. I report results from probit regressions that relate a good's price adjustment decision to several measures of synchronization. These measures are the fraction of other prices that change in any given period in (a) the same manufacturer and product category, (b) the same manufacturer but a different product category, (c) the same product category but different manufacturers, and (d) all other goods within the store. I include these different partitions since different managers may be responsible for setting the prices of different goods and it is unclear ex ante at what level within the store the economies of scope are most important.

These regressions also control for two good-specific variables: (i) the good's markup gap (the absolute log deviation of the good's markup, $\mu_{it} = \frac{p_{it-1}}{c_{it}}$, from

[35]Lach and Tsiddon (2007) and Levy et al. (1997) also provided empirical evidence of economies of scope in price adjustment.

TABLE VI

SYNCHRONIZATION IN PRICE CHANGES[a]

	All Price Changes	Regular Price Changes		Initiate Sales	
	I	I	II	I	II
Own markup gap	0.96	0.17	0.10	0.61	0.62
Temp_{t-1}	0.60	0.04	0.04	–	–
Fraction of price changes:	*Posted*	*Posted*	*Regular*	*Posted*	*Sales*
Same manufacturer and category	0.83	0.03	0.12	0.59	0.60
Same manufacturer, other category	0.05	<u>−0.01</u>	0.02	0.04	0.06
Same category, other manufacturers	<u>0.01</u>	<u>0.00</u>	0.02	<u>0.02</u>	<u>−0.01</u>
Storewide (all other goods)	0.27	<u>0.01</u>	0.09	0.13	0.18
Number of observations	450,182	450,182	450,182	364,541	364,541

[a]The marginal effect on the probability of price changes is reported. Observations are weighted (by each UPC's revenue share). The fraction of price changes is computed for all UPCs other than the UPC in question, weighting incidences of price changes using each UPC's revenue share. Observations are excluded if less than 5 products are available in a group in a particular period. All coefficients are statistically significantly different from 0 at the 1% level (unless underlined). Posted prices are the actual (transactions) prices set by the retailer.

its time series mean) and (ii) an indicator variable for whether the past price was a temporary price.

I report in Table VI marginal effects of a change in the independent variables on the price adjustment decision.[36] The markup gap is strongly correlated with a good's likelihood of a price change: a 10% increase in the markup gap raises the probability of a price change by 9.6%. This suggests that pricing is indeed state-dependent, since in the Calvo model the adjustment hazard is constant.[37]

Price changes are synchronized, especially for goods in the same manufacturer and product category. When the fraction of all other changes increases from 0 to 1, the likelihood that the price in question also changes increases by 0.83. Synchronization is also present at the store level: an increase in the storewide fraction of price changes from 0 to 1 increases the likelihood of a price change for any individual good by 0.27.

There is much less correlation between changes in the *regular* price of a good with all *posted* price changes for other goods: the marginal effects are close to 0.[38] In contrast, there is much more synchronization of a given regular price change with other regular price changes. The strongest correlation is in a given manufacturer and product category (the marginal effect is 0.125), but there is also strong storewide synchronization (0.091). Hence there is little correlation between changes in regular prices and sales (which make up most price changes), but much stronger correlation between regular price changes. This

[36]Standard errors, not reported here, are typically very small. I underline those coefficients that are not statistically significant from 0 at a 1% level.

[37]Eichenbaum, Jaimovich, and Rebelo (2011) conducted a similar test of state dependence.

[38]The markup is now computed using the old regular price.

disconnect is evidence that temporary discounts and regular price changes are performed at different levels (consistent with assumption I make in the model) and there is little interaction among the two types of price changes.

The last columns of Table VI focus on sales alone. Because these make up most of the price changes in the store, the evidence for synchronization is similar to that for all price changes.

Direct Test of Economies of Scope

Synchronization alone is necessary but not sufficient evidence of economies of scope, since it can arise for other reasons, including correlated shocks. A stronger test of my theory is one that recognizes that if economies of scope are indeed present, any individual good's adjustment decision depends not only on that good's desired price change, but also on the other goods' desired price changes. This is the mechanism that generates small price changes in the model, since a good may experience a small price change as long as other goods need larger price changes.

To test this mechanism, recall that in the presence of economies of scope, a good's decision to adjust, for example, its regular price, depends on the vector of markups of all the goods the retailer sells:

$$\text{adjust good } i \quad \text{if } V^R(\boldsymbol{\mu}^R_{-1}) > V^{T,N}(\boldsymbol{\mu}^R_{-1}).$$

A second-order approximation (around the optimal markup, so that first-order terms vanish by the envelope condition) to the above inequality yields

$$\text{adjust good} \quad i \text{ if } \alpha_0+\alpha_1(\ln\mu_{it}-\ln\mu_i^*)^2+\alpha_2\sum_{j\neq i}(\ln\mu_{jt}-\ln\mu_j^*)^2 > \phi^R, \tag{1}$$

where $\mu_i^* = \frac{\gamma}{\gamma-1}$ is the optimal markup. Because the menu-cost, ϕ^R, is shared by all goods, the adjustment decision is a function of the markup gap of all the goods, not of any individual good i. In contrast, absent economies of scope, the adjustment decision for any good i is a function of that good's idiosyncratic state and not of the state of all other goods:

$$\text{adjust good } i \quad \text{if } \alpha_0+\alpha_1(\ln\mu_{it}-\ln\mu_i^*) > \phi^R.$$

In Table VII, I present results of regressions in which I test the null of no scope economies. Under this null, the price adjustment decision is independent of the markup gap of any good other than good i, that is, $\alpha_2 = 0$. I thus estimate probit regressions of equation (1). I once again combine goods into four nonoverlapping partitions, as above, to gauge the level of aggregation at which economies of scope are more important.

These regressions, by using the markup gap directly, account for the possibility that unobserved correlated costs or preference shocks are responsible

TABLE VII

DIRECT TEST OF ECONOMIES OF SCOPE[a]

	All Price Changes	Regular Price Changes
Own markup gap	0.94	0.15
Temp_{t-1}	0.62	0.06
Average markup gap of other goods	*All*	*All*
Same manufacturer and category	0.30	0.23
Same manufacturer, other category	<u>0.06</u>	<u>0.02</u>
Same category, other manufacturers	0.66	<u>0.02</u>
Storewide (all other goods)	0.52	1.17
Number of observations	450,182	450,182

[a]The marginal effect on the probability of price changes is reported. Observations are weighted (by each UPC's revenue share). The fraction of price changes is computed for all UPCs other than the UPC in question, weighting incidences of price changes using each UPC's revenue share. Observations are excluded if less than 10 products are available in a given manufacturer/category group in a particular period. All coefficients are statistically significantly different from 0 at the 1% level (unless underlined).

for the synchronization in price changes. The markup gap for good i is indeed strongly correlated with the mean markup gap of all other goods within the store, but both are controlled for in these regressions. In other words, an economy with correlated cost shocks but no economies of scope would predict a coefficient $\alpha_2 = 0$, as evident in the decision rules above. Hence, this direct test allows me to determine whether synchronization in price changes is solely driven by correlated cost shocks or whether economies of scope play a role as well.

As shown in Table VII, the average gap of all other goods is indeed an important determinant of any given good's price adjustment decisions. This is especially true for regular price changes where the average markup gap of goods in the same manufacturer category has a marginal effect (0.23) on the price adjustment decision that is greater than that of that good's own markup gap (0.15). Even stronger is the dependence on the average markup gap of all goods within the store (the marginal effect is 1.17). This evidence thus rejects the null hypothesis of no scope economies. It also suggests that, especially for regular prices, economies of scope are more important at the level of the store, rather than at the level of individual product categories.

B. *Evidence on Costs*

A second assumption I have made is that retailers infrequently receive relatively large cost shocks so that the distribution of their markup gap (absent a price change) has much mass near zero, while some markups are very far away from their optimal value. In other words, the distribution of markups in the model exhibits excess kurtosis.

TABLE VIII

DISTRIBUTION OF THE MARKUP GAP AND CHANGES IN COSTS[a]

	Markup Gap		
	Posted Prices	Regular Prices	Change in Costs
5%	−0.31	−0.21	−0.05
10%	−0.19	−0.13	−0.02
25%	−0.07	−0.06	−0.003
Median	0.00	−0.01	0.000
75%	0.06	0.05	0.004
90%	0.15	0.12	0.02
95%	0.21	0.18	0.05
Kurtosis	7.66	9.30	29.31
Number of observations	1,133,947	1,133,947	1,150,922

[a]Statistics weigh observations by revenue share of each product. I exclude observations with |gap| and |Delta cost| > 100%.

Such kurtosis is also present in the grocery store data I look at. I report in Table VIII a number of moments of the distribution of the markup gap for both posted and regular prices. I discuss the evidence for regular prices, although that for posted prices is very similar. Notice that the 25th percentile of the distribution of the markup gap is equal to −0.06, while the 75th percentile is equal to 0.05. Hence, a large number of desired price changes are fairly small. In contrast, the 5th percentile of this distribution is equal to −0.21, while the 95th percentile is equal to 0.18: hence a small number of goods experience very large markup gaps. As a result, the distribution of markup gaps shows large kurtosis: 7.66 for posted prices and 9.30 for regular prices.

Table VIII also reports moments of the distribution of changes in the wholesale costs at which Dominick's purchases its goods. The distribution of changes in costs has much mass near 0 (the 25th and 75th percentiles are −0.003 and 0.004, respectively) and since a small fraction of changes in costs are very large (the 1st and 99th percentiles are equal to ±0.12), the kurtosis of the distribution of changes in costs is equal to 29.[39]

C. *Relationship to Other Work*

The key result in this paper is driven by the fact that the distribution of the size of price changes is highly dispersed. This is neither a new fact, nor specific to data from grocery stores. My contribution in this paper is to study the aggregate consequences of this dispersion, not document the facts per se.

[39]See also the evidence in Eichenbaum, Jaimovich, and Rebelo (2011) using admittedly better cost data. They also found that costs (both high- and low-frequency measures) change infrequently.

Klenow and Kryvtsov (2008) reported that 40% of price changes are less than 5% in absolute value in the BLS price data, a data set in which prices change by 9.5% on average. They also showed that heterogeneity in the size of price changes across sectors is, alone, insufficient to account for this large number of small price changes. Kashyap (1995) studied prices for products sold using retail catalogues and also documented that many price changes are small: 44% of price changes in his data set are less than 5% in absolute value. The kurtosis of price changes is 15.7 in his data. Kackmeister (2005) reported that 33% of price changes are less than 10% in absolute value in an environment where the average magnitude of price changes is 20% in a study of prices in retail stores.

D. *Limitations of My Analysis and Extensions*

My model misses several features of the data that are worth exploring in future work. First, I do not allow consumers to store the goods. To the extent to which they can do so, their ability to take advantage of sales would be greatly increased. Recall, however, that the model I study (through an appropriate choice of a demand elasticity) accounts for the fraction of goods that are sold during sales. Hence, although an extension of the model that allows for storability of goods would surely impart richer microlevel dynamics, my conjecture is that such an extension (appropriately parameterized to match the price and quantity facts) would not overturn my main results.

Second, I assume that goods sold by different retailers are imperfectly substitutable, whereas in the data different retailers sell identical goods. An alternative would be a Hotelling-type setting in which consumers face fixed costs of reaching stores that sell identical goods. In such an environment, temporary shocks to demand would generate a motive for temporary price discounts.[40] Combined with a rigidity in the regular price, as in Kehoe and Midrigan (2008), temporary changes would also revert to the pre-sale price, as in the data. An extension along these lines is also an exciting avenue for future research.

6. CONCLUSION

Simple menu-cost models fail to account for two features of the microeconomic price data: the dispersion in the size of price changes and the fact that many price changes are temporary price discounts. I study an economy with economies of scope in price adjustment—a more flexible specification of the distribution of firm-level uncertainty—and a two-tier price setting technology (for regular and posted prices) that is capable of replicating these microlevel facts.

[40]Warner and Barsky (1995).

I find that in this economy the real effects of money are much greater than those in a simple menu-cost economy that fails to account for the microlevel facts. This result is primarily accounted for by the heterogeneity in the size of price changes in my model. When price changes are very dispersed in absolute value, the measure of marginal firms whose adjustment decisions are sensitive to aggregate shocks is small. Hence the selection effect is much weaker and the model produces real effects of a similar magnitude to those in Calvo-type models.

REFERENCES

AGUIRREGABIRIA, V. (1999): "The Dynamics of Markups and Inventories in Retailing Firms," *Review of Economic Studies*, 66, 275–308. [1156]

BARSKY, R., M. BERGEN, S. DUTTA, AND D. LEVY (2000): "What Can the Price Gap Between Branded and Private Label Products Tell Us About Markups?" Mimeo, University of Michigan. [1156]

BASU, S. (1995): "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *American Economic Review*, 85, 512–531. [1170]

BILS, M., AND P. KLENOW (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947–985. [1139]

BURDETT, K., AND K. JUDD (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969. [1140]

BURSTEIN, A. (2006): "Inflation and Output Dynamics With State-Dependent Pricing Decisions," *Journal of Monetary Economics*, 53, 1235–1257. [1141]

BURSTEIN, A., AND C. HELLWIG (2008): "Prices and Market Shares in a Menu Cost Model," Mimeo, UCLA. [1141,1167]

BUTTERS, G. (1977): "Equilibrium Distribution of Sales and Advertising Prices," *Review of Economic Studies*, 44, 465–491. [1140]

CABALLERO, R., AND E. ENGEL (1993): "Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy," *Review of Economic Studies*, 60, 95–119. [1141]

——— (2007): "Price Stickiness in *Ss* models: New Interpretations of Old Results," *Journal of Monetary Economics*, 54, S100–S121. [1141,1165-1167]

CAPLIN, A., AND J. LEAHY (1991): "State-Dependent Pricing and the Dynamics of Money and Output," *Quarterly Journal of Economics*, 106, 683–708. [1141]

CAPLIN, A., AND D. SPULBER (1987): "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, 102, 703–725. [1141]

CHARI, V. V., P. KEHOE, AND E. MCGRATTAN (2002): "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?" *Review of Economic Studies*, 69, 533–563. [1172]

CHEVALIER, J., A. KASHYAP, AND P. ROSSI (2003): "Why Don't Prices Rise During Periods of Peak Demand? Evidence From Scanner Data," *American Economic Review*, 93, 15–37. [1140, 1156]

CHRISTIANO, L., M. EICHENBAUM, AND C. EVANS (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45. [1156]

DANZIGER, L. (1999): "A Dynamic Economy With Costly Price Adjustments," *American Economic Review*, 89, 878–901. [1141]

DOTSEY, M., R. KING, AND A. WOLMAN (1999): "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics*, 114, 655–690. [1141,1142,1171]

DUTTA, S., M. BERGEN, AND D. LEVY (2002): "Price Flexibility in Channels of Distribution: Evidence From Scanner Data," *Journal of Economic Dynamics and Control*, 26, 1845–1900. [1143]

EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO (2011): "Reference Prices and Nominal Rigidities," *American Economic Review* (forthcoming). [1141,1159,1173,1176]
GERTLER, M., AND J. LEAHY (2008): "A Phillips Curve With an *Ss* Foundation," *Journal of Political Economy*, 116, 533–572. [1143,1147,1149,1156,1157]
GOLOSOV, M., AND R. E. LUCAS, JR. (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–199. [1139,1143,1149,1153,1157,1165]
HANSEN, G. (1985): "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 16, 309–327. [1153]
HOCH, S., X. DREZE, AND M. E. PURK (1994): "EDLP, Hi-Lo, and Margin Arithmetic," *Journal of Marketing*, 58, 16–27. [1143]
HOSKEN, D., AND D. REIFFEN (2004): "Patterns of Retail Price Variation," *RAND Journal of Economics*, 35, 128–146. [1159]
IRELAND, P. (2009): "On the Welfare Cost of Inflation and the Recent Behavior of Money Demand," *American Economic Review*, 99, 1040–1052. [1171]
KACKMEISTER, A. (2005): "Yesterday's Bad Times Are Today's Good Old Times: Retail Price Changes in the 1890s Were Smaller, Less Frequent, and More Permanent," Finance and Economics Discussion Series, Federal Reserve Board. [1177]
KASHYAP, A. (1995): "Sticky Prices: New Evidence From Retail Catalogues," *Quarterly Journal of Economics*, 110, 245–274. [1177]
KEHOE, P., AND V. MIDRIGAN (2008): "Temporary Price Changes and the Real Effects of Monetary Policy," Working Paper 14392, NBER. [1140,1141,1144,1151,1159,1160,1165,1168,1177]
KHAN, A., AND J. THOMAS (2007): "Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies," *American Economic Review*, 97, 1165–1188. [1155]
KLENOW, P., AND O. KRYVTSOV (2008): "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent US Inflation?" *Quarterly Journal of Economics*, 123, 863–904. [1139,1143, 1159,1164,1176,1177]
KLENOW, P., AND J. WILLIS (2006): "Real Rigidities and Nominal Price Changes," Working Paper 06-03, FRB Kansas City. [1155]
KRUSELL, P., AND A. SMITH (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896. [1155]
LACH, S., AND D. TSIDDON (2007): "Small Price Changes and Menu Costs," *Managerial and Decision Economics*, 28, 649–656. [1140,1142,1172]
LAZEAR, E. (1986): "Retail Pricing and Clearance Sales," *American Economic Review*, 76, 14–32. [1140]
LEVY, D., M. BERGEN, S. DUTTA, AND R. VENABLE (1997): "The Magnitude of Menu Costs: Direct Evidence From Large US Supermarket Chains," *Quarterly Journal of Economics*, 112, 781–826. [1142,1161,1172]
MIDRIGAN, V. (2011): "Supplement to 'Menu Costs, Multiproduct Firms, and Aggregate Fluctuations'," *Econometrica Supplemental Material*, 79, http://www.econometricsociety.org/ecta/Supmat/6735_data and programs.zip; http://www.econometricsociety.org/ecta/Supmat/6735_miscellaneous.pdf. [1144]
MIRANDA, M., AND P. FACKLER (2002): *Applied Computational Economics and Finance*. Cambridge, MA: MIT Press. [1155]
NAKAMURA, E., AND J. STEINSSON (2008): "Five Facts About Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123, 1415–1464. [1142,1143,1169]
——— (2009): "Price Setting in Forward-Looking Customer Markets," Mimeo, Columbia University. [1154,1170]
NEVO, A. (1997): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69, 265–306. [1156]
PELTZMAN, S. (2000): "Prices Rise Faster Than They Fall," *Journal of Political Economy*, 108, 466–502. [1144]
ROGERSON, R. (1988): "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, 21, 3–16. [1153]

ROMER, C., AND D. ROMER (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055–1084. [1156]
SCARF, H. (1959): "The Optimality of (S, s) Policies in the Dynamic Inventory Problem," in *Mathematical Methods for the Social Sciences*, ed. by K. Arrow, S. Karlin, and P. Suppes. Cambridge, MA: MIT Press, 196–202. [1156]
SHESHINSKI, E., AND Y. WEISS (1977): "Inflation and Costs of Price Adjustment," *Review of Economic Studies*, 44, 287–303. [1141]
——— (1992): "Staggered and Synchronized Price Policies Under Inflation," *Review of Economic Studies*, 59, 331–359. [1156]
SMETS, F., AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles. A Bayesian DSGE Approach," Working Paper 722, ECB. [1156]
SOBEL, J. (1984): "The Timing of Sales," *Review of Economic Studies*, 51, 353–368. [1140]
STOKEY, N., AND R. E. LUCAS, JR. (1989): *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press. [1156]
VARIAN, H. (1980): "A Model of Sales," *American Economic Review*, 70, 651–659. [1140]
WARNER, E., AND R. BARSKY (1995): "The Timing and Magnitude of Retail Store Markdowns: Evidence From Weekends and Holidays," *Quarterly Journal of Economics*, 110, 321–352. [1140, 1177]
ZBARACKI, M., M. RITSON, D. LEVY, S. DUTTA, AND M. BERGEN (2004): "Managerial and Customer Costs of Price Adjustment: Direct Evidence From Industrial Markets," *Review of Economics and Statistics*, 86, 514–533. [1141,1147,1148,1161]
ZBARACKI, M., D. LEVY, AND M. BERGEN (2007): "The Anatomy of a Price Cut: Discovering Organizational Sources of the Costs of Price Adjustment," Mimeo, University of Western Ontario. [1141,1147]

Dept. of Economics, New York University, 19 West 4th Street, 6th Floor, New York, NY 10012, U.S.A.; virgiliu.midrigan@nyu.edu.

Manuscript received September, 2006; final revision received February, 2010.