Menu Costs, Multi-Product Firms, and Aggregate Fluctuations[†]

Virgiliu Midrigan[*]

Ohio State University

January 2006

Abstract

I employ a large set of scanner price data collected in retail stores to document that (i) although the average magnitude of price changes is large, a substantial number of price changes are small in absolute value; (ii) the distribution of non-zero price changes has fat tails; and (iii) stores tend to adjust prices of goods in narrow product categories simultaneously. I extend the standard menu costs model to a multi-product setting in which firms face economies of scale in the technology of adjusting prices. The model, because of its ability to replicate this additional set of micro-economic facts, can generate aggregate fluctuations much larger than those in standard menu costs economies.

JEL classifications: E31, E32.
Keywords: state-dependent pricing, multi-product firms.

[†]I am indebted to George Alessandria, Bill Dupor, Paul Evans, Patrick Kehoe and Mario Miranda for valuable advice and support, as well as Michael Dotsey, Narayana Kocherlakota, Ellen McGrattan, Martin Schneider and seminar participants at the Ohio State University, Minneapolis Fed and Philadelphia Fed for comments and suggestions. Any errors are my own.

[*]Department of Economics, The Ohio State University, 1945 North High Street, Columbus, OH 43210. Phone: (614) 390-5809. Email: midrigan.1@osu.edu

1. Introduction

It has been well documented, using both survey evidence, but also direct observation, that individual goods prices are sticky. The latest major piece of evidence supporting the notion that prices adjust sluggishly is a study by Bils and Klenow (2004) who find, based on a dataset of prices collected by the BLS, that half of the consumer goods prices in the US economy adjust less frequently than every 4.3 months. Whether these firm-level rigidities have important macroeconomic implications is however still an open question.

Most recent work analyzing the consequences of nominal rigidities assumes that firms employ ad-hoc policy rules and does not explicitly model the source of price stickiness. These, time-dependent models postulate that the timing of price changes is exogenous and unresponsive to the state of the world. Information-gathering costs or institutional restrictions are presumed to give rise to this behavior, but these frictions are, with a few exceptions[1], rarely modeled. The fact that these models lack micro-foundations makes them inappropriate for the study of many interesting policy questions, but also reduces the number of dimensions along which the theory can be tested.

Recently the profession has witnessed a growing interest in an alternative class of models in which agents solve fully-specified problems and nominal rigidities arise endogenously, due to fixed physical (menu) costs of changing prices. These, state-dependent or (S, s) pricing models can be traced back to the work of Barro (1972) and Sheshinski and Weiss (1977, 1983), but only recent advances in numerical solution techniques have enabled researchers to study dynamic, general equilibrium versions of these economies. Their aggregate implications are, nevertheless, not well understood. In particular, the ability of firm-level nominal rigidities to generate business cycle fluctuations from nominal shocks crucially depends in these models on the distributional assumptions made to aggregate the economy.

[1] See Bonomo and Carvalho (2005) and the references therein.

The predictions of menu-cost models range from stark neutrality[2] to cases in which the economy is virtually indistinguishable from time-dependent setups[3]. Golosov and Lucas (2004) study the properties of a model with firm-level disturbances capable to match the fact that the magnitude of price changes is large in the US economy: 10% on average, much larger than what can be explained by aggregate shocks alone. They find that the model produces very little output volatility from monetary shocks, a result similar in spirit to that of Caplin and Spulber (1987). Klenow and Kryvtsov (2004) reach an opposite conclusion. They document that there is little evidence of across-firm synchronization in the US price data, contrary to what standard menu-cost models predict. They find that a model with time-varying costs of price adjustment that can replicate this feature of the data behaves identically to a time-dependent sticky price model and produces large output variability from monetary shocks.

This paper revisits the question of whether menu costs of price adjustment can, in fact, generate a monetary transmission mechanism. I argue that a truly micro-founded model must be rendered consistent with the price adjustment practices observed at the firm level before one can proceed to study its aggregate predictions. I start by documenting several salient micro-economic features that characterize firm pricing behavior. To this end, I employ a large set of scanner price data collected in a number of grocery stores over a twelve-year period. In addition to the large frequency and magnitude of price changes, documented by Klenow and Kryvtsov (2004), I document two additional features of the data. First, a large number of non-zero price changes are small in absolute value. Second, the distribution of price changes, conditional on adjustment, exhibits excess kurtosis.

These two facts seem, at a first glance, inconsistent with menu-cost models. Firms that face

[2]Caplin and Spulber (1987), Golosov and Lucas (2004), a version of the model of Gertler and Leahy (2005).

[3]Klenow and Kryvtsov (2004), another calibration of the model in Gertler and Leahy (2005). See also Burstein (2003), Dotsey, King and Wolman (1999), Danziger (1999), Caplin and Leahy (1991) for studies that explore the consequences of fixed costs of resetting prices.

fixed costs of adjustment only reprice when the losses from not doing so are large, and thus tend to adjust prices by large amounts. As Lach and Tsiddon (2005) argue, however, extensions of the model to a multi-product setting in which firms face interactions in the costs of price adjustment of various goods[4] can explain the large number of small price changes. Consider for example, the problem of a restaurant whose prices are quoted on a single menu. If a single item on the menu is subject to a large idiosyncratic shock and needs repricing, the restaurant might find it optimal to pay the fixed cost and reprint the menu. Conditional on having payed this fixed cost, changing any other price on the menu is costless: the restaurant will the reprice all its other items, even if some need only small price changes.

If economies of scale in the technology of adjusting prices of multi-product firms are indeed the source of the large number of small price changes observed in the data, one should observe that prices of goods within a store adjust in tandem. I extend the evidence of Lach and Tsiddon (1996) for data collected in Israel and show, using a discrete choice model of a store's price adjustment practices, that within-store synchronization is indeed a pervasive feature of the micro-price data. Prices of goods in narrow product categories, as well as those produced by a given manufacturer, tend to adjust simultaneously, even after controlling for the effect of wholesale price changes and aggregate marginal cost disturbances that might trigger concomitant adjustment.

I next formulate, calibrate and quantitatively study the properties of a model in which a two-product firm faces a fixed cost of changing its entire menu of prices, but, conditional on paying this cost, zero additional costs of resetting any given price on the menu. I calibrate the distribution of idiosyncratic technological disturbances, the size of the fixed costs of price adjustment, as well as the persistence of the technology processes, by requiring the model to accord with the features of the data enumerated above. I find that the model, because of its ability to replicate this additional

[4]See also Sheshinski and Weiss (1992).

set of micro-economic facts, can generate aggregate fluctuations of the same magnitude as in time-dependent economies.

To understand the intuition behind this result, recall that the main reason standard state-dependent models with idiosyncratic marginal cost shocks[5] under-perform their time-dependent counterparts is the fact that the identity of adjusters in models with menu costs varies endogenously in response to aggregate disturbances[6]. Most firms that adjust in times of, say, a monetary expansion, are firms whose incentive to increase prices arising from the aggregate disturbance is reinforced by an idiosyncratic cost shock that triggers a desired price change in the same direction. This 'selection' effect ensures that the aggregate price level is much more responsive to nominal shocks than in time-dependent models in which the timing of price changes is exogenous. I find, however, that matching the excess kurtosis of price changes and the large number of small price changes observed in the data requires that technology cost shocks be highly leptokurtic. This feature of the calibration, as well as the fact that some price changes arise independently of the state of the world, because of the multi-product aspects of the model, reduces the role of self-selection and therefore the responsiveness of the aggregate price level to monetary shocks.

A second reason the state-dependent model produces business cycle fluctuations of similar magnitude as those in time-dependent models is a property of menu-cost economies that has received little attention in earlier work. One feature of all sticky price models is front-loading of firm prices in response to future expected increases in marginal costs. If monetary disturbances are persistent, a monetary expansion today is expected to be followed by additional increases of the money stock in future periods. Forward-looking firms take these forecasts into account and respond stronger to

[5] Golosov and Lucas (2004), Danziger (1999).

[6] A second feature characteristic to menu cost models are endogenous fluctuations in the fraction of adjusters in response to aggregate disturbances, a feature that, as Klenow and Kryvtsov (2004) observe, is absent in the data. Golosov and Lucas (2004) point out, however, that models with marginal cost shocks volatile enough to account for the large magnitude of price changes observed in the data generate little across-firm synchronization.

the nominal shock than they would otherwise in a static world. I show in this paper that front-loading in response to future disturbances is mainly a feature of time-dependent models: firms in a state-dependent world are less willing to forego current profits in order to ensure that expected future changes in marginal costs are reflected in the current price of the firm. Intuitively, when a state-dependent firm finds itself with a suboptimal price in a future period, it has the option to pay the menu cost and reprice: it's losses are bounded therefore by the size of the menu cost. In contrast, a time-dependent firm pays dearly every time its price is suboptimal: because it has to wait for an exogenously fixed number of periods before it gets to reset its price, it will incur much larger losses from a suboptimal price relative to what a state-dependent firm would. A time-dependent firm's incentive to offset future deviations of its price from the optimum is therefore larger than that of a state-dependent firm and it adjusts more aggressively in response to persistent shocks to the growth rate of the money supply.

This paper proceeds as follows. Section 2 discusses the data used in the empirical work, and documents its salient features. Section 3 presents results of ordered discrete choice models in which I provide evidence of within-store synchronization of price changes. Section 4 discusses the model economy. Section 5 quantitatively evaluates its performance. Section 6 concludes. Appendices discuss the non-linear solution techniques used to solve the functional equations that characterize the equilibrium of the model economy and several aspects of the data.

2. Data

I conduct inference using two sources of publicly available sets of scanner price data, maintained by the Kilts Center for Marketing at the University of Chicago Graduate School of Business[7]. The first dataset was assembled by AC Nielsen and consists of daily observations on the purchasing practices of a panel of households in Sioux Falls (South Dakota) and Springfield (Missouri). I

[7]The data is available online at http://gsbwww.uchicago.edu/kilts/research/index.shtml

use this household level data to construct a panel of weekly price series spanning more than two years (January 1985 to March 1987), 31 stores and 115 products in six different product categories (ketchup, tuna, margarine, peanut butter, sugar and toilet tissue)[8].

The second source of data is a by-product of a randomized pricing experiment conducted by the Dominick's Finer Foods retail chain in cooperation with the Chicago GSB. Nine years (1989 to 1997) of weekly store level data on the prices of more than 4500 products for 86 stores in the Chicago area are available. The products available in this database range from non-perishable foodstuffs (frozen and canned food, cookies, crackers, juices, sodas, beer), to various household supplies (detergents, softeners, bathroom tissue), as well as pharmaceutical and hygienic products.

I discuss, in the data appendix, several aspects regarding the construction of price series. In particular, I time-aggregate weekly data into monthly observations in order to calculate statistics that can be used to evaluate the performance of a model economy in which the length of the period is a month. For Dominick's data, which sets prices on a chain-wide basis, I construct a chain-wide price using the price of the store that has the least number of missing observations for a particular good. Following Golosov and Lucas (2004), I filter out temporary price cuts (sales) that last less than four weeks. I could alternatively incorporate into the model some of the frictions that have been proposed to explain this pattern of retail price variation[9], but this would increase the model's complexity considerably, without producing additional insights. In particular, none of the empirical facts I am about to document are an artifact of my decision to purge the data of sales and time-aggregate the data.

[8]The actual number of observations is larger in the original dataset, but I discard stores/goods with a large number of missing observations. The criteria for inclusion in the sample are discussed in the appendix.

[9]Informational frictions on consumer's side of the market (Varian 1980), demand uncertainty (Lazear 1986), or thick-market explanations (Warren and Barsky, 1995), to name a few.

A. The Size and Frequency of Price Changes

Figure 1 presents histograms of the distribution of price changes, $\log\left(\frac{p_t}{p_{t-1}}\right)$, conditional on adjustment, for the two sets of data, pooled across all goods/stores/months in each sample. I truncate these distributions, by eliminating the top and bottom 1% of observations, in order to ensure that results are not driven by outliers. Superimposed on each histogram is the density of a normal distribution with the same mean and variance as that of the distribution of price changes[10]. Table 1 reports moments of these distributions, again computed using the truncated sample of observations. Several facts emerge in the data.

Fact 1: A large number of price changes are small in absolute value.

Consistent with the evidence presented by Klenow and Kryvtsov (2004), the average size of price changes is large[11]: stores in the AC Nielsen data adjust prices by 10.4% on average, while those in Dominick's sample do so by 7.7%. Notice however, in Figure 1, that a large number of price changes are close to zero. I define, in the data and in the model of the next section, a "small" price change as any price change whose magnitude is less than one-half of the mean of the absolute value of price changes in the data. Roughly 30% of price changes in both datasets are below this cutoff (5.2% and 3.8%, respectively) in both datasets of prices I work with.

Fact 2: The distribution of price changes exhibits excess kurtosis.

Notice, in Figure 1, that the number of price changes in the vicinity of zero is greater than that predicted by a normal distribution, while the tails are fatter than those of a normal. As Table 1 indicates, the kurtosis of price changes is 3.5 and 5.4, respectively, larger than that of a Gaussian distribution[12].

[10]The histogram is scaled so that its cumulative density is also equal to 1.

[11]The excessive volatility of individual goods' prices has also been documented for countries other than the US. See Dhyne et. al (2005) for a survey of findings from studies of European micro-price data.

[12]Kurtosis is defined as the ratio of the fourth central moment to the square of the variance. The kurtosis of the normal according to the convention I employ is then equal to 3.

Fact 3: Prices in grocery stores change frequently.

Table 2 presents an additional set of facts that I will use below in order to calibrate the model. It has been widely documented[13] that prices in retail stores adjust frequently. The two datasets I employ here are no exceptions. Despite the fact that I overestimate the duration of price spells by aggregating weekly data to monthly and eliminating a large number of temporary price cuts, the average price spell lasts 4 months in the AC Nielsen data and 5.2 months in the case of Dominick's prices.

Fact 4: Price changes are transitory.

Let $\hat{p}_t$ be the price (in logs) of a good in period t, expressed in deviations from a time trend. If marginal cost shocks are transitory, one would expect two observations of the firm's price, $\hat{p}_t$, sufficiently distant in time, to lie close to each other. In contrast, if shocks are highly persistent, the firm's price wonders away from the mean, and differs considerably from the price the firm has set in the past. Let

$$D_k = \frac{\frac{1}{\mathcal{N}(G_k)} \sum_{t \in G_k}^{T} |\hat{p}_t - \hat{p}_{t-k}|}{\frac{1}{\mathcal{N}(G_a)} \sum_{t \in Ga}^{T} |\hat{p}_t - \hat{p}_{t-1}|}$$

be the mean absolute difference in a good's (detrended) price in periods that are k-months apart, relative to the average absolute value of non-zero price changes, where G_k is the set of time-periods for which prices were recorded in t and $t-k$, G_a the set of periods in which the product has experienced a non-zero price change, and $\mathcal{N}$ the number of elements of a given set. Note the similarity of these statistics, which I call deviance ratios, to the variance ratios popularized by Cochrane (1988) in non-parametric tests of non-stationarity. These deviance ratios are larger, the more persistently $\hat{p}_t$ moves in a given direction, and although they have no structural interpretation,

[13] Kackmeister (2005), Dutta, Bergen and Levy (2002).

they can be used, in conjuction with the model to be presented below, to infer the persistence of marginal cost shocks. The last rows of Table 2 present the average values of this statistic in the data, at 12- and 24-month horizons. These ratios are close to 1, suggesting that shocks are not too persistent[14].

An alternative measure of persistence is the probability that the next price change will have the same sign as the current one. As Table 2 reports, these probabilities[15] are low (32% and 41%, respectively), suggesting that shocks that trigger price changes are transitory and price changes tend to be reversed. Although useful in providing information about the persistence of idiosyncratic shocks, this alternative statistic will only be used as an "over-identifying" check on the model, as it is sensitive to the definition of sales employed to purge data of temporary price cuts.

B. Ex-ante heterogeneity?

One might argue that the distributional features documented above (and in particular, the large number of small price changes) merely reflect ex-ante heterogeneity (in menu costs, volatility of shocks) across goods or stores in the sample. I test whether this is indeed the case using variance decompositions in which I gauge the importance of month, product[16], and store-specific effects in explaining the variability of the magnitude and frequency of price changes reported above. Specifically, I estimate

$$y_{it}^s = c + d_i + d_s + d_t + e_{it}^s,$$

[14]Given the short span of AC Nielsen's price series, deviance ratios are only reported for 12-month horizons in this data.

[15]Calculated as an equally weighted average of the probabilities of two consecutive positive/negative price changes. Equal weights (as opposed to weights based on the long-run probability of a price increase/decrease) are used in order to account for the upward trend in prices in Dominick's data. Use of long-run weights results in a probability of two consecutive price changes in the same direction equal to 46% in Dominick's data, but this number overstates the persistence of the price series as 65% of price changes in Dominick's data are positive.

[16]The number of non-zero price changes for a given product is small in Dominick's data in which I have collapsed the prices of the different stores into a single, chain-wide price. I therefore estimate product-category×manufacturer, as opposed to individual good effects for this dataset.

where $d_i, d_s,\ d_t$ are good, store, and month-specific effects and y_{it}^s is the size of price changes, $|\Delta \log (p_{ist}) |$, or the duration of price spells that end in a given period. As Table 3 indicates, month or store-specific heterogeneity accounts for less than 10% of the variation of the frequency and size of price changes in the data. Good-specific effects are somewhat more volatile, but nevertheless responsible for less than 16% of the variation in the sample.

C. Relationship with other evidence

I have documented two salient features of the distribution of price changes in grocery stores that will prove important in the calibration of the model economy of the next section: (i) a large number of price changes are small, and (ii) the distribution of price changes is leptokurtic. A potential objection to these findings is that they are specific to the datasets in question. I argue below that this is not the case.

Klenow and Kryvtsov (2004) report that 40% of price changes are less than 5% in absolute value in their dataset of BLS-collected price data covering all goods and services used in the construction of the CPI, a dataset in which prices change by 9.5% on average. Kashyap (1995) uses a dataset of prices for products sold in retail catalogues and also documents that many price changes are small: 44% of price changes in his dataset are less than 5% in absolute value. The kurtosis of price changes, conditional on adjustment, is 15.7 in the data and falls to 6.2 if one excludes the top and bottom 1% of observations[17]. Kackmeister (2005) presents a histogram of the distribution of price changes in a dataset of prices in retail stores: one-third of price changes are less than 10% in absolute value in an environment where the average magnitude of price changes is 20%. It is clear then that the two features of the data I document above are pervasive in different environments, and this underscores the necessity of rendering sticky price models consistent with this set of facts.

[17]These numbers are based on my own calculations using the data published in Kashyap (1995).

3. Evidence of synchronization

The large number of small price changes observed in the micro-price data might lead one to conclude against state-dependent pricing models. Several extensions to the standard state-dependent model have been proposed in earlier work in order to render it consistent with the data. One might assume time-varying adjustment costs, as in Caballero and Engel (1999) or Dotsey, King and Wolman (1999). Alternatively, as Kashyap (1995) has suggested, one might allow fluctuations in the degree of market power possessed by firms, arising from variation in consumer search costs over time[18]. In this paper I explore an alternative route. As Lach and Tsiddon (2005) have argued, an extension of the state-dependent model to a multi-product environment in which firms have large average costs of adjusting an entire menu of prices, but small marginal costs of adjusting any given price, can also generate a large number of small price changes. If interactions in the cost of price changes across different goods are indeed at play, one should observe that prices within stores adjust in tandem, and idiosyncratic disturbances explain little of a given product's adjustment decisions. I next ask whether within-store synchronization is indeed a feature of the data.

Lach and Tsiddon (1996) provide evidence that stores synchronize price adjustments of various products using a dataset of prices collected in Israel[19]. They find that the variability (across stores) of the fraction of products whose prices change in a given period is larger than what would be expected if price adjustment decisions were independent across goods. I look for evidence of within-store synchronization using an alternative approach, one that allows me to control the effect of marginal cost disturbances that might be correlated across goods sold in a particular store, and allows me to quantify the importance of within-store synchronization in determining a given good's price adjustment decision. To this end, I estimate a discrete choice model in which I relate the probability that a store adjusts its price to marginal cost shocks affecting the particular product, as

[18] See also Benabou (1992).
[19] See also Fisher and Konieczny (2000).

well as the fraction of other prices that are changed within the store. Because the two sets of data differ in nature, I discuss the empirical tests of product synchronization separately.

Dominick's

Assume that a good's frictionless optimal price is, (in logs), $p^*_{it} = \gamma_1 c_{it} + \gamma_2 w_t + u_{it}$, where c_{it} is the wholesale price of the good, w_t a chain-wide component to the good's marginal costs (wages, energy costs, etc.), and u_{it} an unobserved disturbance. With menu costs of price adjustment, charging p^*_{it} in each period is suboptimal. I assume instead that Dominick's follows an (S, s) price adjustment policy. Let p_{it} be the chain's actual price and assume that the optimal price adjustment decision x_{it} is

$$x_{it} = \begin{cases} 1, & \text{if } p^*_{it} - p_{i\tau} > S_{it} \\ 0, & \text{if } s_{it} \leqslant p^*_{it} - p_{i\tau} \leqslant S_{it} \\ -1, & \text{if } p^*_{it} - p_{i\tau} < s_{it} \end{cases}$$

where 1 denotes a price increase, -1 a price decrease, and 0 lack of adjustment; τ is the date of the previous price change, s and S are the adjustment thresholds, allowed to vary across products and time. Using the optimal price function assumed above, and assuming that the firm sets a price equal to its frictionless optimum every time it adjusts: $p_{i\tau} = \gamma_1 c_{i\tau} + \gamma_2 w_{i\tau} + u_{i\tau}$, this price adjustment rule can be rewritten as

$$x_t = \begin{cases} 1, & \text{if } \gamma_1 \Delta c_{it} + \gamma_2 \Delta w_{it} + \Delta u_{it} > S_{it} \\ 0, & \text{if } s_{it} \leqslant \gamma_1 \Delta c_{it} + \gamma_2 \Delta w_{it} + \Delta u_{it} \leqslant S_{it} \\ -1, & \text{if } \gamma_1 \Delta c_{it} + \gamma_2 \Delta w_{it} + \Delta u_{it} < s_{it} \end{cases}$$

where, say, $\Delta c_{it} = c_{it} - c_{i\tau}$ is the growth rate of the product's wholesale price since the previous

price adjustment[20]. I assume that $\Delta u_{it} \sim N(0,1)$ as the model's scale and location are unidentified.

To test the hypothesis of synchronization, I parameterize the upper and lower thresholds as linear functions of three measures of within-store synchronization: (i) the fraction of all remaining goods whose prices change in a given month; (ii) the proportion of price changes in the respective good-category (recall that data on 29 product categories, ranging from analgesics to toothpastes is available); as well as (iii) the proportion of prices of goods produced by the manufacturer of the product in question that changes in a given month[21]. All these measures of synchronization are computed based on the adjustment decision of all goods other than i in a given group (otherwise a simultaneity bias would affect our estimates), and I exclude those observations for which any of these statistics are calculated based on fewer than five observations in a given period.

These measures of within-store synchronization make intuitive sense. Levy, Dutta, Bergen and Venable (1998) use a unique store-level dataset for five supermarket chains and report, in great detail, the steps undertaken during a price change process. They report that the bulk (60%, or 79 labor-hours per week/store) of the labor effort used to adjust prices is spent on price tag changes and verification, of which most time (50-60%) goes into finding specific items on shelves. One would thus expect that economies of scale in changing prices are larger for products located in adjacent shelves/aisles. In fact, as Levy et. al. report, "The tag changes are done by aisle, where they are sorted by commodity." An ideal measure of relevant within-store synchronization would then be the fraction of prices adjusted in a given aisle, or in the vicinity of a given product. In the absence of such data, I use the fraction of price changes within a category group, or produced by a given manufacturer, as these items are usually placed in adjacent locations within the store.

In addition, I allow both the scale and location of the adjustment thresholds to differ across months and product-categories, in order to control for heterogeneity across goods/time-periods. I

[20] Cecchetti (1986) employs a similar (one-sided) model in order to study the price adjustment of magazine prices.
[21] Identified based on the first 5 digits of the upc code.

do so by allowing fixed product-category and month effects in the two threshold equations[22]. The specification of the threshold equations is therefore:

$$S_{it} = F_{it}\beta^1 + D^1,$$

$$s_{it} = F_{it}\beta^2 + D^2,$$

where D^1 and D^2 are sets of 11 month and 28 product-group dummy variables, and F_{it} is a vector of the measures of synchronization discussed above. The multi-product extension of the menu-cost model predicts that all elements of β^1 are negative and those of β^2 are positive, i.e., the (S, s) region is expected to narrow down when the proportion of the remaining prices changed within a store increases, thereby increasing the probability of a price change.

Turning to the equation for the latent variable, $\gamma_1 \Delta c_{it} + \gamma_2 \Delta w_{it} + \Delta u_{it}$, my measure of the idiosyncratic cost component c_{it} is the wholesale price at which Dominick's purchases the goods in its inventory[23]. I use the hourly wage rate in the retail sector and the energy and food CPIs to proxy for economy-wide disturbances to a product's desired price. I also include the duration of a given price spell as an explanatory variable in the latent variable equation. I do so because of an asymmetry in the duration of price spells that end with price increases/decreases in the data. Price spells that end with price increases are shorter than price spells that end in price decreases (3.5 months vs. 4.6 months in the AC Nielsen data, 5.1 months vs. 5.3 months in Dominick's data[24]),

[22]Random effects are a more appealing alternative to modeling heterogeneity in non-linear models because of their computational advantage. I use a fixed effects approach instead to guard against the possibility of endogeneity bias. If prices in a given product-category adjust more frequently than in others, the probability that a given good in this category adjusts is higher, and so is the fraction of other products that adjust in this group. In practice, however, as Table 3 reports, there is little ex-ante heterogeneity in the duration of price spells across different product groups in the data. A random effects specification implemented using a Butler and Moffitt (1982) quadrature approach (results are available from the author upon request), produces hence similar results to those reported in text. Note also that the incidental parameters problem does not arise here as the number of observations within a group is large: 6000 on average.

[23]Wholesale prices are measured imperfectly, as Dominick's reports an average acquisition cost for items in its inventory, as opposed to a replacement cost. Given that we work with monthly data, and inventories are usually replenished more frequently than that in retail stores (inventory/sales ratios are less than 1 at the monthly frequency, according to Bureau of Census data), this is unlikely to affect our results. See Peltzman (1999) for a discussion of this issue.

[24]These statistics are calculated for those observations that are used in estimation (those for which measures of

an asymmetry that presumably arises because some temporary price cuts last more than 4 weeks and are treated, according to the definition of sales I employ as regular price changes, even though they are not. Unaccounted for, this asymmetry can bias the coefficients of those components of the chain's marginal costs that are characterized by trend growth.

The first column of Table 4(A.) reports constrained estimates of the specification above in which I exclude the measures of within-store synchronization from the threshold equations. I report marginal effects of changes in the explanatory variables on the probabilities of a price decrease and a price increase, respectively.

Although statistically significant, disturbances to a product's wholesale price and changes in the wage rate in the retail sector have only a small effect on a product's price adjustment decision[25]. A 10% increase in the good's wholesale price since the previous price adjustment increases the probability of a price increase by only 1.3% and decreases the probability of a price cut by only 0.8%. Similarly, a 10% increase in the retail wage increases the probability of a price increase by 3.5% and lowers that of a price cut by only 2%. I report, in the table, a goodness-of-fit statistic used to evaluate the model's performance. Let $\hat{P}_{it}$ be the probability that a price change occurs, predicted by the model. Let $\hat{a}_{it} = 1$ if $\hat{P}_{it} > 0.5$, and $\hat{a}_{it} = 0$, otherwise, be the predicted price adjustment decision. According to this definition, the model predicts a price adjustment if the probability of a price change is greater than 50%. The goodness-of-fit statistic employed is the proportion of price changes in the data explained by the model: $\Pr(\hat{a}_{it} = 1 | a_{it} = 1)$, where a is the actual adjustment decision. As the table reports, the model explains none of the observed price changes in the data, reinforcing our conclusion that marginal cost shocks alone cannot explain Dominick's price adjustment decisions.

within-store synchronization can be computed using data on at least 5 observations in a given period).

[25] Although included in all specifications, I do not report the effect of changes in the food and energy CPIs as their effect is statistically insignificant.

I allow next (column II of Table 4(A.)) the good-specific thresholds to vary with the three measures of within-store synchronization. Note first that there is little evidence of store-wide synchronization of the prices of different goods: the estimated marginal effects are close to 0. In contrast, prices of goods in a given product or manufacturer category do tend to adjust in tandem. An increase in the fraction of remaining prices that change in a given product category from 0 to 1 increases the probability that a given product will experience a price cut by 5% and that a price increase by 29%, thereby increasing the probability of a price change by 34%. Synchronization is even stronger for goods in a given manufacturer category. An increase in the fraction of price changes of the remaining goods produced by a given manufacturer from 0 to 1 increases the probability that the good in question will also adjust by 55%. The goodness-of-fit statistic increases considerably once measures of within-store synchronization are accounted for: the model now correctly explains 17% of the actual price changes observed in the data. In fact, these results provide only a lower bound of the importance of within-store synchronization. As discussed earlier, I work with Dominick's chain-wide price series, defined for each good as the price of the store with the least number of missing observations. To the extent to which stores enjoy some freedom in the regular price they charge, results would have been even stronger if data on a single store's price would have been used to calculate measures of within-store synchronization.

AC Nielsen

An advantage of this second source of price data is the presence of an additional dimension, as observations on the prices of goods within, but also across, 31 different stores are now available. Let i index goods, s stores and t time-periods. The measures of synchronization I employ are (i) the fraction of a particular good's prices changed by other stores in period t (across-store synchronization); (ii) the fraction of remaining prices that change in store s (within-store synchronization); as well as (iii) the proportion of price changes within a good-category in a given store. All

these statistics are again computed excluding the decision of the good/store whose price adjustment is modeled. Moreover, because some of the stores in the data belong to a single chain, I compute measures of across-store synchronization for stores other than those in the chain a particular store belongs to[26].

Information on wholesale prices is unavailable for the AC Nielsen data. I use instead the average price of the store's competitors (excluding again those stores that belong to a particular chain) as a measure of good-specific disturbances to the marginal cost of selling a product. Results of specifications similar to those discussed above are reported in Table 4(B.).

As the table indicates, the conclusions obtained in the case of Dominick's data carry through to this sample as well. Changes in the average competitors' prices and aggregate marginal cost disturbances again explain none of the adjustment decisions. An increase in the average price of other stores by 10% only increases the probability of a price increase by 4%. Similarly, a 10% increase in the wage in the retail sector only increases the probability of a price increase by 17%. In contrast, the model's performance improves considerably when measures of synchronization are included in the threshold equations.

As in the case of Dominick's data the fraction of price changes within the entire store has little effect on a given product's adjustment decision. In contrast, synchronization within a product-category is important: an increase in the fraction of adjusters within a product-category from 0 to 1 increases the probability that a given product will also experience a price cut by 34%, and a price increase by 62%. The goodness-of-fit statistic improves considerably: the model correctly predicts 22.5% of the price changes observed in the data. Finally, the data provides evidence of minor, but statistically significant, across-store synchronization: an increase in the fraction of other stores that adjust a particular product's price from 0 to 1 increases the probability that a given store will also

[26]I do not calculate the fraction of products belonging to a given manufacturer whose prices change in a given period as only 115 goods are present in this dataset and too few goods per manufacturer are available.

adjust the price of this good, albeit by only 4%. I thus conclude that the result that products within a good-category synchronize is not merely driven by aggregate or good-specific demand or marginal cost disturbances (these would lead other stores to also adjust prices), or store-specific shocks (all prices of a particular store would adjust in tandem if this were the case).

These empirical results do not, on their own, provide evidence to support the hypothesis that stores in our data face increasing returns to scale in their price adjustment technology: within-store synchronization can also arise in the presence of interactions in the firm's profit function between the prices of various goods, as in the model of Sheshinski and Weiss (1992). This alternative source of synchronization is more difficult to reconcile, however, with the large number of small price changes observed in the data. As will become clear below, any source of interactions in a multi-product firm's profit function will, as long as it can generate a substantial number of small price changes, deliver implications similar to those of the model in the next section.

4. A Multi-Product State-Dependent Pricing Model

A. Model Economy

Throughout, let s_t denote the event realized at time t, $s^t = \{s_0, s_1, ..., s_t\}$ the history of events up to this period and $\pi(s^t)$ the probability of a particular history as of time 0. The economy is populated by a continuum of consumers and a continuum of monopolistically competitive firms, both of mass 1. Consumers are identical, while firms (indexed by z) differ according to their productivity level, and hence the prices they charge for otherwise identical products. Each firm sells two products, indexed by 1 and 2. I first discuss the problem of the representative consumer, that of the firm, and then define an equilibrium for this model economy.

Consumers

Consumers' preferences are defined over leisure and a continuum of imperfectly substitutable goods. The consumer sells part of her time endowment to the labor market and invests her wealth in

one-period shares in firms. In equilibrium, identical consumers own equal shares of all the economy's firms. The representative consumer's problem is to choose, given prices, how to allocate her income across the different goods available for consumption and how much to work:

$$\max_{\{c^1(z;s^t),c^2(z;s^t)\},n(s^t),b} \sum_{t=0}^{\infty} \beta^t \sum_{s^t} \pi(s^t) U(c(s^t), n(s^t)),$$

subject to

$$\int_0^1 \left[p^1(z,s^t)c^1(z,s^t) + p^2(z,s^t)c^2(z,s^t)\right] dz = w(s^t)n(s^t) + \Pi(s^t),$$

where

$$c(s^t) = \left(\int_0^1 \left(c^1(z,s^t)^{\frac{\theta-1}{\theta}} + c^2(z,s^t)^{\frac{\theta-1}{\theta}}\right) dz\right)^{\frac{\theta}{\theta-1}}$$

is an aggregator over the different varieties of goods that the household consumers, $n(s^t)$ is the supply of labor, $w(s^t)$ the nominal wage rate, $\Pi(s^t)$ the profits the consumer receives from her ownership of firms, $p^1(z,s^t)$ and $p^2(z,s^t)$ are the prices of each good and θ is the elasticity of substitution across goods. Notice that I have assumed that the elasticity of substitution across goods sold by a single firm is equal to the elasticity of substitution across goods sold by different firms.

Firms

Firms produce output using a technology linear in labor:

$$y^i(z,s^t) = \phi^i(z,s^t) l^i(z,s^t),\ i = 1,2,$$

where the firm's technology, $\phi^i(z,s^t)$, evolves according to

$$\log \phi^i(z,s^t) = \rho \log \phi^i(z,s^{t-1}) + \varepsilon^i(z,s^t),\ i = 1,2,$$

and $\varepsilon(z,s^t) \in [\varepsilon_{\min}, \varepsilon_{\max}]$ is a random variable, uncorrelated across firms, goods and time-periods. Firms operate along their consumers' demand schedules, derived as solutions to the consumer's problem discussed above:

$$c^i(z,s^t) = \left(\frac{p^i(z,s^t)}{P(s^t)}\right)^{-\theta} c(s^t),$$

where $P(s^t)$ is the price index in this economy, defined as a consumption-weighted average of the prices in this economy:

$$P(s^t) = \left(\int_0^1 \left[p_t^1(z,s^t)^{1-\theta} + p_t^2(z,s^t)^{1-\theta}\right] dz\right)^{\frac{1}{1-\theta}}.$$

I assume that firms face fixed menu costs of resetting prices. Any time at least one (or both) of the two prices change, the firm must hire ξ additional units of labor. Let $q(s^t) = \beta^t \frac{U_c(c(s^t),n(s^t))}{U_c(c(s^0),n(s^0))}$, where U_c is the marginal utility of consumption, denote the t-period stochastic discount factor. The firm's problem is to maximize

$$\sum_{t=0}^{\infty} \sum_{s^t} \pi(s^t) q(s^t) \Pi(z,s^t),$$

where

$$\Pi(z,s^t) = \sum_{i=1,2} \left(\frac{p^i(z,s^t)}{P(s^t)}\right)^{-\theta} \left(\frac{p^i(z,s^t)}{P(s^t)} - \frac{w(s^t)}{\phi^i(z,s^t)P(s^t)}\right) c(s^t) -$$
$$-\xi \frac{w(s^t)}{P(s^t)} \mathcal{I}_{p^1(z,s^t) \neq p^1(z,s^{t-1}) \text{ or } p^2(z,s^t) \neq p^2(z,s^{t-1})},$$

and $\mathcal{I}$ is an indicator function. The last term of this expression is the increase in the firm's wage bill if it decides to adjust any of its two prices.

B. Equilibrium

I introduce money by assuming that nominal spending must be equal to the money stock[27]:

$$\int_0^1 \sum_{i=1,2} p^i(z,s^t) c^i(z,s^t) dz = M(s^t)$$

The money supply growth rate $g(s^t) = \frac{M(s^t)}{M(s^{t-1})}$ evolves over time according to an AR(1) process:

$$\log g(s^t) = \delta \log g(s^{t-1}) + \eta(s^t),$$

[27]See Rotemberg (1987) for a transaction technology that gives rise to this particular specification of money demand.

where η is an iid $N(0, \sigma_\eta^2)$ disturbance.

The equilibrium is a collection of prices and allocations: $p^i(z, s^t), w(s^t), P(s^t), c^i(z, s^t), c(s^t)$, $n(s^t), l^i\left(z, s^t\right), y^i(z, s^t)$ such that, taking prices as given, consumer and firm allocations, as well as firm prices solve the consumer and firm problems, respectively, and the labor, goods, and money markets clear.

C. Computing the Equilibrium

I normalize all nominal variables by the money stock in the economy, e.g., $\tilde{P}(s^t) = \frac{P(s^t)}{M(s^t)}$, in order to render the state-space of this problem bounded. Let $\tilde{p}^i_{-1}(z, s^t) = \frac{\tilde{p}^i(z, s^{t-1})}{M(s^t)} \in \mathcal{P}$ be a firm's (normalized) last period's price and $\Phi = [\frac{\varepsilon_{\min}}{1-\rho}, \frac{\varepsilon_{\max}}{1-\rho}]$ the support of the distribution of technology levels in the economy. The aggregate state of this economy is an infinite-dimensional object, consisting of the growth rate of money: $g(s^t)$, but also of the endogenously varying joint distribution of last period's firm prices and technology levels. Let $\boldsymbol{\mu} : \mathcal{P}^2 \times \Phi^2 \to [0, 1]$ denote this distribution and Γ it's law of motion: $\boldsymbol{\mu}' = \Gamma(g, \boldsymbol{\mu})$. Finally, let $\boldsymbol{\phi} = (\phi^1, \phi^2)$ be a vector of a firm's technology levels and $\mathbf{p}_{-1} = \left(\tilde{p}^1_{-1}, \tilde{p}^2_{-1}\right)$ collect the firm's last period's nominal prices.

Let $V^a(\boldsymbol{\phi}; g, \mu)$ and $V^n(\mathbf{p}_{-1}, \boldsymbol{\phi}; g, \boldsymbol{\mu})$ denote a firm's value of adjusting and not adjusting its nominal prices, as a function of its last period's prices and current technology, as well as the aggregate state of the economy. These two functions satisfy the following system of functional equations:

$$V^a(\boldsymbol{\phi}; g, \boldsymbol{\mu}) = \max_{\mathbf{p}} \left(\sum_{i=1,2} \left(\frac{\tilde{p}^i}{\tilde{P}} - \frac{\tilde{w}}{\phi^i \tilde{P}} \right) \left(\frac{\tilde{p}^i}{\tilde{P}} \right)^{-\theta} c - \xi \frac{\tilde{w}}{\tilde{P}} + \beta \int \frac{U_{c\prime}}{U_c} V(\mathbf{p}'_{-1}, \boldsymbol{\phi}'; g\prime, \boldsymbol{\mu}\prime) dF(\varepsilon^1, \varepsilon^2, \eta) \right)$$

$$V^n\left(\mathbf{p}_{-1}, \boldsymbol{\phi}'; g, \boldsymbol{\mu}\right) = \sum_{i=1,2} \left(\frac{\tilde{p}^i_{-1}}{\tilde{P}} - \frac{\tilde{w}}{\phi^i \tilde{P}} \right) \left(\frac{\tilde{p}^i_{-1}}{\tilde{P}} \right)^{-\theta} c + \beta \int \frac{U_{c\prime}}{U_c} V(\mathbf{p}'_{-1}, \boldsymbol{\phi}'; g\prime, \boldsymbol{\mu}\prime) dF(\varepsilon^1, \varepsilon^2, \eta),$$

where $V = \max(V^a, V^n)$ is the firm's value function and $\mathbf{p}$ is a vector of nominal prices the firm chooses every time it adjusts. The laws of motion for the state variables are:

$$\boldsymbol{\mu}\prime = \Gamma(g, \boldsymbol{\mu}), \phi^{i\prime} = \phi^{i\rho} \exp\left(\varepsilon^i\right), g\prime = g^{\delta} \exp(\eta)$$

$$\tilde{p}_{-1}^{i\prime} = \begin{cases} \frac{\tilde{p}^i}{g} \text{ if adjust} \\ \frac{\tilde{p}_{-1}^i}{g} \text{ otherwise} \end{cases}$$

Note that $\tilde{w}$, $\tilde{P}$, c are functions of the aggregate state as well, i.e., $\tilde{w} = \tilde{w}(g, \boldsymbol{\mu})$ etc., although I do not explicitly indicate this dependence in order to conserve on space. The unknowns in this problem are the following functions: $V^a(), V^n(), c(), \tilde{w}(), \tilde{P}(), \Gamma()$. To solve this system of functional equations, I (i) approximate the infinite-dimensional distribution $\boldsymbol{\mu}$ with a small number of its moments, following the suggestion of Krusell and Smith (1997); (ii) replace the unknown functions with a linear combination of orthogonal polynomials; and (iii) solve for the unknown coefficients on these polynomials by requiring that the system of six functional equations (the Bellman equations, as well as the equilibrium conditions) be exactly satisfied at a finite number of nodes along the state-space. A technical appendix discusses the solution method in more detail.

5. Quantitative Results

A. Calibration and Parametrization

I parameterize the utility function as

$$U(c, n) = \log(c) - \psi n.$$

This specification follows Hansen (1985) by assuming indivisible labor decisions implemented with lotteries. I set the length of the period to one month, and therefore choose a discount factor $\beta = .997$. I choose ψ to ensure that in the absence of aggregate shocks households supply $1/3$ of their time to the labor markets. To calibrate the process characterizing the growth rate of the money supply, I estimate an AR(1) process for the growth rate of M1 for the US economy for 1985-1997, the years for which the micro-price data used to calibrate the model is available.

I next discuss the choice of θ, the elasticity of substitution across varieties. It has been widely documented that price-cost margins are large in the retail industry. Nevo (2001) uses a

structural econometric model in order to estimate demand elasticities for ready-to-eat cereals sold in a representative sample of supermarkets. His estimates of markup ratios range from 1.4 to 2. Barsky et. al. (2000) use an indirect method to measure markup ratios for Dominick's products. They argue that the ratio of the price of a branded product to that of a generic good provides a lower bound for the markup ratio on nationally branded products. They find that these lower bounds range from 1.4 to 2.1 for most of the categories of products sold by Dominick's. Chevalier, Kashyap and Rossi (2003) estimate price elasticities using the quantity and price data for Dominick's dataset. Most of their elasticity estimates range between 2 and 4. Based on this evidence, I set $\theta = 3$, implying a markup ratio of 1.5. Table 6 summarizes the choice of parameter values I assign the model.

The rest of the parameters are calibrated: ξ – the size of the fixed costs incurred by the firm when it changes its menu of prices, ρ – the parameter that governs the persistence of marginal cost shocks, as well as the distribution of technology shocks. I choose these parameters in order to match the salient properties of the micro-price data discussed in Section 2. I target an average duration of price spells of 4.5 months, an average value for the two datasets; an average size of price changes of 9%; a standard deviation of price changes of 12%; and a kurtosis of 4.5. I also require that the model generates 30% 'small' (less than one-half of the mean) price changes, as well as a 24-month 'deviance ratio' of 1.02. Additional "over-identifying" checks will be used to gauge the persistence of marginal cost shocks in the data. Table 5 reports the choice of moments used to calibrate the model economy.

To calibrate the distribution of idiosyncratic technology shocks, I assume that shocks ε_t are drawn from the following parametric family of distributions:

$$\varepsilon_t = \begin{cases} -b_t \varepsilon^{\max}, & \text{with} \quad p = \frac{1}{2} \\ b_t \varepsilon^{\max}, & \text{with } 1 - p = \frac{1}{2} \end{cases}$$

where b_t is a random variable drawn from a Beta distribution with parameters α_1 and α_2. The distribution of technology shocks is thus symmetric around zero, and flexible enough to enable the model to reproduce the distributional features of the data.

B. Results

Benchmark Model

I solve for the unknown parameters: $\rho, \xi, \alpha_1, \alpha_2, \varepsilon^{\max}$, by minimizing the sum of squared log-deviations of the model-generated moments from the six targets in Table 5. The last rows of Table 6 (the column labeled Benchmark model) reports the calibrated parameter values.

Marginal cost shocks are fairly transitory: $\rho = 0.5$. Firms pay a menu cost equal to 1.2% of their steady-state labor bill (0.8% of revenue) every time they undergo a new price change, a number close to that reported by Levy et. al. (1997) in a study of the price adjustment costs of five large supermarkets. The distribution of technology shocks is highly leptokurtic, with a kurtosis in excess of 20 and a variance of 2.7×10^{-3} ($\alpha_1 = 0.05$, $\alpha_2 = 1.30$, $\varepsilon^{\max} = 0.4$).

Before proceeding to analyze the model's performance, I briefly discuss the consequences of the assumption that the firm faces a fixed cost of changing an entire menu of prices, as opposed to a given price on the menu. Figure 2 plots the firm's value of adjusting its price, as well as the value of inaction, in the $(\tilde{p}^1_{-1}, \tilde{p}^2_{-1})$ space (prices are expressed as log-deviations from the optimum). Conditional on adjusting and paying the menu cost, the firm's value is independent of its last period's prices. The value of inaction is falling however as the two prices deviate away from their optima. The region in the center, in which both prices are close to their optima and the value of inaction is greater than that of adjustment, defines the inaction region. The firm's adjustment decision thus depends, in a multi-product setting, on how far both prices are away from their respective optima.

Figure 3 performs a small comparative statics exercise. I illustrate in this figure the size of the inaction regions for different values of the productivity level of the firm. Similar to what Golosov

and Lucas (2004) find, firms are more willing to adjust their prices in periods when their technology is higher: firms prefer to make hay (set optimal prices) while the sun is shining. Note in the left panel of the figure that differences in productivity levels across the two goods a firm sells induce differences in the weight a particular good is assigned in an adjustment decision; because the first good is produced with a worse technology, the firm is willing to tolerate larger deviations of this good's price from the optimum than it does for the second product. Finally, note in the figure why interactions in the costs of adjusting prices can generate a number of small price changes: as long as a given product the firm sells is subject to a large marginal cost disturbance, the firm will adjust both prices, independently of how large a price change it desires for the second product.

I next evaluate the model's performance quantitatively. Note, in Table 5, that the model is successful at matching the salient properties of the microeconomic data documented in Section 2, with all model-based moments close to their targets. In particular, the kurtosis of the distribution of price changes is 4.3, and 33% of price changes are less than 4.4% in absolute value ($\frac{1}{2}$ the mean of absolute value of non-zero price changes). The model also does well in matching other, 'over-identifying' restrictions used to infer the accuracy of the estimate of technology shock persistence. The 12-month deviance ratio is equal to 0.88, a value equal to the average of that in the two datasets of scanner prices. The fraction of two consecutive price changes in the same direction is greater, however, in the model (0.52), than in the data (0.37), suggesting that I over-estimate the persistence of shocks, but this statistic is sensitive to the definition of 'sales' employed.

Given the model's success in matching microeconomic features of the data, I next turn to its aggregate implications. Table 7 (Benchmark, SDP column) reports the volatility and persistence of HP-filtered output in simulations of the model. For comparison, I also solve a Calvo-type time-dependent model, identical in all respects to the original model, in which the timing of price changes is outside the control of the firm. The multi-product, menu-costs setup generates business cycle

fluctuations from monetary disturbances almost as large as those in the time-dependent model: the standard deviation of output is equal to 0.61% (0.75% in the Calvo setup). Business cycles are equally persistent in the two models: the autocorrelation of output is equal to 0.94.

Standard State-Dependent Pricing Model

I next compare the results above to those one obtains in a standard menu cost economy[28] with idiosyncratic marginal cost disturbances. I abstain from the multi-product features discussed earlier and assume that single-product firms face fixed costs of adjusting their prices. I assume, in addition, as is standard in earlier work, that the firm's technology shocks, ε_t, are drawn from a Gaussian distribution[29] with mean 0 and variance σ^2. Three parameters must be calibrated in this model: ξ, ρ, σ^2. I set ρ equal to 0.5, as in the benchmark model[30]. The other two parameters are jointly chosen so that the frequency of price changes, and the mean absolute value of non-zero price changes are equal to those in the data: ($\xi = 0.96\%$ of the SS labor bill, $\sigma^2 = 2.2 \times 10^{-3}$).

As the third column (Standard SD) of Table 5 indicates, the standard menu cost model fails to accord with the micro-price data along two dimensions: it generates no small price changes (no price change is less than 4.5% in absolute value) and produces a kurtosis of price changes much smaller than that in the data (1.3). I plot, in Figure 4, the distribution of price changes, conditional on adjustment, implied by the standard model, as well as the multi-product model calibrated above. In contrast to the multi-product model, which produces a unimodal, leptokurtic distribution of price changes, similar to that observed in the data, the standard model generates a bi-modal distribution, with no price changes close to zero, clearly failing to capture higher-order moments of the data.

The aggregate performance of the standard model, is, as a result of its inability to accord with the micro-economic evidence regarding the distribution of non-zero price changes, lackluster

[28] As exposited, for example, in Golosov and Lucas (2004).
[29] Similar results obtain if we were to assume uniformly distributed shocks. See Gertler and Leahy (2005).
[30] The results we are about to report are insensitive to how persistent technology shocks are.

at best. The model generates output fluctuations that are 5 times less volatile than those in a time-dependent model (the standard deviation of HP-filtered output is only 0.15%, compared to 0.73% in a Calvo model). Business cycles are also less persistent (an autocorrelation of only 0.75) than in the time-dependent model. These results accord to those of Golosov and Lucas (2004), Caplin and Spulber (1986), and Gertler and Leahy (2005), who find that standard state-dependent models generate, despite nominal rigidities at the firm level, small (if any) business cycle fluctuations from monetary disturbances.

C. Counterfactual Experiments

Two features of the standard menu-cost model are responsible for its inability to produce business cycle fluctuations. First, as illustrated by Dotsey, King and Wolman (1999), firms synchronize price changes in response to large aggregate shocks and render the aggregate price level more flexible than in time-dependent economies. Second, the identity of adjusters also varies endogenously in response to nominal disturbances. In times of, say, a monetary expansion, firms are more likely to adjust if their idiosyncratic disturbances also call for price increases and reinforce the desire to reprice stemming from the aggregate shock. Because the distribution of idiosyncratic shocks, conditional on adjustment, varies in response to aggregate disturbances, the response of the price level to a money shock is amplified in menu-cost models relative to time-dependent economies[31].

I quantify the role of synchronization and self-selection using counterfactual experiments in which I shut down each of these two effects. In the first experiment I assume that, conditional on adjustment, firms price using the policy rules that are optimal in a state-dependent world. Instead of allowing them to endogenously choose the timing of their price changes, I select a constant fraction of firms to adjust prices, in an iid fashion, following the timing assumption of a Calvo model. I

[31]This second mechanism is similar to the one that arises in regression equations in which the sample is non-random: failure to deal with sample selection leads to biased estimates of ordinary least squares estimates.

aggregate firm decision rules according to the equilibrium conditions in order to generate time-series of aggregate prices and quantities. This counterfactual, by holding constant the fraction and identity of the adjusting firms, allows us to gauge the combined role of firm synchronization and self-selection in reducing output variability in menu cost models.

Note in Table 7 that the standard deviation of output in this counterfactual is 2.1 times larger than in the original Benchmark model with multi-product firms. In contrast, shutting down synchronization and self-selection in the standard single-product model generates output fluctuates that are 9 times larger than originally. Clearly, synchronization/self-selection has a much stronger effect in the standard menu-costs models than in a model with multi-product firms and disturbances sufficiently leptokurtic to match the distribution of price changes observed in the data.

I next solve a second counterfactual, in order to pinpoint the exact source of the monetary neutrality in the standard SDP model. If synchronization is the sole reason behind the lack of output variability in the state-dependent model, we would expect that a counterfactual simulation in which we shut down the self-selection effect (by selecting adjusters in an exogenous fashion), but turn on the synchronization effect (by allowing the fraction of adjusters to vary endogenously, as in the original model) to produce as little output variability as in the original menu-cost models. As Table 7 indicates however, this is not the case. Allowing the fraction of adjusters to vary in response to aggregate shocks produces almost as much variability of output as in the previous counterfactual. This suggests that synchronization plays a small role and most of the standard model's weak macroeconomic performance is due to variation in the identity of adjusters in response to aggregate disturbances. This is not surprising, since as Golosov and Lucas (2004) show, most of the adjustment in models with large marginal cost shocks is driven by idiosyncratic disturbances and there is little synchronization in response to aggregate shocks.

To understand why self-selection plays a much more important role in the standard model

(output variability would have been 8.6 times higher in its absence) than in the multi-product model (where it only reduces the volatility of output in half), consider the heuristic example in Figure 5. I assume, in the left panel of the figure, that technology shocks are drawn from a Gaussian density, and, in the right panel, the leptokurtic density calibrated in the benchmark model. Suppose, for simplicity, that the (S, s) bands of price adjustment are symmetric and firms adjust prices whenever they desire price changes greater than 10%. In the absence of aggregate shocks, only those firms whose technology shocks exceed 10% in absolute value will find it optimal to adjust their prices. The inaction region (in the ε_t space) is therefore $[s_0 = -0.1,\ S_0 = 0.1]$, illustrated with solid lines in the two panels of Figure 5.

Consider next the effect of a monetary expansion that increases all firms' desired prices by 5%. This aggregate disturbance changes the range of technology shocks for which firms find it optimal to pay the menu costs from $[s_0 = -0.1,\ S_0 = 0.1]$ to $[s_{high} = -0.05,\ S_{high} = 0.15]$. Consider for example, a firm whose technology shock is $\varepsilon_t = -0.07$. In the absence of an aggregate disturbance, this shock is, on its own, too small to trigger an adjustment as it only generates a desired price increase of 7%. Combined with an additional 5% desired price increase stemming from the monetary shock, the firm's desired price increase is 12%, large enough to force a price change.

The effect aggregate shocks have on the identity of adjusters depends however on the kurtosis of the distribution of technology shocks. When the distribution is Gaussian, the measure of 'marginal' firms whose adjustment decision is affected by the aggregate shock is large (the shaded area in the figure). In contrast, when the distribution is leptokurtic, technology shocks that trigger adjustment are spread further away from the (S, s) bands and aggregate shocks have a small effect on the identity of adjusters.

Gertler and Leahy (2005) illustrate a similar point using a state-dependent model in which technology shocks affect, in a given period, only a fraction of firms in the economy: the slope of

the Phillips curve that their model generates depends on the fraction of firms that are subject to technology disturbances in a given period, or in our language, the kurtosis of technology shocks in the economy. The discussion above should also make it clear that one can render the state-dependent model's business cycle fluctuations as large as those of a time-dependent model by increasing the kurtosis of the distribution of technology shocks. If technology shocks were, say, drawn from a three point distribution and were either equal to zero, or extremely large in absolute value, firms subject to non-zero technology disturbances would always adjust, and the model would behave similar to a Calvo model[32]. Such a parametrization would, however, generate a distribution of price changes inconsistent with that observed in the data.

Note however that self-selection, although muted in a model with leptokurtic marginal cost disturbances, still plays an important role in the Benchmark model and reduces the volatility of output in half. Why then, does the state-dependent model produce business cycle fluctuations that are of similar magnitude as those in the Calvo model where this effect is entirely absent?

The explanation lies in a feature of state-dependent models that has not been, to my knowledge, quantitatively explored in earlier work. If shocks to the growth rate of the money supply are, as in the data, persistent, a monetary expansion today forecasts additional increases of the money stock in future periods. Forward-looking firms take into account these forecasts and respond much stronger to the nominal shock than they would in a flexible-price world. Several devices have been proposed to remedy this problem: assuming that a fraction of firms is backward looking and behaves suboptimally, automatic indexation of prices to last period's inflation rate, firm-specific factors of production or non-constant elasticities of demand. It turns out however that front-loading in response to future expected marginal cost disturbances is mainly a feature of time-dependent sticky

[32]This is of course, only true is shocks are transitory, or at least not too persistent, as it is the growth rate of technology, rather than the shock itself, that matters for the firm's price adjustment decision.

price models, and plays a much smaller role when the timing of price changes is endogenous[33].

Figure 6 plots the price functions, conditional on adjustment, in the Calvo and menu-cost models, expressed as log-deviations of the optimal price from the one that firms would set in a flexible price world in which prices would respond one-for-one to the monetary disturbance. Clearly, Calvo firms front-load current prices much more aggressively in response to future expected increases in the money supply than state-dependent firms do: a 3% increase in the money growth rate triggers a price increase by Calvo firms that is almost 5% larger than what is optimal in a flexible price world. In contrast, the state-dependent firm's price increases by only 1%.

To understand why state-dependent firms refuse to front-load prices, even though they expect future increases in the growth rate of the money supply, consider Figure 7 in which I plot the two types of firms' values of adjustment and inaction, as a function of one of the two goods' past prices. Note two features of this figure. First, a firm's value is higher under state-dependent pricing because a state-dependent firm, although it adjusts as infrequently as a time-dependent firm, chooses the timing of its price changes optimally. More importantly, a state-dependent firm's value of inaction is much less sensitive to deviations of the past price from the optimum than is the value of inaction of a Calvo firm. This result makes intuitive sense. If a state-dependent firm finds itself with a suboptimal price in a given period, it can always exercise its option to adjust. Its losses are therefore bounded by the size of the menu cost. In contrast, a time-dependent firm pays a hefty price every time its nominal price is suboptimal: because it has to wait for an exogenously fixed number of periods before it gets to reset its price, it will incur much larger losses from a suboptimal price relative to what a state-dependent firm would. A time-dependent firm's incentive to offset future deviations of its price from the optimum is therefore larger than that of a state-dependent firm and it adjusts

[33]Dotsey, King and Wolman (1999) point out this difference in the optimal price functions of time- and state-dependent firms, but find it quantitatively small in their baseline parametrization with no serial correlation in the growth rate of the money supply and time-varying costs of price adjustment.

more aggressively in response to persistent shocks to the growth rate of the money supply.

D. Leptokurtic Shocks or Multi-Product Firms?

I finally isolate the role of the two departures from the standard menu costs model I have introduced and evaluate the contribution of each in improving the model's macroeconomic performance. I first solve a model with economies of scale in the costs of price adjustments, in which idiosyncratic shocks are drawn from a Gaussian distribution. The only parameter I calibrate is the size of the menu costs, chosen to ensure the same frequency of price changes as in the Benchmark model. The rest of the parameters are set equal to their calibrated values in the Benchmark model. Notice in the 4th column of Table 5 that the model fails to match the kurtosis of the distribution of price changes in the data, although it does generate a large number of small price changes (26%)[34]. A departure from the assumption of Gaussian shocks is therefore crucial in reproducing the kurtosis of the distribution of technology shocks in the data. Table 7 presents this model's aggregate implications: although the volatility of output increases relative to that in models with no interactions in the costs of price adjustment (0.25% vs. 0.15%), business cycle fluctuations are much smaller than those in time-dependent models. The reason the model's performance does improve is a weakening of the selection effect, which, as the counterfactual experiments indicate, is almost twice weaker than in the standard single-product economy.

I next solve the problem of a single-product firm in which technology shocks are drawn from the distribution assumed in the Benchmark model. Leptokurtic shocks are capable, on their own, to reproduce the kurtosis of price changes in the data. They fail to generate however a large number of small price changes: only 4% of price changes are now less than 4.5% in absolute value. As Table 7 indicates, leptokurtic shocks increase the volatility of output from 0.15% in the Standard model to

[34]I define, here and in the next calibration, a small price change as a price change whose absolute value is less than 4.5%, the average magnitude of price changes in the Benchmark model.

0.46%, a significant improvement, albeit smaller than in the Benchmark model. Both interactions in the costs of price adjustment, as well as leptokurtic shocks are thus a necessary ingredient of a model capable of reproducing the microeconomic evidence and generating sizable business cycle fluctuations.

6. Conclusion

This paper has shown that standard state-dependent pricing models are inconsistent with two facts regarding the behavior of individual good's prices: the large number of small price changes and excess kurtosis of price changes in the data. The large number of small price changes can be reconciled with state-dependent models if multi-product firms face interactions in the costs of adjusting prices: I find indeed substantial evidence that prices of products in narrow product categories within grocery stores adjust in tandem.

I then study the general equilibrium properties of a multi-product menu-cost economy calibrated to accord with this micro-economic evidence, and find that the model can, in fact, generate substantial business cycle fluctuations from nominal disturbances. A key feature of the calibration, the leptokurtic distribution of idiosyncratic disturbances, implies, together with the assumption of economies of scale in the price adjustment technology, that the selection effect that plays an important role in standard menu cost economies is much weaker in this setup. This, as well as the fact that state-dependent pricing firms are less willing to front-load prices in response to expected future changes in the stock of the money supply, ensures that a menu-cost model can generate substantial output variability from monetary disturbances.

References

[1] Barro, Robert, 1972, "A Theory of Monopolistic Price Adjustment," Review of Economic Studies, 39(1): 17-26

[2] Barsky, Robert, Mark Bergen, Shantanu Dutta, and Daniel Levy, "What Can the Price Gap Between Branded and Private Label Products Tell Us About Markups?," mimeo

[3] Benabou, Rolland, "Inflation and Efficiency in Search Markets,"Review of Economic Studies, 59(2): 299-329

[4] Bils, Mark and Peter Klenow, 2004, "Some evidence on the importance of sticky prices," NBER working paper 9069

[5] Bonomo, Marco and Carlos Viana de Carvalho, 2005, "Endogenous Time-Dependent Rules and Inflation Inertia," forthcoming, Journal of Money, Credit and Banking

[6] Burstein, Ariel, 2002, "Inflation and Output Dynamics with State-Dependent Pricing Decisions," mimeo

[7] Butler, J.S., and Rober Moffitt, 1982, "A Computationally Efficient Quadrature Procedure for the One Factor Multinomial Probit Model," Econometrica, 50: 761-764

[8] Caballero, Ricardo and Eduardo Engel, "Explaining Investment Dynamics in US Manufacturing: A Generalized (S,s) approach," Econometrica, 67(4): 783-826

[9] Calvo, Guillermo, 1983, "Staggered Prices in a Utility-Maximizing Framework," Journal of Monetary Economics, 12(3): 383-398

[10] Caplin, Andrew and Daniel Spulber, 1987, "Menu Costs and the Neutrality of Money," Quarterly Journal of Economics, 102(4): 703-725

[11] Caplin, Andrew and John Leahy, 1991, "State-Dependent Pricing and the Dynamics of Money and Output," Quarterly Journal of Economics,106(3): 683-708

[12] Cecchetti, Stephen, 1986, "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines," Journal of Econometrics, 31(3): 255-274

[13] Chevalier, Judith, Anil Kashyap and Peter Rossi, 2003, "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," American Economic Review, 93(1): 15-37

[14] Danziger, Leif, 1999, "A Dynamic Economy with Costly Price Adjustments," American Economic Review, 89(4): 878-901

[15] Dhyne, Emmanuel, et. al., 2005, "Price Setting in the Euro Area. Some Stylized Facts from Individual Consumer Price Data," ECB working paper

[16] Dotsey, Michael, Robert King and Alexander Wolman, 1999, "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," Quarterly Journal of Economics, 114(2): 655-690

[17] Dutta, Shantanu, Mark Bergen and Daniel Levy, 2002, "Price flexibility in channels of distribution: Evidence from Scanner Data," Journal of Economic Dynamics and Control

[18] Fisher, Timothy, and Jerzy Konieczny, 2000, "Synchronization of Price Changes by Multi-Product Firms: Evidence from Canadian Newspaper Prices," Economic Letters, 68: 271-277

[19] Gertler, Mark and John Leahy, 2005, "A Phillips Curve with an Ss Foundation," mimeo

[20] Golosov, Mikhail, and Robert E. Lucas, Jr, 2004, "Menu Costs and Phillips Curves," NBER working paper 10187

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

[22] Hodrick, Robert and Edward Prescott, 1997, "Post-War US Business Cycles: An Empirical Investigation," Journal of Money, Credit and Banking, 29(1): 1-16

[23] Hosken, Daniel and David Reiffen, 2004, "Patterns of Retail Price Variation," Rand Journal of Economics, 35(1): 128-146

[24] Kackmeister, Alan, 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

[25] Kashyap, Anil, 1995, "Sticky Prices: New Evidence From Retail Catalogs," Quarterly Journal of Economics, February, 245-274

[26] Khan, Aubhik and Julia Thomas, 2004, "Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies," FRB Philadelphia Working Paper 04-11

[27] Klenow, Peter and Oleksiy Kryvtsov, 2004, "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent US Inflation?" mimeo

[28] Krusell, Per and Anthony Smith, 1997, "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns", Macroeconomic Dynamics, 1(2): 387-422

[29] Lach, Saul and Daniel Tsiddon, 2005, "Small Price Changes and Menu Costs," mimeo

[30] Lach, Saul and Daniel Tsiddon, 1996, "Staggering and Synchronization in Price-Setting: Evidence from Multiproduct Firms," American Economic Review, 86(5): 1175-1196

[31] Lazear, Edward, 1986, "Retail Pricing and Clearance Sales," American Economic Review, 76: 14-32

[32] Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable, 1997, "The Magnitude of Menu Costs: Direct Evidence from Large US Supermarket Chains," Quarterly Journal of Economics, 112(3): 781-826

[33] Miranda, Mario and Paul Fackler, 2002, "Applied Computational Economics and Finance," MIT press

[34] Nevo, Aviv, 1997, "Measuring Market Power in the Ready-to-Eat Cereal Industry," Econometrica, 69(2): 265-306

[35] Peltzman, Sam, 2000, "Prices Rise Faster Than They Fall," Journal of Political Economy, 108(3): 466-502

[36] Sheshinski, Eytan and Yoram Weiss, 1977, "Inflation and Costs of Price Adjustment," Review of Economic Studies, 44(2): 287-303

[37] Sheshinski, Eytan and Yoram Weiss, 1983, "Optimum Pricing Policy Under Stochastic Inflation: the Multi-Product Monopoly Case," Review of Economic Studies, 59: 331-359

[38] Sheshinski, Eytan and Yoram Weiss, 1992, "Staggered and Synchronized Price Policies Under-Inflation," Review of Economic Studies, 50(3): 513-529

[39] Varian, Hal, 1980, "A Model of Sales," American Economic Review, 70(4): 651-659

[40] Warner, Elizabeth and Rober Barsky, 1995, "The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays," Quarterly Journal of Economics, 110(2): 321-352

[41] Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen, 2004 "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," Review of Economics and Statistics, 86(2): 514-533

Table 1: Distribution of price changes conditional on adjustment

	AC Nielsen	Dominick's
Δp		
mean. %	0.0	1.5
standard deviation, %	13.2	10.4
kurtosis	3.5	5.4
fraction positive changes	0.50	0.65
$\lvert\Delta p\rvert$		
mean, %	10.4	7.7
fraction "small" changes	0.28	0.35
# obs.	7560	23029

Notes:

1. p is the natural logarithm of a store's price
2. Statistics are unweighted, based on all non-zero regular price changes, excluding top and bottom 1%.
3. I define a small price change as any change whose absolute value is lower than 1/2 the mean of $|\Delta p|$.

Table 2: The frequency and persistence of price changes

	AC Nielsen	Dominick's
Duration of Price Spells, months		
mean	4.0	5.2
median	3	3
iqr	3	6
Persistence of price changes		
Deviance Ratio: 12 months	0.96	0.80
Deviance Ratio: 24 months	N/A	1.02
Prob{sgn(Δp') = sgn(Δp)}	0.32	0.41

Notes:

1. Prob{sgn(Δp') = sgn(Δp)} is the probability that the next price change will have the same sign as the current one, computed as an equally weighted average of the probabilities of two consecutive positive/negative price changes.
2. The Deviance ratios are the mean absolute difference in a good's (detrended) price in periods that are 12 (24) months apart, relative to the mean absolute value of non-zero price changes.

Table 3: Fraction of variance explained by ex-ante heterogeneity

	AC Nielsen	Dominick's
Absolute value of price changes		
store	0.04	-
product	0.14	0.16
month	0.05	0.03
Duration of price spells		
store	0.03	-
product	0.12	0.08
month	0.08	0.10

Notes:

1. Fraction of variance attributed to store/product/month fixed effects reported.
2. Product category × manufacturer, as opposed to upc-specific effects are included in Dominick's data.

Table 4: Ordered Probit Marginal Effects

A. Dominick's

	I.		II.	
	price decrease	**price increase**	**price decrease**	**price increase**
Δ wholesale price	-0.08	0.13	-0.07	0.11
	(0.003)	(0.01)	(0.002)	(0.01)
Δ retail wage	-0.21	0.35	-0.25	0.44
	(0.09)	(0.15)	(0.08)	(0.14)
Fraction of price changes:				
within store			0.02	-0.02
			(0.01)	(0.02)
within product category			0.05	0.29
			(0.004)	(0.01)
within manufacturer category			0.10	0.45
			(0.003)	(0.01)
% adjustments explained	0.0		17.1	
Log-likelihood	-62301		-52341	
# obs.	151614		151614	

Notes: marginal effects on probability of a price change reported. Standard errors reported in parantheses.
month and product-category dummies included in threshold equations

Table 4: Ordered Probit Marginal Effects

B. AC Nielsen

	I.		II.	
	price decrease	**price increase**	**price decrease**	**price increase**
Δ avg. competitor price	-0.43	0.37	-0.32	0.34
	(0.02)	(0.14)	(0.02)	(0.02)
Δ retail wage	-1.99	1.73	-1.63	1.73
	(0.49)	(0.43)	(0.43)	(0.46)
Fraction of price changes:				
across stores			0.02	0.02
			(0.01)	(0.01)
within store			0.07	-0.07
			(0.02)	(0.03)
within product category			0.34	0.62
			(0.01)	(0.02)
% adjustments explained	0.0		22.5	
Log-likelihood	-18672		-15825	
# obs.	36400		36400	

Notes: marginal effects on the probability of a price change reported. Standard errors reported in parantheses.
month, product-category and store dummies included in threshold equations

Table 5: Calibration targets

Moments:	**Targets**	**Benchmark**	**Standard SD**	**Multi-product, Gaussian shocks**	**Single-product, Leptokurtic Shocks**
Duration of price spells	4.5	**4.6**	**4.6**	**4.5**	**4.6**
mean (|Δp|), %	9.0	**8.8**	**9.0**	7.4	9.9
std (Δp), %	12.0	**11.8**	9.3	8.4	12.5
kurtosis(Δp)	4.5	**4.3**	1.3	1.7	4.0
Pr(|Δp|<mean(|Δp|)/2)	0.30	**0.33**	0.00	0.26	0.04
Deviance ratio: 24 months	1.02	**1.01**	0.84	0.89	1.12
"Over-identifying" restrictions					
Deviance Ratio: 12 months	0.88	0.88	0.82	0.85	1.02
Probability that next price change is in the same direction	0.37	0.52	0.30	0.35	0.53

Note: entries in bold are the moments used to pin down parameter values for the various calibrations

Table 6: Parameter Values

		Benchmark	Standard SD	Multi-product, Gaussian shocks	Single-product, Leptokurtic Shocks
Parameters not explicitly solved for:					
Common:					
β	discount factor	0.997			
δ	persistence of money shocks	0.79			
σ^2_η	variance of money shocks	$2.05*10^{-5}$			
ψ	marginal disutility from work	2.4			
θ	elasticity of substitution	3			
Calibration specific					
ρ	persistence of technology shocks	-	0.50	0.50	0.50
σ^2_ε		-	-	$2.2*10^{-3}$	-
α_1	Beta(α_1,α_2)	-	-	-	0.05
α_2	Beta(α_1,α_2)	-	-	-	1.3
ε_{max}	upper bound of distb of shocks	-	-	-	0.40
Parameters explicitly solved for					
σ^2_ε	variance of technology shocks	-	$2.2*10^{-3}$	-	-
ξ	menu cost, % of SS labor bill	1.20	0.96	1.00	0.67
α_1	Beta(α_1,α_2)	0.05	-		
α_2	Beta(α_1,α_2)	1.30	-		
ρ	persistence of technology shocks	0.50	-	-	-
ε_{max}	upper bound of shock distribution	0.40	-		

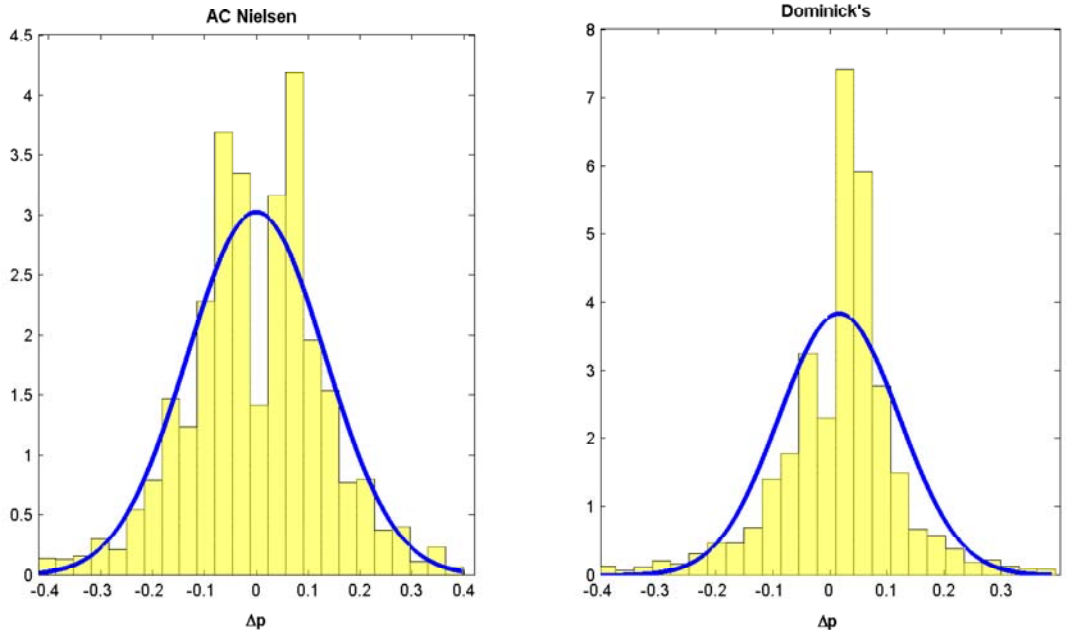
Table 7: Aggregate implications

	Benchmark		Standard SDP		Multi-product Gaussian shocks	Single-product Leptokurtic Shocks
	SDP	Calvo	SDP	Calvo	SDP	SDP
Business Cycle Statistics						
$\sigma(y)$	0.61	0.75	0.15	0.73	0.25	0.46
$\rho(y)$	0.94	0.93	0.75	0.92	0.82	0.92
Counterfactual Experiments						
$\sigma(y)$ with Calvo timing (relative to original model)	2.1	-	9.2	-	4.8	3.0
$\sigma(y)$ with no self-selection (relative to original model)	1.8	-	8.6	-	4.4	2.5

Notes:

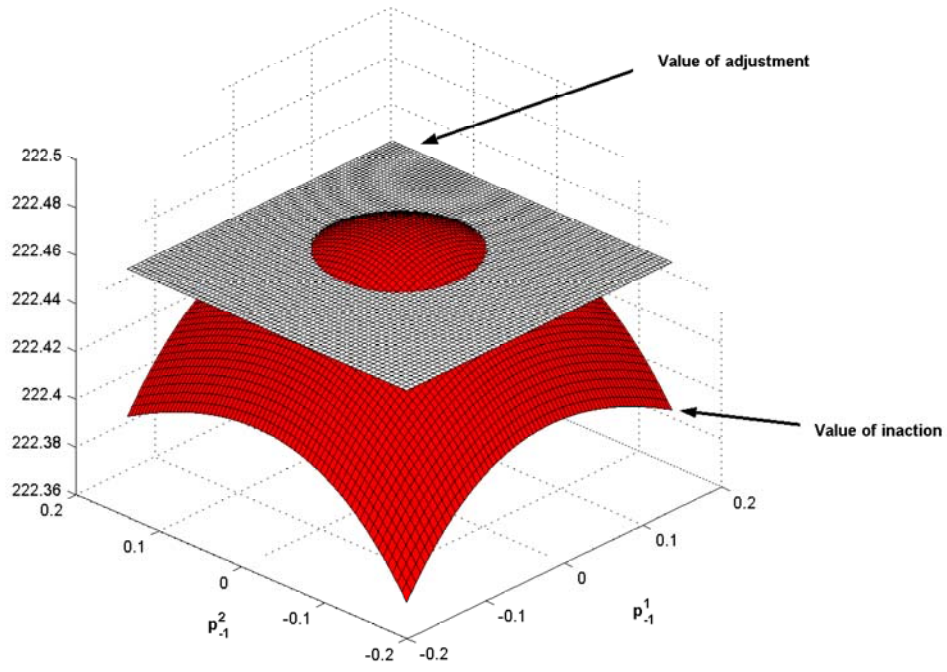
1. output data was detrended using an HP(14400) filter
2. Δ_t is the average price change of adjusting firms
3. F(ε|adjustment) is the distribution of technology shocks conditional on adjustment

Figure 1: Distribution of price changes conditional on adjustment



Note: superimposed is the pdf of a Gaussian distribution with equal mean and variance

Figure 2: Value Functions



Note: all other arguments at steady-state levels

Figure 3: Inaction (Ss) regions for multi-product firms

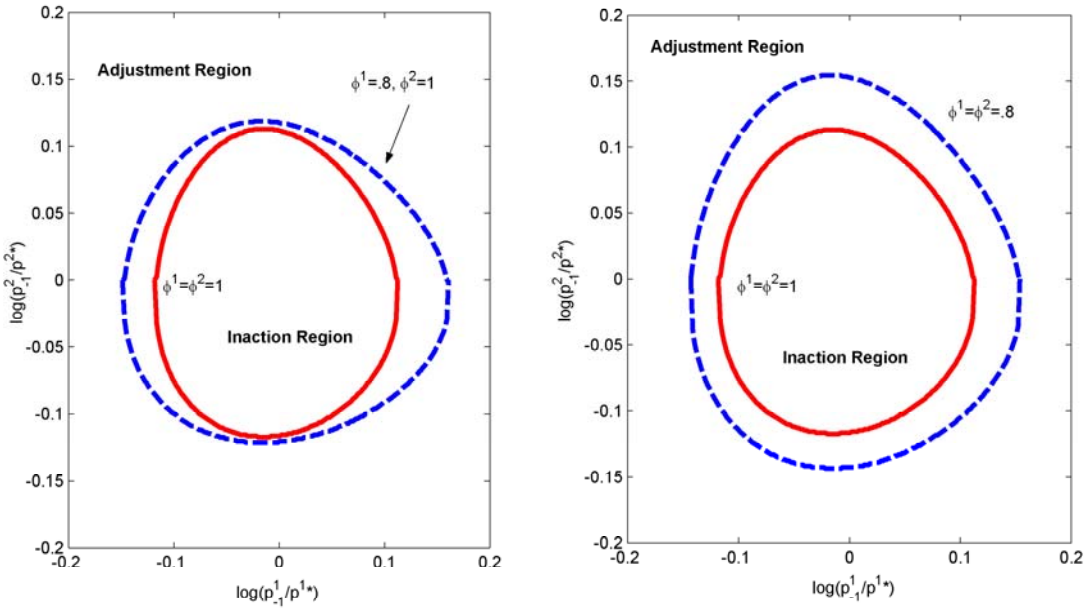
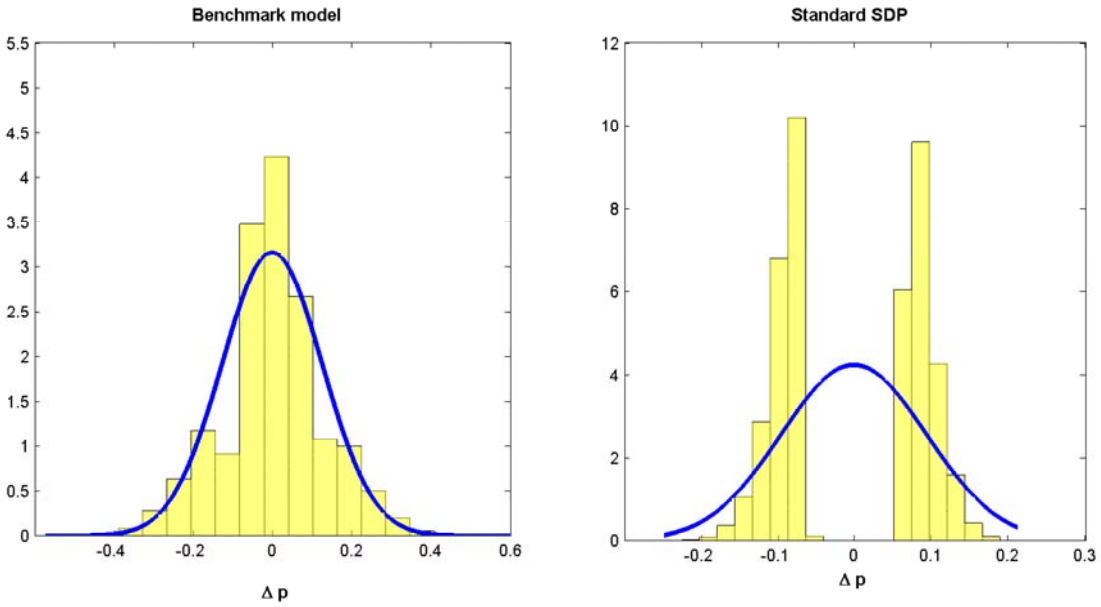


Figure 4: Distribution of price changes conditional on adjustment



Note: superimposed is the pdf of a Gaussian with the same mean and variance

Figure 5: Variation in distribution of technology shocks across adjusters

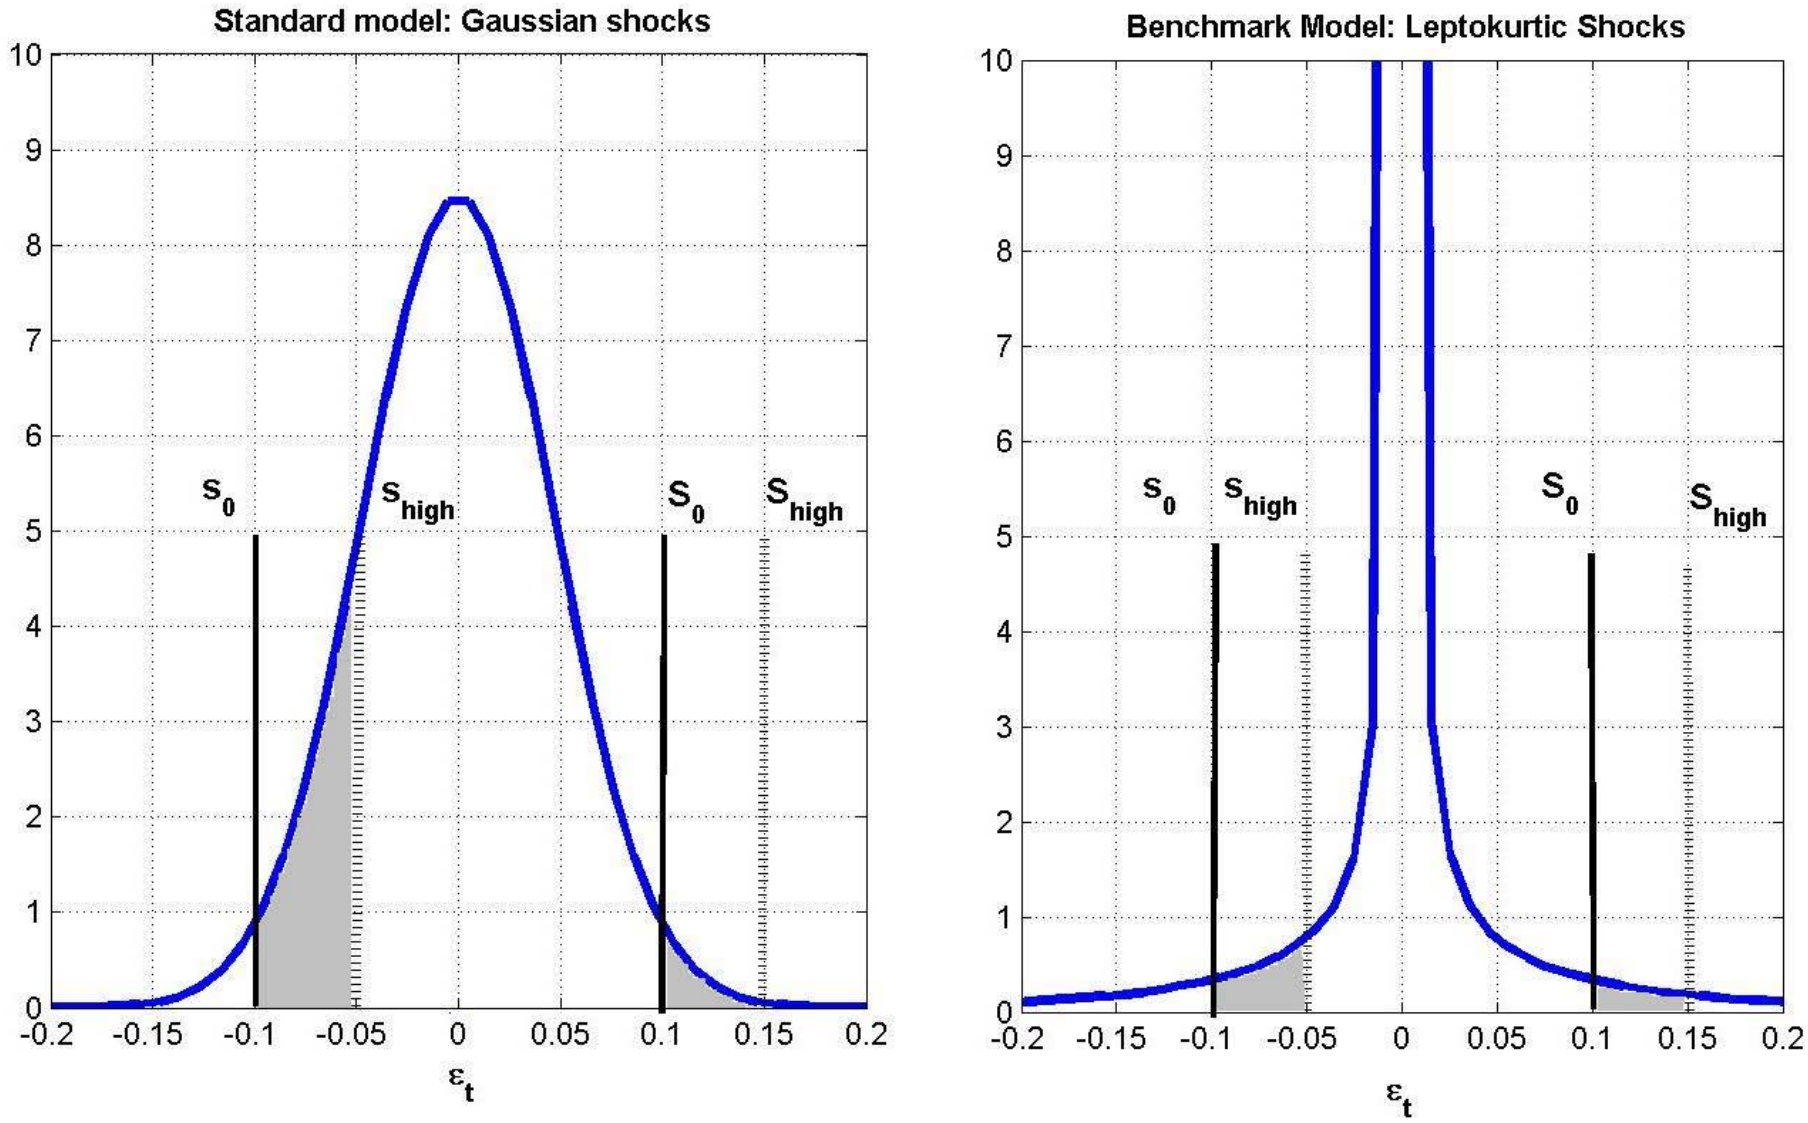
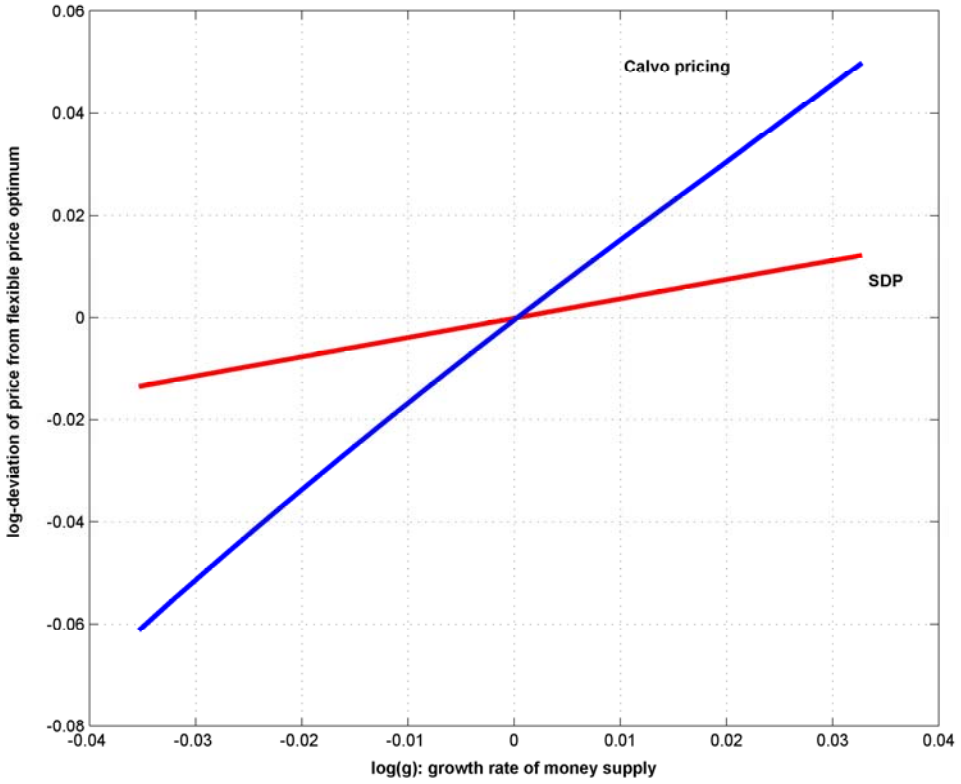
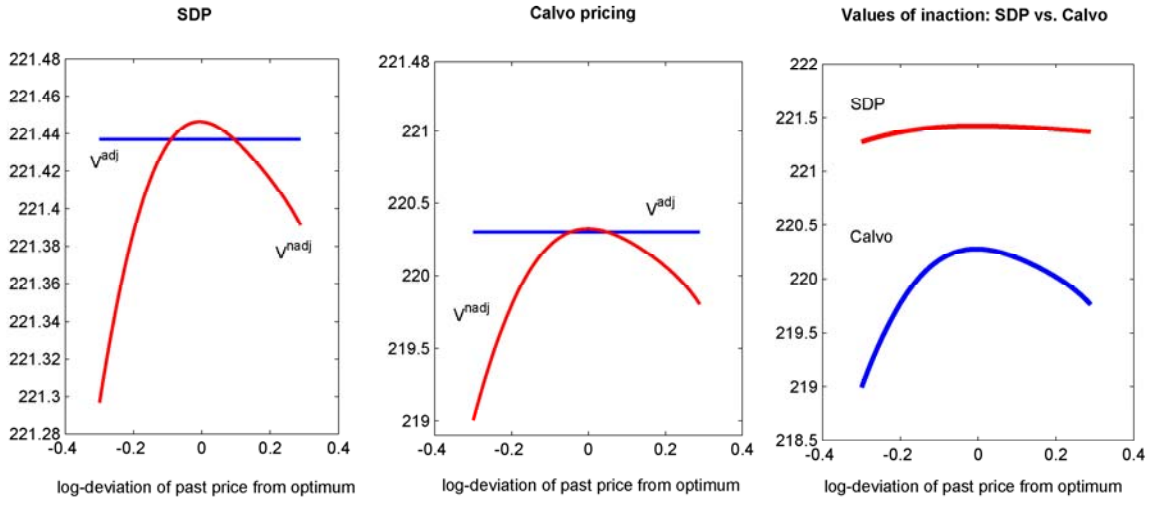


Figure 6: Price functions conditional on adjustment



Note: all other arguments at steady-state values

Figure 7: Value functions



Note: all other arguments set at steady-state levels

Data Appendix
A. Construction of Price Series
AC Nielsen

I use the household-level data to construct weekly price series according to the following algorithm. For each store/good in the sample, I calculate the number (if any) of units sold at a particular price during the course of the week. If the store sells the product at a single price during the week, I assign this value to the weekly price series. If more than one price is available, the weekly price is the price at which the store sold the largest number of units. In case of a tie in the number of units sold at a particular price, the weekly price is the highest price at which the store sells in a given week[1].

Given that I use scanner price data, price observations for a particular store/good are only available when a customer purchases the product in a particular week. The original prices series are therefore frequently interrupted by gaps. I ignore the gaps if these last for four weeks or less, and use the latest available price before the gap to fill in the price series. Because I study the frequency and size of price changes, I require an uninterrupted prices series. If gaps larger than four weeks are present in the data, I keep only the longest spell of uninterrupted price observations, and discard the rest. The original dataset has observations for 124 weeks, 38 stores, and 225 products, but many of these series have a large number of missing observations. To ensure that I work with a more balanced panel, I discard the first 6 and last 5 weeks of the data in which the number of products available is small (less than 60 on average across stores). I also discard goods for which fewer than 10 stores (one-fourth of the original number of stores) have price data at least two-third of the time. Finally, I discard stores for which fewer than 30 products (one-fourth of the number of products that make the cutoff above) are non-missing at least two-third of the time. The resulting sample is one with 113 weeks of data on prices of 115 products in 31 stores.

[1]Intuitively, if the number of units sold at two prices is equal, the highest price is likely to have been in effect for a longer time-period as consumers are more likely to buy at the lowest price.

Dominick's

The Kilts Center for Marketing makes available weekly price quotes for 86 of Dominick's stores. As noted above, this dataset is a by-product of a series of randomized pricing experiments conducted by Dominick's from 1992 to 1993. I only work with the prices of those stores/product categories that were part of control groups in order to avoid treating price changes arising due to experiments as regular price changes[2].

Dominick's stores are divided into three groups: high, low and medium-price stores, depending on the extent of local competition. Prices within, but also across groups, are strongly correlated, as Dominick's sets prices on a chain-wide basis[3]. Given that gaps in prices are a common occurrence for this dataset, I collapse store-wide prices into a single, chain-wide price, in order to reduce the number of missing observations. I work with medium-price stores only, as these account for the largest share of Dominick's stores. From this set of observations, I let the chain-wide price be the price of the store that has the largest number of observations[4]. For each gap present in this series, I fill in the gaps with the price of another store in the chain, whose pricing most closely resembles the price of the original store, provided that data for this store is available during this period, and the two store's prices coincide in the periods immediately before and after the gap[5]. My metric of the similarity of two stores' price policies is the number of periods in which the two stores set identical prices for a given product. On average 1% of price series are imputed using another store's price. The prices of the stores used to fill in missing data coincide with the price of the original store in an average of 96% of time periods for which data on both stores is available. I discard from the sample those goods for which less than 100 uninterrupted weekly observations are available (4979

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

[3]See Peltzman (2000) for a discussion of Dominick's pricing practices.

[4]Prior to this step, I eliminate gaps smaller than 4 weeks following the algorithm used for the AC Nielsen data.

[5]This last constraint is dictated by our unwillingness to confound changes in stores with changes in the price of a particular store.

out of the original 5700 goods make this cutoff).

B. Time-Aggregation and Treatment of Sales

Retail prices are characterized by a large number of temporary price markdowns (sales). Kackmeister (2005) reports that 40% of price changes arise due to sales in a dataset of forty-eight products sold in retail stores during 1997-1999. Hosken and Reiffen (2004) find that 60% of the price decreases in their sample of twenty products sold in 30 locations during 1988-1997 are followed by a price increase in the following month. Several hypothesis have been advanced to explain this pattern of retail price variation, stressing informational frictions on consumer's side of the market (Varian 1980), demand uncertainty (Lazear 1986), or thick-market explanations (Warren and Barsky, 1995), to name a few. Instead of incorporating one (or more) of these explanations into the model economy, I follow Golosov and Lucas (2004) and filter out temporary price sales from the price series. This decision leads me to overestimate the importance of nominal rigidities, as I artificially increase the duration of price spells I ask the model to match, but my goal is to compare the performance of two competing sticky price models, rather than compare the models' performance to the data.

Although I eventually time-aggregate the weekly data into monthly observations, I first filter out sales using the original weekly data[6]. I eliminate sales according to the following algorithm[7]. For any price decrease, I check whether this price change is reversed in one of the four weeks following the original price cut. This definition eliminates both V-shaped price changes (price decreases immediately followed by price increases), but also gradual price decreases, provided these are eventually followed by a price increase after at most four weeks following the first price cut. If a sale is deemed to have taken place, I replace the "sales" price with the price in effect in the period

[6]The alternative choice (of time-aggregating the data first and then eliminating sales) can produce spurious price changes if stores periodically put their prices on sale, at regular intervals.

[7]Dominick's dataset includes a "sales" variable that we could in principle use to eliminate temporary markdowns from our price series. This variable is however coded inconsistently and leaves out many temporary price cuts. I therefore choose to eliminate sales manually.

immediately before the sale. Figure A1 (left panel) illustrates how the algorithm works. The thin line in the figure is the original price series, while the thick line is the "regular" price. Note for example that the first "sale" was implemented gradually, with the original price decrease followed by no price change in the first week of the "sale", an additional price decrease in the second week, and finally a price increase four weeks after the original sale. The drawback of this definition is that it does not eliminate all temporary markdowns in case a price cut is gradually reversed. For example, the final price cut illustrated in the figure was followed by two consecutive price increases, and I have artificially introduced a new "sale" using the algorithm discussed above. To address this problem I repeat the algorithm above an additional three times, in order to eliminate sales that have been gradually implemented. The right panel of Figure 1 illustrates the resulting series of "regular prices" following the last iteration. Note that a single "regular" price change remains after all gradual price reversals are taken into account. Although this remaining price decrease is also reversed, it does not constitute a sale according to our definition because it lasts more than four weeks.

Note a final issue that will play an important role in our discussion of the size and frequency of price changes. By eliminating temporary price cuts, I have introduced an artificial small price change in the regular price series. Any time temporary price reductions are not completely reversed, or followed by price changes larger than the original price cut, setting the regular price equal to the price prior to the sale will artificially introduce a number of small price changes that are otherwise absent in the actual price data. Given that a key statistic in the micro-price data I use to calibrate the menu costs model is the fraction of small price changes in the data, I ignore artificially generated price changes arising due to the filtering of sales and only work with those changes in the regular price that have actually been observed in the original data.

Finally, given that most quantitative studies of sticky price models calibrates them to the

monthly or quarterly frequency, I assume a period of one month in the model presented in text. I therefore time-aggregate weekly observations into monthly data, by constructing the monthly series using price data collected in the first week of the month.

Technical Appendix

The typical approach used in solving state-dependent pricing or inventory models is the simulation technique suggested by Krusell and Smith (1997) and applied by Willis (2002) and Khan and Thomas (2004) to models with non-convexities. I depart slightly from the standard method and use a solution technique free of simulations, one that draws heavily on collocation, a residual-based functional approximation method discussed at length in Miranda and Fackler (2002). A simulation-free solution technique used to solve models with heterogeneous agents was originally suggested by DenHaan (1997) in the context of an uninsurable idiosyncratic risks model.

Recall that the firm's problem is to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{U_c(c_t,n_t)}{U_c(c_0,n_0)} \Pi_t(\tilde{p}_t,)$$

where (for simplicity, I discuss the problem of a single-product firm)

$$\Pi_t(\tilde{p}_t) = \left(\frac{\tilde{p}_t}{\tilde{P}_t}\right)^{-\theta} \left(\frac{\tilde{p}_t}{\tilde{P}_t} - \frac{\tilde{w}_t}{\phi_t \tilde{P}_t}\right) c_t - \xi \frac{\tilde{w}_t}{\tilde{P}_t} \mathcal{I}\left(\tilde{p}_t \neq \frac{\tilde{p}_{t-1}}{g_t}\right) .$$

In equilibrium, $\tilde{w}_t$ is constant at the steady-state level because of the preference structure assumed, and $\tilde{P}_t c_t = 1$.The unknown aggregate functions are $c(g, \boldsymbol{\mu})$, aggregate consumption as a function of the growth rate of the money supply, g and $\boldsymbol{\mu}$, the joint distribution of last period's firm prices and current technology, as well as Γ, the law of motion of $\boldsymbol{\mu}$.

Letting $\Xi = \{g, \boldsymbol{\mu}\}$ denote the aggregate state of the world, one can rewrite the firm's problem recursively as:

$$V^{adj}(\tilde{p}_{-1}, \phi; \Xi) = \max_{\tilde{p}} U_c \left(\Pi(\tilde{p}) - \xi \frac{\tilde{w}}{\tilde{P}}\right) + \beta E V(\tilde{p}, \phi'; \Xi') \tag{V1}$$

$$V^{nadj}(\tilde{p}_{-1}, \phi; \Xi) = U_c \Pi\left(\tilde{p}_{-1}\right) + \beta E V\left(\frac{\tilde{p}_{-1}}{g}, \phi'; \Xi'\right), \tag{V2}$$

where $V = \max\left\{V^{adj}, V^{nadj}\right\}$ is the firm's value, V^{adj} and V^{nadj} is the value of adjustment and not adjustment, respectively and $\Pi(\tilde{p}) = \left(\frac{\tilde{p}}{\tilde{P}}\right)^{-\theta}\left(\frac{\tilde{p}}{\tilde{P}} - \frac{\tilde{w}}{\phi\tilde{P}}\right)c$. $\boldsymbol{\mu}$ evolves according to $\boldsymbol{\mu}' = \Gamma(g, \boldsymbol{\mu})$. The unknowns in this problem are $V^{adj}()$, $V^{nadj}()$, as well as $c()$ and $\Gamma()$. Following Krusell and Smith (1997), I approximate $\boldsymbol{\mu}$ with one moment. In particular, I have found that $\hat{\mu}_t = \int \tilde{p}_{t-1}(z)\phi_t(z)dz$ yields a large degree of accuracy. In the multi-product case $\boldsymbol{\mu}$ is the joint distribution of the two past nominal prices of the firm. I approximate this distribution once again with its first moment: $\hat{\mu}_t = \frac{1}{2}\int \sum_{i=1}^{2} \tilde{p}^i_{t-1}(z)\phi^i_t(z)dz$

Given initial guesses for $c()$ and $\Gamma()$, I solve the functional equations in (V1-V2) using collocation. Specifically, I approximate each of the two value functions using a linear combination of N Chebyshev polynomials. To solve for the $2N$ unknown coefficients, I require that (V1) and (V2) hold at $2N$ nodes in the state space. This condition yields $2N$ equations I use to solve for the unknown coefficients. I solve the firm's maximization problem in (V1) using a Newton-type routine and evaluate the expectations on the RHS of the Bellman equation by discretizing the distribution of shocks and integrating using Gaussian quadrature.

To solve for the aggregate functions $c()$ and $\Gamma()$, I replace them with a linear combination of Chebyshev polynomials and solve for an equilibrium at each node used to discretize the state-space. For each aggregate node $(g_i, \hat{\mu}_i)$, I solve the firm's problem and recompute aggregate variables c_i and $\hat{\mu}'_i$. To calculate these objects, I need to integrate individual firms' decision rules. Given that I only use one moment of the joint distribution of idiosyncratic states, I assume away all variability in $p_{-1}\phi$ (and also that a firm's past prices are independent of each other). I discretize the cross-sectional distribution of ϕ and calculate, for each mass point in this distribution, the associated p_{-1} consistent with the assumption that $p_{-1}\phi$ is degenerate at $\hat{\mu}_i$ and the law of motion for ϕ_t. For each node $\{p_j, \phi_j\}$, I solve firm decision rules and integrate using Gaussian quadrature. Given aggregate quantities c_i and μ_i, I retrieve a new set of coefficients that characterize aggregate functions. For example,

letting c be an $M \times 1$ vector of aggregate consumption that satisfy the equilibrium conditions at each node used to discretize the aggregate state-space, Φ be a $M \times K$ matrix of K Chebyshev polynomials evaluated at the M nodes, I find the K unknown coefficients γ_c by solving $\Phi\gamma_c = C$. This set of coefficients for all aggregate variables is used to re-solve the firm's problem, obtain a set of new aggregate variables at each node and calculate a new set of $\gamma = [\gamma_c, \gamma_\Gamma]$. I stop when the difference between two successive sets of γ is sufficiently low: $||\gamma^{k+1} - \gamma^k||_\infty < 10^{-5}$.

To evaluate the accuracy of this solution method, I plot, in Figure A2, a time-series of aggregate consumption predicted by the approximant $\hat{c}(\mu, \Gamma)$ for a simulation of stochastic forcing processes, as well as the actual aggregate consumption calculated by integrating firm decision rules. I do so for the Benchmark model, the state-dependent setup. Note that the two series are close to each other: the variability of the actual consumption series explained by the approximant is 94%, suggesting that the additional, higher-order moments that we assume away explain little of the fluctuations of aggregate variables in simulations of the model economy. Aggregate functions are even more accurate than the ones illustrated in the figure in the Calvo and single-product state-dependent models.

In Figure A2 I ask how accurate are the solutions to the firm's problem. I plot the left and right-hand side of the Bellman equation in V2, holding constant all other state-variables but $\tilde{p}_{-1}$, at a large number of nodes (larger than that used to pin down the coefficient on the basis functions). Note that the two value functions (predicted and actual), are close to each other. The difference in the two (the residuals) are, as the right panel of Figure A3 indicates, small in absolute value (less than 5×10^{-3}) and oscillate around zero.

Figure A1: Eliminating Sales

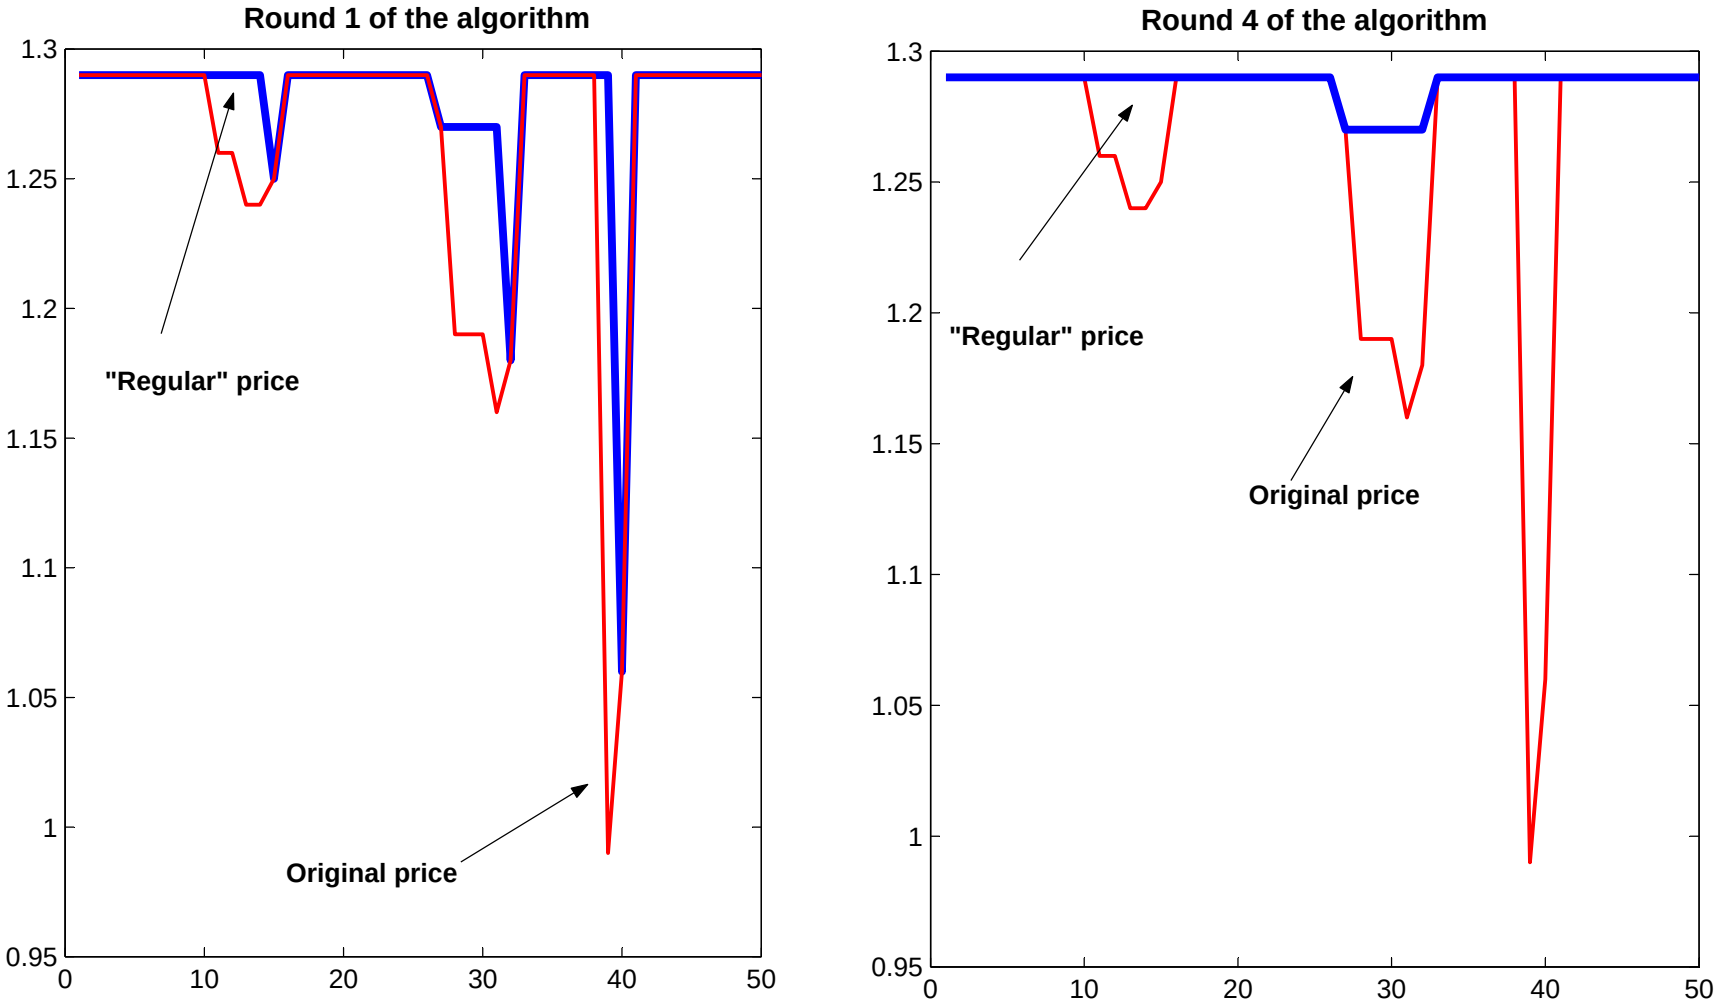


Figure A2: Accuracy of approximants

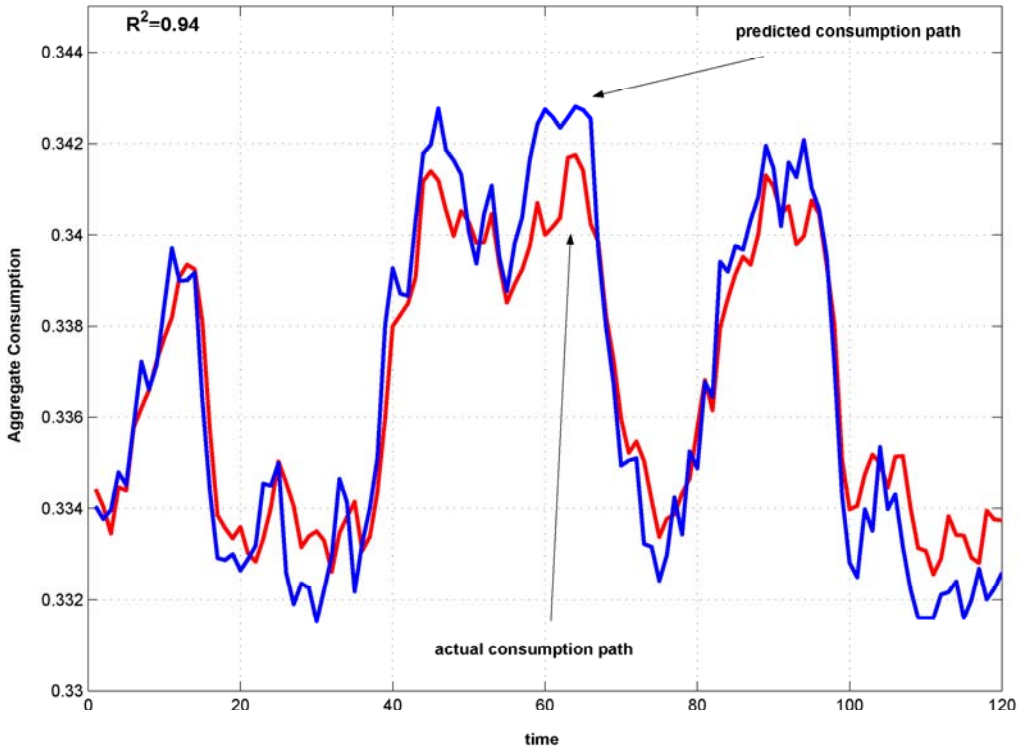


Figure A3: Accuracy of the solution of the firm's problem

