

THE QUARTERLY JOURNAL OF ECONOMICS

Vol. CXXIII

August 2008

Issue 3

STATE-DEPENDENT OR TIME-DEPENDENT PRICING: DOES IT MATTER FOR RECENT U.S. INFLATION?*

PETER J. KLENOW AND OLEKSIY KRYVTSOV

In the 1988–2004 microdata collected by the U.S. Bureau of Labor Statistics for the Consumer Price Index, price changes are frequent (every 4–7 months, depending on the treatment of sale prices) and large in absolute value (on the order of 10%). The size and timing of price changes vary considerably for a given item, but the size and probability of a price change are unrelated to the time since the last price change. Movements in aggregate inflation reflect movements in the size of price changes rather than the fraction of items changing price, because of offsetting movements in the fraction of price increases and decreases. Neither leading time-dependent models (Taylor or Calvo) nor first-generation state-dependent models match all of these facts. Some second-generation state-dependent models, however, appear broadly consistent with the empirical patterns.

I. INTRODUCTION

In time-dependent sticky price models, the timing of individual price changes is exogenous. A firm set its price every n th period (Taylor 1980) or randomly (Calvo 1983). The Taylor and Calvo models feature exogenous staggering of price changes across firms in the economy, and therefore a fixed fraction of firms adjusting their prices each period. Moreover, there is no selection as to who changes their price in a given period.

*This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. John Greenlees, David Johnson, Walter Lane, Rob McClelland, and Teague Ruder provided invaluable assistance and guidance in using BLS data. The views expressed here are those of the authors and do not necessarily reflect the views of the Bank of Canada or the BLS. Rodrigo Barros, Adam Cagliarini, Ryoji Hiraguchi, Romans Pancs, and Benjamin Malin provided superlative research assistance. We thank numerous seminar participants, an editor, and three referees for helpful comments. Send comments and questions to Pete@Klenow.net and okryvtsov@bankofcanada.ca.

© 2008 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

The Quarterly Journal of Economics, August 2008

In state-dependent sticky price models, firms choose when to change prices subject to “menu costs.” Price changes may be bunched or staggered, depending on the importance of common vs. idiosyncratic shocks and other factors. Those changing prices are those with the most to gain from doing so.

Monetary shocks typically have longer lasting real output effects in time-dependent pricing (TDP) models than in state-dependent pricing (SDP) models.¹ Related, prices respond more rapidly to monetary impulses in many more SDP models than TDP models. In SDP models, a positive monetary shock boosts the fraction of firms changing prices and/or the average size of those price changes. In Dotsey, King, and Wolman (1999, hereafter DKW) it is predominantly the fraction that responds, whereas in Golosov and Lucas (2007, hereafter GL) it is almost wholly the average size of changes (more increases and fewer decreases).

Because their positive and normative implications can differ so much, it is important to empirically distinguish between TDP and SDP models. To this end, we deploy the microdata underlying the Consumer Price Index (CPI) compiled by the BLS. The data set consists of monthly retail prices of individual goods and services at specific outlets (excluding shelter) from January 1988 through January 2005. We document the following properties of this data set.

Prices change frequently for the median category—every four months if one includes sale prices, every seven months if one excludes sale prices. Price changes are usually big in absolute terms (averaging around 10%), although a large subset are much smaller (5% or less). Even for individual items, price durations and absolute price changes vary considerably over time. For the typical item, plots of hazards vs. age and size vs. age are pretty flat. The variance of aggregate inflation can be attributed almost entirely to the size of price changes (the intensive margin) rather than the fraction of items changing price (the extensive margin). Underneath the comparatively calm overall fraction, the fraction of price increases swells and the fraction of price decreases subsides when inflation rises.

None of the leading TDP or SDP models we examine can explain all of these empirical regularities. The Taylor model collides

1. See Caplin and Spulber (1987), Caplin and Leahy (1991, 1997), Caballero and Engel (1993, 2006), Dotsey, King, and Wolman (1999, 2006), Gertler and Leahy (2006), and Golosov and Lucas (2007), among others.

with the variable length in price spells for individual items. The Taylor and Calvo models predict bigger absolute price changes for older prices, when no such pattern exists in the data. They also have trouble generating enough movement in the fraction of price increases vs. decreases with inflation. The DKW model produces no large price changes, and predicts far too big a role for the extensive margin in inflation movements. The GL model does not generate enough small price changes. As we discuss briefly in the conclusion, second-generation SDP models such as DKW (2006), Midrigan (2006), and Gertler and Leahy (2006) fare better vis à vis our evidence.

In Section II, we describe the U.S. CPI data set. In Section III, we present a series of stylized facts about this data set. In Section IV, we compare the predictions of leading TDP and SDP models to these facts. We conclude in Section V.

II. BLS MICRODATA SET ON CONSUMER PRICES

To construct the non-shelter portion of the CPI, the BLS surveys the prices of about 85,000 items a month in its *Commodities and Services Survey*.² Individual prices are collected by 400 or so BLS employees visiting 20,000 retail outlets a month, mainly across 45 large urban areas. The outlets consist of grocery stores, department stores, auto dealerships, hospitals, etc. The survey covers all goods and services other than shelter, or about 70% of the CPI based on BLS consumer expenditure weights.

The BLS selects outlets and items based on household point-of-purchase surveys, which furnish data on where consumers purchase commodities and services. The price collectors have detailed checklists describing each item to be priced—its outlet and unique identifying characteristics. They price each item for up to five years, after which the item is rotated out of the sample.

The CPI Research Database (CPI-RDB), maintained by the BLS Division of Price and Index Number Research, contains all prices in the Commodities and Services Survey since January 1988. We use the CPI-RDB through January 2005, and will refer to this as “1988–2004.”

2. The BLS conducts a separate survey of landlords and homeowners for the shelter portion of the CPI. The sources for this section are the BLS *Handbook of Methods* (U.S. Department of Labor, 1997, Chap. 17) and unpublished documentation for the CPI-RDB (to be described shortly).

II.A. Frequency of BLS Pricing

The BLS collects consumer prices *monthly* for food and fuel items in all areas. The BLS also collects prices monthly for all items in the three largest metropolitan areas (New York, Los Angeles, and Chicago). The BLS collects prices for items in other categories and other urban areas only *bimonthly*.³ About 70% of observations in our pooled sample over the 1988–2004 period are monthly price quotes, and the remaining 30% are bimonthly. We concentrate our analysis on the top three areas. Because our focus is on the endogenous timing of price changes, we prefer a longer sample of 205 monthly observations to a pair of broader samples with bimonthly observations.

II.B. Temporary Price Discounts (“Sales”)

According to the BLS, a “sale” price is (a) temporarily lower than the “regular” price, (b) available to all consumers, and (c) usually identified by a sign or statement on the price tag. Roughly 11% of quotes in the sample are sale prices. Sales are especially frequent for food items, where they comprise 15% of all quotes (vs. 8% of nonfood price quotes). Chevalier, Kashyap, and Rossi (2003) also observe frequent sales in their analysis of scanner data from grocery stores. They report that sales often generate V shapes, as the price goes down and then returns to the regular (presale) level in the next period. In the BLS data, about 60% of sales exhibit this pattern. A number of papers, such as Golosov and Lucas (2007), Midrigan (2007), and Nakamura and Steinsson (2008), exclude sale prices from their analysis. To facilitate comparison, we report many statistics for both “posted prices” (based on both regular and sale prices) and “regular prices” (based only on regular prices).

II.C. Forced Item Substitutions

Forced item substitutions occur when an item in the sample has been discontinued from its outlet and the price collector identifies a similar replacement item in the outlet to price going forward. This often takes the form of a product upgrade or model changeover. The monthly rate of forced item substitutions hovers around 3% in the sample. Essentially all item substitutions involve price changes—these are new items at new prices. Shapiro and Wilcox (1996) and Moulton and Moses (1997) point out that

3. In Philadelphia and San Francisco, the BLS priced items monthly through 1997 and bimonthly thereafter.

substitutions account for about half of the inflation rate (a third if one excludes apparel). As stressed by Bils and Klenow (2004), these price changes materially boost the overall frequency of price changes. We include these substitution-related price changes in our statistics, with the sole exception being when we show implied durations with and without them to highlight their importance.

II.D. Out-of-Season Items

Although the Commodities and Services Survey attempts to price 85,000 or so items per month, the BLS succeeds in collecting around 74,000 price quotes in the typical month. The 11,000 unavailable quotes per month consist of out-of-season items, temporary stockouts, and permanently discontinued items. The BLS categorizes about 5% of the items they attempt to price as out of season in the average month. Not surprisingly, the out-of-season fraction is particularly high for clothes. We compare prices across missing observations because doing so matters for the overall frequency of price changes and for getting close to the mean inflation rate in the official CPI.

II.E. Stockouts

Even when in season, almost 7% of items are temporarily unavailable in a typical month. The BLS reserves this classification for items that are temporarily out of stock from the outlet or at outlets that are temporarily closed. As we do for out-of-season items, we incorporate prices straddling stockouts.

II.F. Outliers

Although the BLS requires its price collectors to explain large price changes to limit measurement errors, some price changes in the data set appear implausibly large. We exclude price changes that exceed a factor of 10. Such price jumps constitute less than one-tenth of 1% of all price changes.

III. STICKY PRICE FACTS

With our sampling decisions, we arrive at a subsample of the CPI-RDB consisting of between 13,000 and 14,000 price quotes a month in the top three areas (New York, Los Angeles, and Chicago). We call the longitudinal string of prices for a particular item at a particular outlet a “quote-line.” This includes multiple versions of an item due to forced item substitutions. In our subsample, the mean (median) time between the first and last

observation in a quote-line is 43 (45) months. The interquartile range is 24 to 60 months. Recall that items are rotated every five years or more frequently. A quote-line may be shorter than 60 months because it is rotated more frequently (e.g., as with consumer electronics recently), because its outlet closed or discontinued its category, or because it straddles the beginning or end of our sample. The mean (median) number of price quotes in a quote-line is 37 (34) months. The interquartile range is 18 to 54 months. The number of prices falls short of a quote-line's life span because of temporary stockouts and items being out of season (especially important for clothing).

We now document some stylized facts about these prices. In the subsequent section we compare these facts to predictions of some leading TDP and SDP models.

III.A. Price Changes Are Frequent

Using summary statistics from the CPI microdata, Bils and Klenow (2004, hereafter BK) estimated the median duration of prices across categories to be 4.3 months for posted prices, and perhaps 5.5 months for regular prices. These durations were considerably shorter than the prior consensus of roughly a year between price changes. Nakamura and Steinsson (2008, hereafter NS), using the CPI-RDB, reported notably higher median durations of 8–11 months for regular prices. We now estimate our own price durations, with the upshot being that even regular prices appear to change every seven months or so.

We follow both BK and NS in first estimating the frequency of price changes for 300 or so categories of consumption known as Entry Level Items (ELIs). Within these categories are prices for particular items (quote-lines). The BLS has provided us with unpublished ELI weights for each year from 1988 to 1995 and 1999 to 2004, based on Consumer Expenditure Surveys in each of those years. We normalize the non-shelter portion of the weights to sum to 1 in each year. We set the 1996 and 1997 ELI weights to the 1995 weights, and the 1998 weights to their 1999 level. We set the 2005 weights to their 2004 level (for January 2005). The BLS changed ELI definitions in 1998, and so interpolating was not an option. The CPI-RDB contains weights for each price within an ELI. We allocate each ELI's weight to individual prices in each year in proportion to these item weights to arrive at weights ω_{it} that sum to 1 across items (i 's) and months in each year. Thus our statistics weight all years equally.

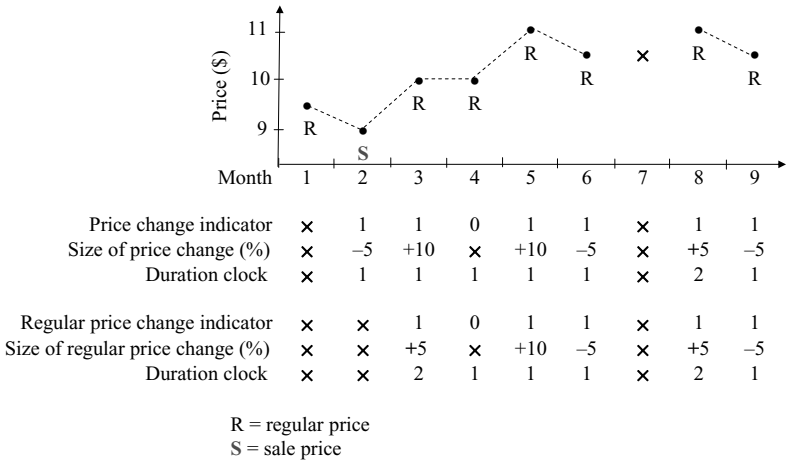


FIGURE I

A Hypothetical Sequence of Prices for an Individual Item in the CPI

Notes: An \times indicates that the variable is unavailable for that month. In the text, for item i in month t , the price change indicator is denoted I_{it} , the size of the price change dp_{it} , and the duration clock (gap in months between observations) γ_{it} . Equation (1) shows how these variables are used to estimate price change frequencies, for example.

As discussed, even if one includes sale prices, there are many gaps in the quote-lines. If we exclude sale prices, this creates additional gaps. To estimate the frequency of price changes using all available observations, we apply a maximum likelihood estimator that takes into account the time between each observed price. Let $\{p_{it}\}$ denote the set of log price observations in quote-line i . Denote γ_{it} as the gap in months between the price at t and the previous observation. Let I_{it} be a price-change indicator; $I_{it} = 1$ if $p_{it} \neq p_{it-\gamma_{it}}$, 0 otherwise. See Figure I, and note that these variables are defined differently depending on whether we are using all (posted) prices or just regular prices. The monthly Poisson rate of price change for a quote-line in category s is λ_s , which is assumed to be common across quote-lines within the category (ELI) and across time. Conditional on the sequence of prices observed, the likelihood of observing the I_{it} values in a category is then

$$L_S = \prod_{i \in S} \prod_t (1 - e^{-\lambda_s \gamma_{it}})^{I_{it}} (e^{-\lambda_s \gamma_{it}})^{1 - I_{it}}$$

The inside product is across prices within a quote-line and the outside product is for different quote-lines in category s . The

log-likelihood function is then

$$\ln L_S = \sum_{i \in S} \sum_t [I_{it} \ln(1 - e^{-\lambda_S \gamma_{it}}) - (1 - I_{it}) \lambda_S \gamma_{it}].$$

The first order condition $\partial \ln L_S / \partial \lambda_S = 0$ implicitly defines the maximum likelihood estimate $\hat{\lambda}_S$:

$$(1) \quad \sum_{i \in S} \sum_t \frac{I_{it} \gamma_{it}}{e^{\hat{\lambda}_S \gamma_{it}} - 1} = \sum_{i \in S} \sum_t (1 - I_{it}) \gamma_{it}.$$

For each category we solve this nonlinear equation for $\hat{\lambda}_S$, applying weights ω_{it} to both sides.

This method assumes two things, whether the gap between price quotes is a single month or longer: if the price does not differ from the previous price then it did not change at any point in between, and if the price does differ from the previous price it may have changed at any point in between (possibly more than once). Given that the price often returns to a previous level because of temporary “sales,” this method will miss some price changes.

Allowing for price changes during gaps in item availability deserves some elaboration. We may want to know how often prices change conditional on availability. When items are not available, one should arguably stop the “duration clock” altogether. Think of the extreme case of items out of season for half the year or more. As we shall see, stopping the duration clock (setting $\gamma_{it} = 1$ regardless of the gap) implies modestly more price flexibility than our latent price approach. An advantage of a ticking duration clock, however, is that the shocks motivating price changes presumably accumulate over time. Thus our baseline estimates treat unavailable items as if they were available but not sampled by the BLS. A fuller treatment would try to model the gaps to determine the best treatment.

Once we have frequencies for each category (ELI), we take weighted means or medians across categories. Because the BLS redefined ELIs in 1998, we estimate the means and medians separately for 1988–1997 and 1998–2004, then weight them by 10/17 and 7/17. Table I provides monthly frequencies and implied durations (inverses of the frequencies). For posted prices, the mean (median) frequency across categories is 36% (27%), and the mean (median) implied duration is 6.8 (3.7) months. We do not report standard errors; sampling error is small given our sample of roughly 3 million prices. The median duration of 3.7 months is

TABLE I
MONTHLY FREQUENCIES AND IMPLIED DURATIONS

Statistic	Posted prices	Regular prices	Like prices
Mean frequency	0.362	0.299	0.313
Median frequency	0.273	0.139	0.168
Implied mean duration	6.8	8.6	7.8
Implied median duration	3.7	7.2	6.0

Notes. Samples run from January 1988 through January 2005 and include data from top three urban areas. All data are from the CPI-RDB. The durations are inverses of the monthly frequencies. Regular prices are posted prices excluding sale prices. Like prices compare a regular (sale) price only to the previous regular (sale) price. Means and medians use weights based on the BLS consumer expenditure surveys and unpublished BLS point-of-purchase surveys.

actually lower than that in BK. The BK sample was 1995–1997 for all areas and weighted items equally, whereas ours is 1988–2004 for the top three areas with item weights.

For regular prices, Table I reports a mean (median) frequency across categories of 30% (14%), and a mean (median) implied duration of 8.6 (7.2) months. As NS emphasize, sales are a bigger fraction of price changes for items near the median (13 of the 27 percentage points) than for the mean item (6 of the 30 percentage points). Hence BK's uniform adjustment for sales, which yielded a median duration of 5.5 months, understated the importance of sales for the median item. We find that taking out sales moves the median duration from 3.7 months up to 7.2 months. The implied mean duration is less affected, rising from 6.8 to 8.6 months.

Focusing on regular prices begs the question of whether one should exclude sale prices for macro purposes. This is not obvious. First, sales may have macro content. Items may sell at bigger discounts when excess inventory builds up or when inflation has been low. Klenow and Willis (forthcoming 2007) find that sale-related price changes are at least as sensitive to inflation as are regular price changes. Second, excluding price changes because they are not macro-related (e.g., because of idiosyncratic shocks or price discrimination) risks confusing sticky nominal prices with sticky information as in Mankiw and Reis (2002). Third, it is awkward to apply a menu cost model to data in which some price changes have been filtered out. Kehoe and Midrigan (2007) eschew this approach, instead modeling one-period price discounts as incurring a lower menu cost. They find that sale-related price changes can contribute to macro price flexibility, but not as much as regular price changes. They focus on grocery stores, in which V-shaped

sales (which end in the price returning to the previous regular price) are most common. In apparel, the other category with endemic sales, sales typically end in regular price changes related to product turnover.

In a variant of the sticky information hypothesis, one might imagine that firms have a sticky plan that involves alternating between a sticky regular price and a sticky sale price. The BLS may sometimes catch the regular price and sometimes the sale price. Because comparing regular and sale prices to each other can obscure such rigidity, in Table I we present a third classification of price changes: those involving “like” prices. We compare a regular price only to the previous regular price, and a sale price only to the previous sale price. We apply maximum likelihood estimation as if the regular and sale prices come from separate quote-lines within the ELI. If sale prices are very sticky in nominal terms or relative to regular prices, then “like” prices will be at least as sticky as regular prices. This is not what we see in Table I. The median frequency of changes in like prices is 17% (vs. 14% for regular prices), and the implied median duration of like prices is 6.0 months (vs. 7.2 months for regular prices).⁴ Evidently a sale price is more likely to differ from the previous sale price than a regular price is to differ from the previous regular price. And many changes in sale price are not interrupted by regular prices (think of repeated markdowns for clearance items). Our evidence says sales are not at a fixed price or discount from the regular price; they have a life of their own.

We summarize by characterizing the median duration of prices as ranging from 4 to 7 months, depending on the treatment of sales (included, excluded, or like).⁵ The *mean* durations are higher at 7 to 8.5 months or so. Which statistic is more meaningful depends on the model. According to Carvalho (2006), stronger strategic complementarities make the mean and higher moments more relevant, and the median less so.

For the median duration, our upper bound of 7 months is outside the 8- to 11-month range reported in NS for regular prices. Table II attempts to shed light on this discrepancy. The bottom line is that NS exclude many prices that we incorporate. The first column of Table II presents medians to facilitate comparison. The first row contains our 3.7-month duration for posted prices. The

4. We do not report statistics on sale prices alone because most categories have no sales at all in the BLS data.

5. If we stop the “duration clock” when prices are not available, the range is 3–6 months rather than 4–7 months.

TABLE II
DURATIONS UNDER VARIOUS EXCLUSIONS

Case	Implied median duration	Implied mean duration
Posted prices	3.7	6.8
No V shapes	5.3	7.7
Like prices	6.0	7.8
Regular prices	7.2	8.6
No substitutions	8.7	10.4
Adjacent prices	9.3	12.0
1998–2004 only	10.6	13.4

Notes. Samples run from January 1988 through January 2005 and include data from the top three urban areas. All data are from the CPI-RDB. The durations are inverses of the monthly frequencies. Means and medians use weights based on the BLS consumer expenditure surveys and unpublished BLS point-of-purchase surveys.

No V shapes: lower middle prices are replaced with identical neighbors.

Like prices: regular (sale) price is compared only to the previous regular (sale) price.

Regular prices: posted prices excluding sale prices.

No substitutions: only regular prices in between item substitutions are compared.

Adjacent prices: only consecutive monthly regular prices in between substitutions.

1998–2004 only: sample is confined to 1998–2004 (otherwise adjacent prices).

next row shows what happens if, every time a middle price is lower than its identical neighbors (a V shape), the middle price is replaced with its neighbors. We do this regardless of the “sale” or “regular” labels. Filling in V shapes raises the median duration from 3.7 months to 5.3 months. Thus purely temporary price discounts do not explain our low durations compared to the pre-BK literature. The next row of Table II repeats the median duration of 6.0 months for like prices. Like prices allow for the possibility that sale prices are sticky even if they do not end with a return to the previous regular price. The subsequent row repeats our regular price duration of 7.2 months. These initial rows drive home the point that our 7-month upper bound involves aggressive filtering of all sales.

The Table II row labeled “No substitutions” compares only regular prices in between item substitutions. The median duration climbs to 8.7 months without substitution-related price changes. About 80% of the time the substituted item’s price differs from the previous item’s price. To us, these are new prices pure and simple.⁶

6. We may be missing some price changes tied to product turnover. Bills (2005) argues that the BLS overstates quality-adjusted price increases at substitutions. A subset of the 20% of substitutions showing no price change may therefore involve true price decreases. Also, we ignore the 0.5% frequency of item exit without replacement. Exiting items probably give way to new items at new prices (e.g., Broda and Weinstein [2007]).

BK included them and stressed their importance to overall price flexibility. Now, it is possible the new prices on new items do not incorporate macro information. But excluding them convolutes sticky nominal prices and sticky information. Moreover, Klenow and Willis (forthcoming 2007) find that substitution-related price changes are very much related to inflation.

NS exclude substitution-related price changes in their 8- to 11-month range. This leads them to report, for example, that regular prices change every 27 months in apparel, where the average item lasts only about 10 months. Excluding substitutions is tantamount to assuming that the frequency of price changes is the same at the point of model turnover as within the product cycle, despite the fact that a new price is arrived at 80% of the time. NS note that substitutions are often seasonal, and therefore arguably time-dependent rather than state-dependent. Golosov and Lucas (2007) and many others have argued that SDP involves greater price flexibility, but this does not seem to be grounds for excluding substitutions any more than excluding all time-dependent price changes.^{7,8}

The next row in Table II reports a median duration of 9.3 months for “adjacent prices.” This statistic compares only consecutive monthly regular prices between substitutions. Their slightly higher duration says price changes are more frequent after items return to stock, come back into season, and return from sales. NS exclude nonadjacent price comparisons in the high end of their 8- to 11-month range. We see no justification for ignoring these data, given that it matters.

The final row of Table II presents a median duration of 10.6 months when the sample is confined to 1998–2004 (and to adjacent regular prices between substitutions). At the high end of their 8- to 11-month range, NS drop the first 10 years of the CPI-RDB. We think it is most informative to include all of the years in the CPI-RDB.

To recap, the median duration rises from 7.2 months for 1988–2004 regular prices to 10.6 months for 1998–2004 adjacent regular

7. For seasonal apparel, one might imagine a saw-toothed pattern, with the price being the same in each January, each February, etc. For the 25 seasonal apparel ELIs, however, prices 12 months apart differ 76% of the time. This implies a monthly hazard of 12% and a duration of 8 months. Comparing only regular prices within product cycle, NS report a duration of 28 months for seasonal apparel items.

8. Regular product turnover may affect the frequency of price changes before and afterward. Analogously, see Hobijn, Ravenna, and Tambalotti (2006) for analysis of inflation around the euro changeover.

prices between substitutions. The mean duration goes from 8.6 months to 13.4 months. One can obtain median durations above 4–7 months only by excluding many price changes beyond all sale-related price changes.

III.B. Price Changes Are Large in Average Absolute Value

How large are price changes? We again start by calculating means within categories (ELIs) for 1988–1997 and 1998–2004 separately. Let

$$dp_{it} \equiv p_{it} - p_{it-\gamma_{it}}$$

be the difference in the log price from the previous observation γ_{it} months ago. Recall from Figure I that these dp_{it} (and γ_{it} gaps) are defined differently for posted vs. regular prices.

Within each ELI-year, we weight price changes (observations with $I_{it} = 1$ and therefore $dp_{it} \neq 0$) in proportion to the price weights in the CPI-RDB. These $\hat{\omega}_{it}$ weights differ from the ω_{it} weights used for frequency estimation, because the $\hat{\omega}_{it}$ are nonzero only for price changes (they are 0 when $I_{it} = 0$). The $\hat{\omega}_{it}$ weights add up to 1 in each year, thus weighting years equally. The ELI mean for a given sample of years is then

$$(2) \quad \overline{|dp|}_s \equiv \frac{\sum_{i \in s} \sum_t \hat{\omega}_{it} |dp_{it}|}{\sum_{i \in s} \sum_t \hat{\omega}_{it}}.$$

In the numerator the inside summation is across price changes within a quote-line and the outside summation is across quote-lines (i 's) in an ELI. The denominator is the sum of the weights for all price changes. We calculate weighted means and medians across ELIs for 1988–1997 and 1998–2004, then take a weighted average of them (with weights 10/17 and 7/17) to arrive at statistics for 1988–2004.

Table III shows that absolute changes in posted prices average 14% with a median of 11.5%. Regular price changes (both the current and previous prices are regular prices) are smaller but still large, with a mean of 11% and median of 10%. Sale-related price changes (either the current or the previous price is a sale price) are very large, with a mean of 25% and median of 23%. “Like” price changes are a weighted average of the regular and sale-related price changes, and they average 12% with a median

TABLE III
ABSOLUTE SIZE OF PRICE CHANGES

Sample	Mean	Median
Posted prices	0.140	0.115
Regular prices	0.113	0.097
Sale-related prices	0.251	0.233
Like prices	0.118	0.110

Notes. Samples run from January 1988 through January 2005 and include data from top three urban areas. All data are from the CPI-RDB. Means and medians use weights based on the BLS consumer expenditure surveys and unpublished BLS point-of-purchase surveys. Regular prices are posted prices excluding sale prices. Sale-related prices: either the current or previous price is a sale price. Like prices: regular (sale) price compared only to the previous regular (sale) price.

of 11%. We summarize Table III as saying that absolute price changes are large, averaging around 10%.⁹

Because prices change every 4 months or so and inflation averages around 0.2% per month in our sample, the average price change is about 0.8%. The combination of small average price changes and large average *absolute* price changes suggests substantial idiosyncratic shocks to marginal cost and/or desired markups. Golosov and Lucas (2007) use this evidence as motivation for an SDP model with idiosyncratic productivity shocks.¹⁰ Midrigan (2006) and Dotsey, King, and Wolman (2006) do so as well. Klenow and Willis (2007) argue that these big price changes undermine the case for strategic complementarities that make it costly for firms to change their relative prices.

Burstein and Hellwig (2006) consider idiosyncratic demand shocks alongside productivity shocks. They use scanner data from Dominick's grocery stores, and report smaller price increases than price decreases. In the CPI-RDB, price increases are 2 to 3 percentage points smaller than price decreases. For regular prices, increases average 10.6% whereas decreases average 13.3%. For sale-related price changes, the difference is only a percentage point.

III.C. Yet Many Price Changes Are Small

Figure II provides a histogram of regular price changes. On the horizontal axis are bins of price changes; on the vertical axis is the percentage of the 1988–2004 weight represented by the

9. Absolute price changes are 1 to 2 percentage points smaller if we exclude the large changes at substitutions. Note that, throughout the paper, we refer to log-changes as “percent” changes.

10. In Appendix I, we describe five moments that we calculated from the CPI-RDB expressly for Golosov and Lucas.

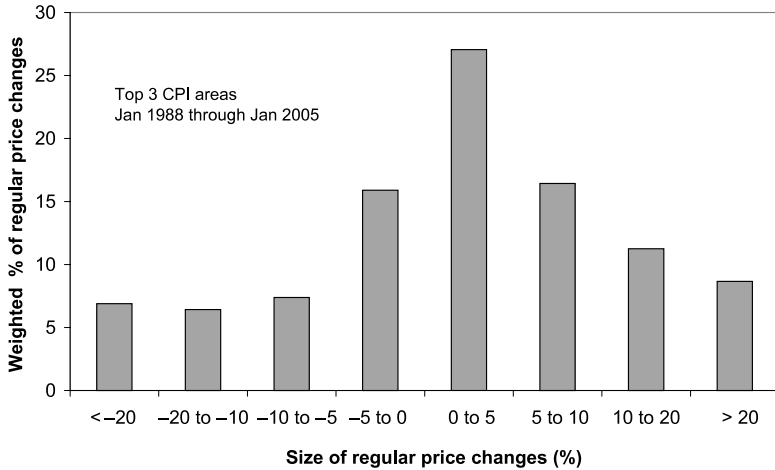


FIGURE II
Weighted Distribution of Regular Price Changes

TABLE IV
FRACTION OF PRICE CHANGES BELOW SIZE THRESHOLDS

Sample	$ dp_{it} < 5\%$	$ dp_{it} < 2.5\%$	$ dp_{it} < 1\%$
Posted prices	0.398	0.234	0.113
Regular prices	0.443	0.254	0.121
Like prices	0.427	0.248	0.120

Notes. Samples run from January 1988 through January 2005 and include data from the top three urban areas. All data are from the CPI-RDB. Regular prices are posted prices excluding sale prices. Like prices compare a regular (sale) price only to the previous regular (sale) price. Entries are weighted mean fractions of price changes that are smaller than 5%, 2.5%, or 1% in absolute value; $|dp_{it}|$ = the absolute size of price changes. Weights are based on the BLS consumer expenditure surveys and unpublished BLS point-of-purchase surveys.

price changes in each bin. The weights on each individual price change are the $\hat{\omega}_{it}$ weights used in (2). The histogram shows that both large and small price changes are common, as are both price increases and price decreases. The corresponding histogram for posted prices is similar, though with thicker tails.

Table IV provides greater detail on small price changes. Around 44% of regular price changes are smaller than 5% in absolute value, 25% are smaller than 2.5%, and 12% are smaller than 1%. Midrigan (2006) finds many small price changes in the Dominick's data as well. He points out that such small price changes are hard to reconcile with the large menu costs needed to rationalize the large average price changes. The GL model, for

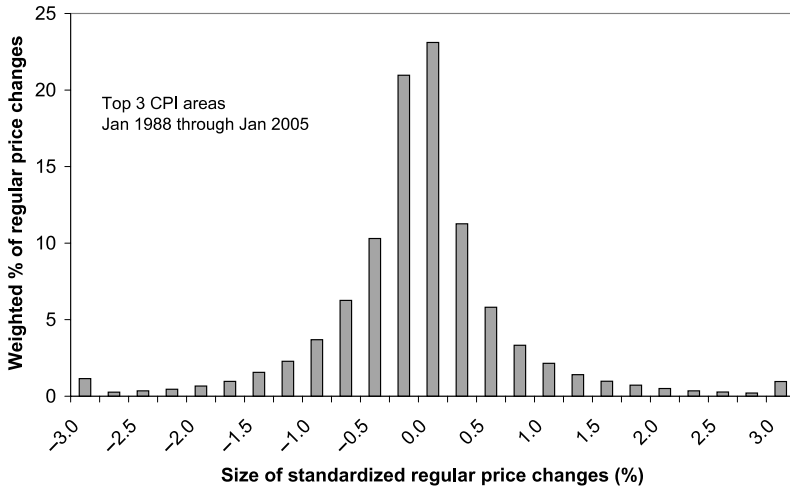


FIGURE III
Weighted Distribution of Standardized Regular Price Changes

example, does not generate many small price changes. Midrigan introduces economies of scope in changing prices (across items for a given manufacturer or a given aisle for a retailer) to explain the coexistence of many small price changes alongside large average changes.

The existence of many large and small price changes might also be rationalized by a wide range of menu costs across items and/or time, as in the DKW model; that is, the small (large) price changes may reflect small (big) menu costs. We visit this possibility in our model section, below. For now, Figure III plots the distribution of *standardized* price changes:

$$(3) \quad z_{it} \equiv \frac{dp_{it} - \mu_{dps}}{\sigma_{dps}}$$

This *z*-statistic is the deviation of a price change from the weighted average price change in the same ELI $s(\mu_{dps})$, divided by the weighted standard deviation of price changes in $s(\sigma_{dps})$. Figure III shows that price changes are often small compared to their average size within the same ELI.

III.D. Price Durations Are Variable, Even for a Given Item

What is the distribution of price spells? We could infer this from our frequency estimates, perhaps taking into account heterogeneity across categories or quote-lines. But rather than impose

constant flat hazards, we now present direct evidence on spell lengths. Several issues arise: gaps in quote-lines, censoring at the ends of quote-lines, and weighting of spells. For our frequency estimates, we allowed for the possibility that prices changed between observations, but that leaves the spell length indeterminate. Thus we go to the “sticky” extreme here and carry forward the last observed price to fill in gaps. If a price has changed, we treat it as having done so at the end of any gap. We did this for stockouts, out-of-season items, and even sale prices.

We dropped the censored spells at the beginnings and ends of quote-lines because we cannot determine their length. We could try to adjust for this censoring, but are worried that doing so would make the distribution of spells look artificially uniform. Also, it would entail making assumptions about hazard profiles. By focusing on completed spells, this distribution could be biased toward short spells. But there are biases in each direction: long spells are more likely to be censored, but short spells are less likely to be in the sample to begin with.¹¹ Also, our filling in of gaps should bias the durations upward.

We are careful to weight each spell equally within quote-lines, but *not* across quote-lines. We take the observation weights ω_{it} (which sum to 1 across all prices in each month) and divide them by 205, the number of months. The resulting weights sum to 1 across all items and months. We sum these $\omega_{it}/205$ weights across time for a given quote-line. We then divide this cumulative quote-line weight equally across each of the quote-line’s completed price spells to arrive at spell weights ϖ_{it} (which are nonzero only if $I_{it} = 1$). If a quote-line has a constant flat hazard, these spell weights correctly yield an average duration equal to the inverse of that hazard. This method does *not* weight quote-lines with frequent price changes more than quote-lines with rare price changes. If a quote-line has only a few price changes, all of its cumulative weight will be spread among these few price changes. Put differently, the weight on each price change will be small for items with many price changes and large for items with few price changes. This keeps our durations from tilting toward items with frequent, short price spells.

Figure IV presents the (weighted) distribution of times between *regular* price changes. Almost half of all weighted spells (47%) are within one quarter. As shown in the right portion of

11. Gabriel and Reiff (2007) discuss the effect of censoring on the estimated distribution of price spells.

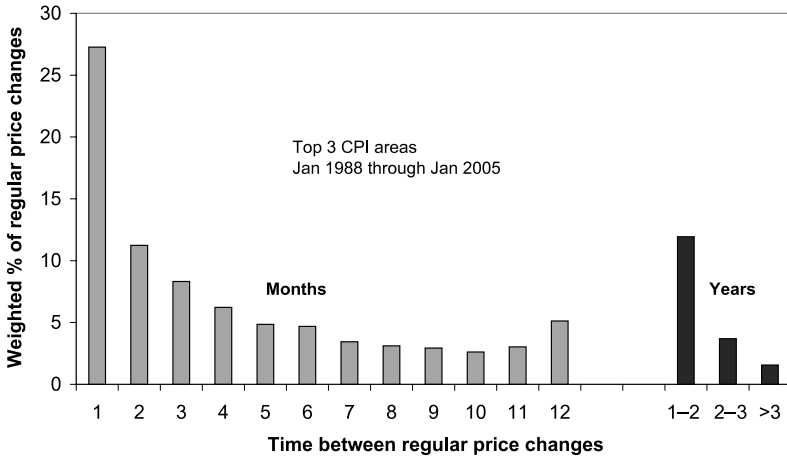


FIGURE IV
Weighted Distribution of Times between Regular Price Changes

the figure, almost one-fifth of spells (17%) last a year or longer. Only 5% of price spells last precisely a year; it will be useful to keep this modest figure in mind when we plot hazard rates below. Spells are shorter for posted prices (which include sale prices), with 57% lasting a quarter or less and 13% lasting more than a year. For both regular and posted prices, these median “direct” durations are considerably shorter than those implied by our frequency estimates in Table I. Thus the direct durations may be biased downward because of censoring of longer spells. We wish to use Figure IV only to establish that durations are highly variable, however, a point that can be made even with a censored sample.

Figure IV makes it clear that there is no dominant spell length. But perhaps different durations predominate for different categories or items, and the pooled figure conceals this pattern. We take the spells and weights underlying Figure IV, subtract the weighted mean duration within ELIs (or quote-lines), sum the weighted squared deviations, and take the square root to arrive at a weighted standard deviation of spells. Table V shows that spells vary considerably even within ELIs and quote-lines. Within ELIs, the standard deviation is 6.2 months for posted prices (compared to the mean duration of 6.8 months) and 7.0 months for regular prices (compared to the mean of 8.6 months). Even within quote-lines, the standard deviation is 4.3 months for posted prices and 5.2 for regular prices.

TABLE V
STANDARD DEVIATION OF COMPLETED SPELLS (IN MONTHS)

Sample	Within ELIs	Within quote-lines
Posted prices	6.2	4.3
Regular prices	7.0	5.2

Notes. Samples run from January 1988 through January 2005 and include data from top three urban areas. All data are from the CPI-RDB. Regular prices are posted prices excluding sale prices. Weights are based on the BLS consumer expenditure surveys and unpublished BLS point-of-purchase surveys.

III.E. Hazard Rates Are Flat (for a Given Item)

Rather than assume a flat profile of price-change hazards, we now estimate hazard rates for each age of prices. We calculate the weighted average frequency of price changes conditional on reaching each age. We use the observation weights ω_{it} (which sum to 1 across all prices in each month) and divide them by 203, the number of months for which we have prices with determinate ages.¹² The resulting weights $\tilde{\omega}_{it}$ sum to 1 across items and months. With this pooled sample, we calculate the weighted average of the price-change indicator for each age of prices τ using the indicator $I\{\tau_{it} = \tau\}$, which is 1 when $\tau_{it} = \tau$ and 0 otherwise:

$$\{\bar{I}_{it} | \tau_{it} = \tau\} \equiv \frac{\sum_i \sum_t \tilde{\omega}_{it} I_{it} I\{\tau_{it} = \tau\}}{\sum_i \sum_t \tilde{\omega}_{it} I\{\tau_{it} = \tau\}}.$$

Figure V provides the monthly hazard rates at each age for regular prices. This pooled hazard is sharply declining in the first several months. In fact, it is monotonically declining except for blips every six months. The increase at one year is more of a spike, with the hazard rate more than doubling to around 22%. Recall from Figure IV, however, that only about 5% of spells last precisely twelve months. So this locally important spike is not so important globally. The hazard profile for posted prices has a similar shape, but is at a higher level.

The declining pooled hazards could simply reflect a mix of heterogeneous flat hazards, that is, survivor bias. Consistent with this interpretation, the hazards flatten out considerably when we split quote-lines into those above vs. below the median frequency

12. We do not know the age of prices in the first, censored spell of each quote-line, and so we do not use these prices in estimating hazard rates. For items that changed price in February 1988, the prices in March 1988 do go into our hazard rates. Hence 203 months (March 1988 through January 2005) are represented.

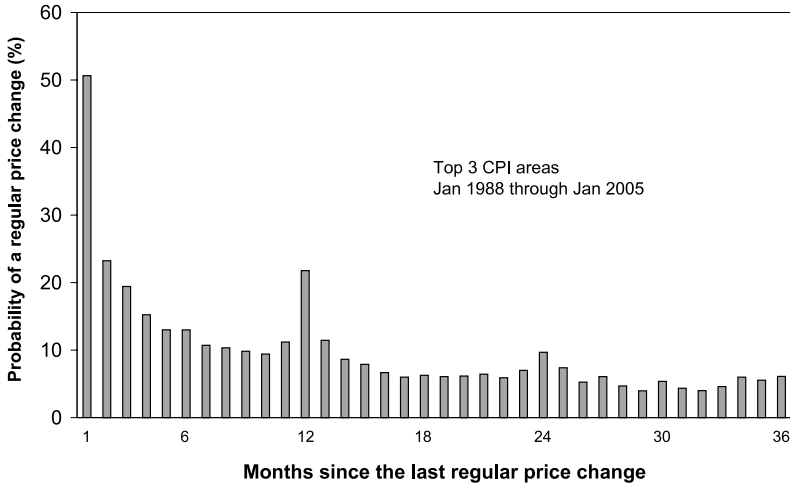


FIGURE V
Weighted Hazard Rates for Regular Prices

and look increasingly flat when we split further into quartiles, deciles, etc. We illustrate this by allowing for ten different fixed effects in the pooled hazards. We calculate the unconditional hazard rate for each quote-line μ_{I_i} , sort the quote-lines into deciles based on μ_{I_i} , and then calculate the unconditional hazard rate for each decile μ_{I_d} , $d = 1, \dots, 10$. Finally, we calculate the weighted mean hazard at each age relative to the decile fixed effect:

$$\left\{ \overline{I_{it}/\mu_{I_d}} | \tau_{it} = \tau \right\} \equiv \frac{\sum_i \sum_t \tilde{\omega}_{it} [I_{it}/\mu_{I_d}] I \{ \tau_{it} = \tau \}}{\sum_i \sum_t \tilde{\omega}_{it} I \{ \tau_{it} = \tau \}}.$$

Figure VI provides the hazard rates (relative to fixed effects) for regular prices in each of the first twelve months, where the vast majority of the duration mass lies. After a modest drop from month 1 to month 2 (the equivalent of going from 24% to 21% in frequency), hazard rates are essentially flat until spiking at twelve months. Allowing for additional fixed effects does not alter the picture much, and the pattern is similar for posted prices. Taking into account survivor bias, therefore, virtually eliminates the downward slope to the hazard profile. This empirical finding stands in contrast to that of Nakamura and Steinsson (2008), who emphasize downward-sloping hazards in the CPI-RDB.




FIGURE VI
Weighted Hazards for Regular Prices vs. Decile Fixed Effects

III.F. The Size of Price Changes Is Unrelated to the Time since the Previous Change (for a Given Item)

We now explore how the absolute size of price changes relates to the duration of a completed spell. As with Figure IV and Table V on variable durations, we use spell weights ϖ_{it} that divide a quote-line’s cumulative weight equally among all of its completed spells. Pooling all items, the weighted average size of price changes for each age of prices τ is

$$\{|\overline{dp_{it}}| \mid \tau_{it} = \tau\} \equiv \frac{\sum_i \sum_t \varpi_{it} |dp_{it}| I\{\tau_{it} = \tau\}}{\sum_i \sum_t \varpi_{it} I\{\tau_{it} = \tau\}}.$$

Figure VII plots average size vs. age for regular prices. Price changes start out small ($\sim 9\%$ for one-month spells) and grow larger ($\sim 14\%$ for ten-month spells). For posted prices the picture is much flatter, with price changes hovering around 14% for nine of the first ten months.

The rising size vs. age pattern for regular prices could reflect heterogeneity in size that is correlated with the heterogeneity in hazards. We check for this by allowing for ten different fixed effects. We calculate the unconditional size of price changes for each quote-line $\mu_{|dp|i}$, sort the quote-lines into deciles based on $\mu_{|dp|i}$, and then calculate the unconditional size of price changes for each




FIGURE VII
Weighted Average Size of Regular Price Changes by Age

decile $\mu_{|dp|d}$, $d = 1, \dots, 10$. We then calculate the weighted mean size of price changes relative to the decile fixed effect at each age:

$$\overline{\{|dp_{it}|/\mu_{|dp|d}|\tau_{it} = \tau\}} \equiv \frac{\sum_i \sum_t \varpi_{it} [|dp_{it}|/\mu_{|dp|d}] I\{\tau_{it} = \tau\}}{\sum_i \sum_t \varpi_{it} I\{\tau_{it} = \tau\}}.$$

Figure VIII confirms that heterogeneity contributed to the rising size vs. age in Figure VII. For regular prices, size edges up modestly (the equivalent of about 1 percentage point), compared to increasing 5 percentage points without decile fixed effects. Evidently, quote-lines with more frequent price changes have smaller absolute price changes. Similarly, there is little relation between size and age for posted prices. The picture changes little when we add more fixed effects.

III.G. *The Intensive Margin Dominates the Variance of Inflation*

Inflation is the product of the fraction of items with price changes (the extensive margin, EM) and the average size of those price changes (the intensive margin, IM):

$$\pi_t \triangleq \sum_i \omega_{it}(p_{it} - p_{it-1}) \triangleq \underbrace{\sum_i \sum_t \omega_{it} I_{it}}_{fr_t} \cdot \underbrace{\frac{\sum_i \sum_t \omega_{it}(p_{it} - p_{it-1})}{\sum_i \sum_t \omega_{it} I_{it}}}_{dp_t}.$$

(4)




FIGURE VIII
Size of Regular Price Changes by Age vs. Decile Fixed Effects

As shown, fr_t is the “fraction” (more accurately, the CPI weight) of items changing price in month t , and dp_t is the weighted-average magnitude of price changes occurring in month t .

Table VI contains summary statistics for regular prices in the top three CPI areas from February 1988 through January 2005. The monthly inflation averages 0.27%, or 3.3% annualized.¹³ The fraction of items changing price averages 27% a month, with a 3% standard deviation. The coefficient of variation is much higher for the average size of price changes, which has a standard deviation of 1.2% and a mean of 1%.

Figure IX displays twelve-month moving averages for π_t , fr_t , and dp_t . EM (fr_t) is relatively stable and not so obviously correlated with inflation (correlation .25 in Table VI). IM (dp_t) is more volatile and comoves almost perfectly with inflation (correlation .99).

How important are EM and IM, respectively, for the variance of inflation? We decompose the variance of inflation over time into (terms involving) the variance of the average price change, the variance of the fraction changing price, and their covariance. We

13. Our regular price inflation rate is higher than the 3.0% headline CPI inflation rate because it excludes clearance prices. For posted prices, our annual inflation rate is 2.5%, which is lower than 3.0% because of our use of geometric weighting in (4) and our exclusion of shelter. The BLS began using geometric weights in 1999 and only within categories. Also, our series is more volatile because it is only for the top three areas.

TABLE VI
TIME SERIES MOMENTS FOR REGULAR PRICES

Variable	Mean	Std dev	Correlation with π	Regression on π	
	(%)	(%)		Coef.	Std. err.
π	0.27	0.36			
fr	26.6	3.2	.25	2.38	0.66
dp	0.98	1.19	.99	3.55	0.04
$fr+$	15.0	2.6	.69	5.48	0.41
$fr-$	11.5	2.5	-.41	-3.10	0.48
$dp+$	8.87	1.10	.19	0.64	0.23
$dp-$	9.37	1.64	-.19	-0.96	0.34
pos	1.33	0.27	.74	0.59	0.04
neg	1.06	0.23	-.60	-0.41	0.04

Notes. Samples run from January 1988 through January 2005 and include data from top three urban areas. All data are from the CPI-RDB. The entries are means, standard deviations, and cross-correlations across time of the monthly values of each variable. The last two columns are the OLS regression coefficient and standard error from regressing the variable in the first column on inflation. The monthly values of the variables are across-item weighted means from month $t - 1$ to t of

- π_t = inflation.
- fr_t = the fraction of items with changing prices.
- dp_t = the size of price changes (note not the absolute value).
- $fr+_t$ = the fraction of items with rising prices.
- $fr-_t$ = the fraction of items with falling prices.
- $dp+_t$ = the size of price increases.
- $dp-_t$ = the absolute size of price decreases.
- pos $_t$ = $fr+_t \cdot dp+_t$.
- neg $_t$ = $fr-_t \cdot dp-_t$.

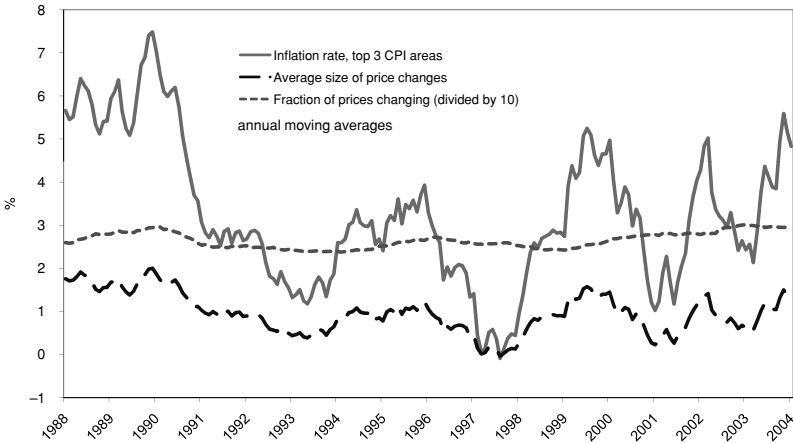


FIGURE IX
Extensive and Intensive Margins of Inflation

TABLE VII
VARIANCE DECOMPOSITIONS

Sample	IM vs. EM (%)		POS vs. NEG (%)	
	IM term	EM terms	POS terms	NEG terms
Posted prices	94	6	59	41
Regular prices	91	9	59	41

Notes. All samples go from January 1988 through January 2005. See equation (5) in the text for the definition of the IM term and the EM terms, and equation (9) for the definition of the POS terms and NEG terms. All data are from the CPI-RDB.

find it useful to take the variance of a first-order Taylor series expansion of $\pi_t \triangleq fr_t \cdot dp_t$ around the sample means \bar{fr} and \bar{dp} to arrive at this exact variance decomposition:

$$(5) \quad \text{var}(\pi_t) = \underbrace{\text{var}(dp_t) \cdot \bar{fr}^2}_{\text{IM term}} + \underbrace{\text{var}(fr_t) \cdot \bar{dp}^2 + 2 \cdot \bar{fr} \cdot \bar{dp} \cdot \text{cov}(fr_t, dp_t) + O_t}_{\text{EM terms}}$$

The higher order terms (O_t) and the covariance term are small both in the data and in the SDP models we simulate. The quantitatively important terms are the variance terms.

Table VII reports variance decompositions based on posted prices and regular prices. The IM term accounts for all of inflation's variance in staggered TDP models, and anywhere from a little to most of it in SDP models.¹⁴ The EM terms involve changes in the extensive margin, which only contribute in SDP models. The results in Table VII are clear. The intensive margin accounts for 94% of inflation's variance for posted prices and 91% for regular prices. We repeated the variance decomposition for a variety of samples: excluding item substitutions, excluding food and energy items, and including urban areas with bimonthly price quotes. Across the samples, the IM term accounted for between 86% and 113% of the variance of inflation. The fraction of items changing

14. From (5), one might think small average price changes would guarantee a small extensive margin, but this was decidedly not the case when we simulated the DKW SDP model in Klenow and Kryvtsov (2005). Even when matching the empirical average price change (dp) and monthly inflation rate, the EM accounted for 70% of the variance of inflation in the DKW model.

price—a key endogenous variable in *some* SDP models—are a relatively unimportant source of fluctuations in inflation.^{15,16}

What is behind high volatility of the average size of price changes (dp_t) and its strong comovement with inflation? It is useful to further decompose inflation and its components into terms due to price increases and decreases. Let I_{it}^+ be a price-increase indicator that is 1 if $p_{it} > p_{it-\gamma_{it}}$ and 0 otherwise. Similarly, let I_{it}^- be a price-decrease indicator that is 1 if $p_{it} < p_{it-\gamma_{it}}$ and 0 otherwise. The weighted fraction of items changing price (fr_t) is the sum of the weighted fractions of items with rising (fr_t^+) and falling (fr_t^-) prices:

$$(6) \quad fr_t \triangleq \sum_i \omega_{it} I_{it} \triangleq \underbrace{\sum_i \sum_t \omega_{it} I_{it}^+}_{fr_t^+} + \underbrace{\sum_i \sum_t \omega_{it} I_{it}^-}_{fr_t^-}.$$

Related objects are the average sizes of price increases and decreases:

$$(7) \quad dp_t = \frac{fr_t^+}{fr_t} \cdot \underbrace{\frac{\sum_i \sum_t \omega_{it} I_{it}^+ (p_{it} - p_{it-1})}{\sum_i \sum_t \omega_{it} I_{it}^+}}_{dp_t^+} - \frac{fr_t^-}{fr_t} \cdot \underbrace{\frac{\sum_i \sum_t \omega_{it} I_{it}^- |p_{it} - p_{it-1}|}{\sum_i \sum_t \omega_{it} I_{it}^-}}_{dp_t^-}.$$

Using (6) and (7), inflation is the sum of terms involving price increases and decreases:

$$(8) \quad \pi_t = \underbrace{fr_t^+ \cdot dp_t^+}_{pos_t} - \underbrace{fr_t^- \cdot dp_t^-}_{neg_t}.$$

Table VI reports that price decreases are less frequent (at 11.5%) than increases (15%).¹⁷ The sizes of price increases and

15. Cecchetti (1985) points out that, if price changes are perfectly staggered, then one can estimate the frequency of price changes from across-good variation of inflation relative to mean inflation. He successfully applies this to magazine prices, meaning their changes are indeed largely staggered—a precursor to our finding. In Cecchetti (1987), he applies the same methodology to estimate the frequency of wage changes.

16. Gagnon (2007) finds a larger role for the extensive margin in Mexico's more volatile 1994–2004 inflation.

17. Thus 57% of price changes are price increases. Nakamura and Steinsson (2008) report that 67% are increases. The difference is that we are reporting the mean fraction, whereas they are highlighting the median fraction.

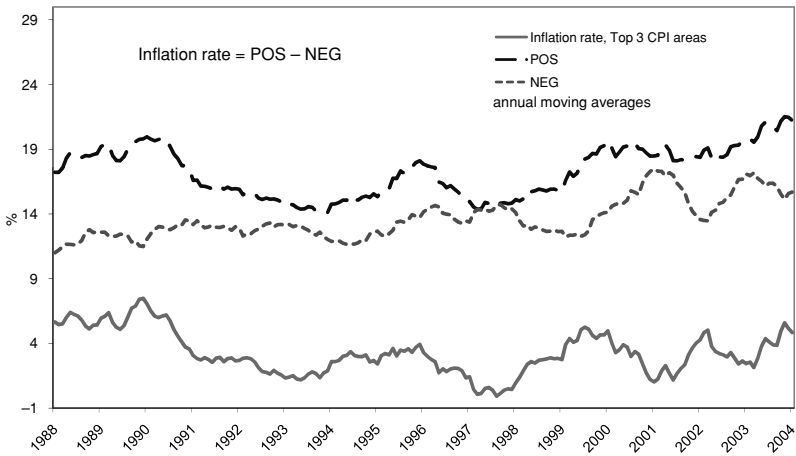


FIGURE X
Inflation due to Price Increases and Decreases (Regular Prices)

decreases are modestly correlated with inflation (0.2 and -0.2 , respectively); the corresponding fractions correlate more strongly with inflation (0.7 and -0.4 , respectively). We also regress these components on inflation. Table VI shows that a 1 percentage point increase in inflation is associated with a 5.5 (-3.1) percentage point change in the fraction of price increases (decreases), and a 0.6 (-1) percentage point change in the size of price increases (decreases). These coefficients are all statistically significant, but fr_t^+ and fr_t^- are more closely tied to inflation than are dp_t^+ or dp_t^- (both statistically and relative to their means).¹⁸

Nakamura and Steinsson (2008) also relate the frequency and size of price increases and decreases to inflation. Like us, they find the fraction of price increases correlates most with inflation. Unlike us, they find no significant relation between inflation and the incidence of price decreases or either size variable. In their interpretation, price increases (pos_t), not price decreases (neg_t), drive inflation movements.

Figure X plots twelve-month moving averages of our π_t , (pos_t), and (neg_t). The (pos_t) term correlates modestly more with inflation (0.7) than does the (neg_t) term (-0.6). When we regress the terms on inflation, the coefficients are 0.59 and -0.41 (see Table VI).

18. All regression results in this subsection are robust to including seasonal dummies, including a time trend, or using the official CPI rather than our top three area inflation rate.

These coefficients mirror a particular variance decomposition:

$$\text{var}(\pi_t) = \underbrace{\text{var}(\text{pos}_t) - \text{cov}(\text{pos}_t, \text{neg}_t)}_{\text{POS term}} + \underbrace{\text{var}(\text{neg}_t) - \text{cov}(\text{pos}_t, \text{neg}_t)}_{\text{NEG term}}.$$

(9)

Here the covariance term is split evenly between the POS and NEG terms.¹⁹ Table VII shows that the POS term indeed accounts for 59% of inflation's variance, and the NEG term 41%. If we include month dummies (to control for seasonality), the inflation decomposition becomes 50% POS, 50% NEG. Hence our series suggest a very important role for both price increases and price decreases in inflation movements.

We know of two reasons our "POS, NEG" results differ from those of NS. First, in their figures they plot cross-category *median* fractions and sizes against inflation, whereas our statistics are for cross-category *means*. We think means are at least as useful as medians because of identities tying the means to actual inflation. Second, their regression results, although for means, are for twelve-month moving averages and for 1988–1997 and 1998–2005 separately. Thus they have only ten independent observations on inflation in one subsample, and eight in another. Our results in Tables VI and VII are for the 204 months across the entire 1988–2004 sample. When we aggregate up to 68 quarters, the roughly equal importance and significance of POS and NEG remain very much intact. It is only when we aggregate up to annual data and split the sample that the components start plummeting in significance. Compared to monthly data for the whole sample, our standard errors are four to ten times higher with annual data and a split sample.

IV. STICKY PRICE MODELS VS. FACTS

In this section we simulate some sticky price models to compare their predictions to the facts presented in the previous section. In particular, we study the patterns of price adjustment in partial equilibrium versions of the Golosov and Lucas (2007), Dotsey, King and Wolman (1999), Calvo (1983), and Taylor (1980) models. For each model we simulate firm pricing decisions in 67

19. The covariance is only -0.055 relative to inflation's variance, and so this matters little.

TABLE VIII
STICKY PRICE ECONOMIES VS. BLS FACTS

Facts	GL	DKW	Calvo	Taylor
Frequent price changes	Yes	Yes	Yes	Yes
Large average absolute price changes	Yes	No	Yes	Yes
Many small price changes	No	Yes	Yes	Yes
Variable price durations	Yes	Yes	Yes	No
Flat hazard rates	Yes	No	Yes	No
Size of price changes does not increase with duration	Yes	No	No	No
Intensive margin dominates the variance of inflation	Yes	No	Yes	Yes
Frequency of price increases and decreases moves with inflation	Yes	Yes	No	No

Notes. The table summarizes the success of sticky price models in matching the empirical evidence from the BLS microdata on retail prices in the top three metropolitan areas of the United States. More details on the empirical evidence can be found in Section III. Each model's performance in replicating the facts is discussed in Section IV.

sectors calibrated to evidence for 67 BLS Expenditure Classes.²⁰ We then calculate the model statistics corresponding to the facts documented in Section III. Readers familiar with these models will want to skip to the scorecard in Table VIII. The bottom line will be that none of these models matches all of the stylized facts we laid out.

IV.A. *GL Model*

A monopolistically competitive firm in a sector is endowed with a constant-returns-to-scale technology that converts l units of labor input into $z l$ units of output.²¹ Here z represents a firm's productivity level. We assume $\ln(z)$ follows an AR(1) process:

$$\ln(z') = \rho_z \ln(z) + \varepsilon_z,$$

where ε_z is a mean zero, normally distributed error with standard deviation σ_z . Let p' (and p) denote the price of the good sold in the current (and previous) period relative to the geometric average price within a sector in that period. Demand for the firm's output in the current period is $p'^{-\theta}$, where the elasticity of goods demand, θ , is a constant, and the aggregate output is fixed at 1.

20. BLS expenditure classes are a higher level of aggregation than strata (~200) or ELIs (~300). The parameter values for these simulations are available in an appendix on our websites.

21. Here and throughout, the sector indices are implicit.

Let φ denote a fixed cost of implementing a price change (“menu cost”) expressed in units of labor. The firm begins the current period with relative price p , inherited from the previous period. After realizing its current productivity level z , the firm chooses whether to adjust its price. If it changes its price, the firm pays the fixed labor cost at wage w , $w\varphi$, and chooses the new relative price p' . Otherwise, the firm keeps its previous price, so that the new relative price is decreased by the rate of inflation π . Formally, the problem of the firm can be written as follows:

$$(10) \quad V_A(p, z) = \max_{p'} p'^{-\theta} [p' - w/z] + \beta \int V(p', z') Q(z, dz'),$$

$$(11) \quad V_N(p, z) = (p/\pi)^{-\theta} [p/\pi - w/z] + \beta \int V(p/\pi, z') Q(z, dz'),$$

$$V(p, z) = \max [V_A(p, z) - w\varphi, V_N(p, z)],$$

where $V(p, z)$ is the value before the adjustment decision, $V_A(p, z)$ is the value of adjusting one's price, and $V_N(p, z)$ is the value of not adjusting one's price; $Q(z, \cdot)$ denotes the cumulative density function (CDF) of future shocks z' conditional on the current realization z . Appendix II describes the numerical solution method we apply to the firm's problem.

IV.B. The DKW Model

In the original Dotsey, King, and Wolman (1999) model, the fixed cost, φ , is not constant but rather i.i.d. across firms and over time, drawn from a differentiable CDF $G(\varphi)$, where $0 \leq \varphi \leq \varphi_{\max}$. The firm's problem is now written as

$$V_A(p, z) = \max_{p'} p'^{-\theta} [p' - w/z] + \beta \int V(p', z', \varphi) Q(z, dz') G(d\varphi),$$

$$V_N(p, z) = (p/\pi)^{-\theta} [p/\pi - w/z] + \beta \int V(p/\pi, z', \varphi) Q(z, dz') G(d\varphi),$$

$$V(p, z, \varphi) = \max [V_A(p, z) - w\varphi, V_N(p, z)],$$

The i.i.d. assumption on menu costs implies that values after the adjustment decision, $V_A(p, z)$ and $V_N(p, z)$, do not depend on the realization of the current fixed cost. The original DKW model does not have idiosyncratic shocks to the firm's productivity, and so all adjusting firms in a given period choose the same price.

IV.C. *The Calvo Model*

As with the GL model, we allow idiosyncratic productivity shocks here. With this modification, a Calvo firm can be approximated by a DKW firm with a fixed cost CDF putting all weight on $\varphi = 0$ and $\varphi = \varphi_{\max}$. When the realized fixed cost is $\varphi = \varphi_{\max}$ and φ_{\max} is large, the firm will (almost) never change its price. Conversely, a firm will (almost) always adjust its price when the menu cost is zero. Such exogenous TDP contrasts with the SDP in the GL and DKW models, in which a firm's decision to change its price depends on the state (p, z) .

IV.D. *The Taylor Model*

Under Taylor price setting, a firm is allowed to adjust its price every n th period. As in the Calvo model, the price adjustment decision does not depend on the state.

To facilitate the comparison of the sticky price models to the BLS data, we simulate prices set by an optimizing firm for 204 consecutive months. The number of firms simulated corresponds to the number of price quotes in the corresponding BLS expenditure category divided by 204. Several parameters in the firm's problem are calibrated to moments in the regular price data for the corresponding BLS category: the inflation rate in the sector, π , equals the average inflation rate in the BLS category; the menu cost, φ , and two parameters of the firm's productivity process, ρ_z and σ_z , are chosen to fit the frequency of price changes plus the standard deviation and serial correlation of a firm's relative price (conditional on price change) for the expenditure category.²² The remaining parameters are common across simulated sectors. The monthly discount factor $\beta = 0.97^{1/12}$, corresponding to a 3% annual real interest rate. The goods demand elasticity $\theta = 10$, implying a desired markup of 11%. The wage w does not affect price dynamics.

We confront the alternative models with the following stylized facts in the BLS data: (a) price changes are frequent; (b) price changes are typically large; (c) many price changes are small; (d) price spells vary in length; (e) hazard rates are essentially flat; (f) the size of price changes does not increase with duration; and (g) the intensive margin dominates the variance of inflation;

22. For the DKW model we assume a uniform CDF of menu costs and calibrate φ_{\max} . For the Calvo model we set the probability of the menu cost being zero to the frequency of price changes in the BLS category.

the extensive margin contributes little because of offsetting movements in the fraction of price increases vs. decreases.

All of the models can generate the first fact (price changes are frequent) by setting the menu cost low enough. The last fact (the intensive margin drives inflation) cannot be reproduced by models in this section because they abstract from aggregate inflation fluctuations. We cite the results from general equilibrium simulations in Golosov and Lucas (2007) and Klenow and Kryvtsov (2005) to assess models' success in matching the importance of the intensive margin for the variance of inflation.

IV.E. Golosov and Lucas vs. the Facts

The GL model can generate large price changes, and indeed was calibrated to do so using moments we provided from the CPI-RDB. In our simulations of this model the average absolute price change is 14.6%, in the neighborhood of the value for posted prices. Midrigan (2006) stresses, however, that the GL model fails to generate enough small price changes. The fraction of price changes between -5% and $+5\%$ is around 9% in our simulations vs. 40–44% in the BLS data. This model fails to generate enough small price changes because it has a single, large cost of price adjustment. A model with a range of menu costs, including many small ones as in Calvo or DKW, stands a better chance. Rather than the stochastic menu costs of those models, we entertain menu costs specific to each of our 67 sectors. Figure XI plots the distribution of “standardized” price changes as in (3). There remains a “missing middle” of small price changes in the model vs. the data. This is true even if we entertain a much lower elasticity of substitution ($\theta = 2$), as will be helpful for matching some of the facts. We stopped at calibrating 67 sectors to the data, but it is possible that a GL model with more heterogeneity could fit this fact. As mentioned, Midrigan (2006) proposes a different fix, namely economies of scope in changing prices.

Because of the endogenous timing of price adjustment and random shocks to the firm's productivity, the GL model succeeds in predicting variable spell durations. The standard deviation of durations in our simulations is 10.4 months, which is actually larger than the 7.0 months or so within ELIs (see Table V). A model with more sectors would presumably come closer to the data. Still, Figure XII shows that the model matches the empirical distribution of spell durations fairly well with $\theta = 10$.

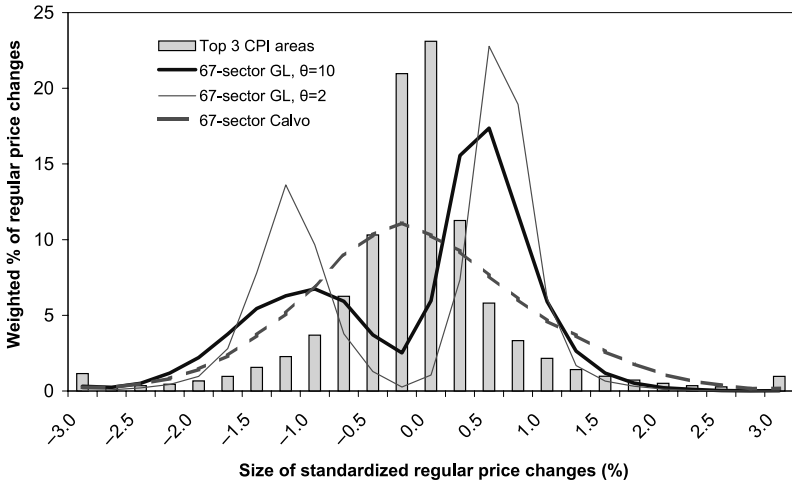


FIGURE XI
Standardized Price Changes: Models vs. Evidence

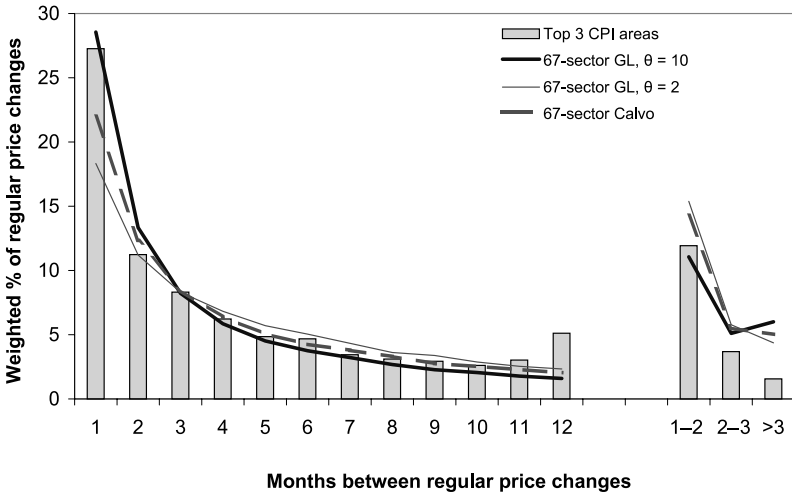


FIGURE XII
Distribution of Price Durations: Models vs. Evidence

Figure XIII compares hazard rate deviations for the GL model to those in the top three CPI areas. The fixed effects in the GL model are at the sector level, whereas in the data they are in deciles of quote-lines. With $\theta = 10$ the GL model predicts *declining* hazards, in contrast to the flat hazards in the data. The

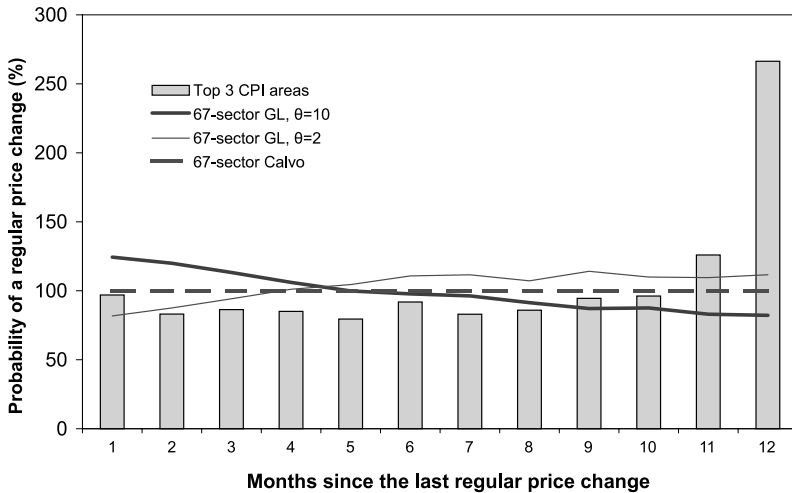


FIGURE XIII
Hazard Rates vs. Decile Fixed Effects: Models vs. Evidence

probability of a price change is about 10 percentage points higher at one month than at one year. The declining hazards may come as a surprise because one expects hazards to rise as persistent productivity shocks accumulate. The reason they are downward sloping is because the width of the inaction region, or “sS band,” is not fixed. When a firm’s idiosyncratic productivity is persistently high, it chooses a low price and sells a high quantity. A high level of profits is at stake in getting the price right, and so the sS band is narrow. At low productivity levels, less is at stake and so the sS band is wider.²³ This lends a form of survivor bias to the hazard profile: young prices are more common when the sS band is narrower and the hazard rate is higher; old prices are more common when the sS band is wider and the hazard rate is lower. Consistent with this explanation, the model hazards slope upward when we lower the elasticity to $\theta = 2$ in Figure XIII.²⁴ At a lower elasticity the scale effect of higher productivity on profits is much weaker. Our model has a deterministic inflation trend, but no aggregate shocks. Such shocks might be more persistent than idiosyncratic shocks, lending some upward slope to the hazard.

23. See Figure 1 in Golosov and Lucas (2007).

24. Also consistent with our interpretation, we find a strong negative relationship in the model between a firm’s relative price and the probability of its price changing when $\theta = 10$, and a much weaker relationship when $\theta = 2$.




FIGURE XIV

Size vs. Age (vs. Decile Fixed Effects): Models vs. Evidence

This effect would probably be modest because idiosyncratic forces drive most price changes.

Figure XIV compares models and data on the size of price changes (relative to fixed effects) against the age of the price. The GL model with $\theta = 10$ predicts bigger absolute price changes for older prices. Price changes are about 5 percentage points larger after one year than after one month. The explanation for this property is again the varying width of the sS bands. Young prices are disproportionately from the narrow portion of the band, where price changes are small; old prices are disproportionately from the wide portion of the band, where price changes are big. When $\theta = 2$, the scale effect on the width of sS bands is much weaker. The size vs. age profile becomes as flat as a pancake in Figure XIV when $\theta = 2$.²⁵

Finally, Golosov and Lucas (2007) report that their general equilibrium rightly implies that the intensive margin drives inflation movements. The intensive margin term in their variance decomposition accounts for around 85% of inflation's variance. Their overall fraction of price changes is stable because prices are changed mostly in response to large idiosyncratic shocks rather

25. We find a strong positive relationship in the model between a firm's relative price and the size of its price changes when $\theta = 10$, and a much weaker relationship when $\theta = 2$.

than the smaller aggregate shocks. In their model, just as in our data, the fraction of price increases is positively correlated with inflation, and the fraction of price decreases is negatively correlated with inflation. Aggregate shocks affect inflation predominantly through this “selection effect.”

The first column of Table VIII summarizes the ability of the GL model to fit the facts from the CPI-RDB. With $\theta = 2$, the model succeeds in matching all of the facts except for the sizable fraction of small price changes in the data.²⁶

IV.F. DKW vs. the Facts

In Klenow and Kryvtsov (2005), we simulate the original general equilibrium version of the Dotsey, King, and Wolman (1999) model. Because of its range of menu costs, the model can produce a lot of small price changes and a wide span of spell durations. The variance of inflation is dominated by the extensive margin in the DKW calibration, however, in stark contrast to the evidence. Also, because the original DKW model does not have idiosyncratic shocks, it is incapable of generating large average absolute price changes. Furthermore, without idiosyncratic shocks, the gap between the firm’s actual and desired prices is increasing in the age of the actual price. This makes for increasing hazard rates and rising absolute size of price changes, contradicting two of our facts.

IV.G. Calvo vs. the Facts

Our Calvo model with large idiosyncratic shocks to a firm’s price has large average absolute price changes. At the same time, its variable menu costs (sometimes zero, sometimes prohibitive) generate a lot of variation in the size of price changes. The mean and standard deviation of the absolute size of price changes are 12.8% and 10.8%, respectively. The former statistic is lower and the latter is higher than in the GL model, because many price changes are small in the Calvo model. The pooled distribution of price changes is close to the bell-shaped empirical distribution. Small price changes occur whenever the previous relative price is near the desired price and the realized menu cost is zero. Thus, the Calvo model fits the size distribution of price changes much better than the GL model. Still, Figure XI shows that the Calvo

26. Instead of a low demand elasticity, one could neutralize the scale effect by making the menu cost proportional to (say) sales rather than being fixed in terms of labor. See, for example, Gertler and Leahy (2006).

model does not have as many small standardized price changes as seen in the data.

Random productivity and menu costs imply variable price spell durations. The standard deviation of durations between price adjustments is 8.7 months, compared to 7.0 months within ELIs. Figure XII illustrates how the model stacks up against the empirical distribution of spell durations. The model falls short by 5 percentage points for one-month durations, and predicts around 7 percentage points too many spells exceeding one year.

Calvo hazard rates for individual sectors are flat by construction. Figure XIII shows this to be an empirical victory. Related, the intensive margin term in the variance decomposition is 100% by definition. This is compared to the 90% or more seen in the data. The frequency of price increases vs. decreases, however, does not move much with inflation in the Calvo model, in contrast to the data.

Like the GL model, the Calvo model predicts bigger absolute size of price changes for older prices. Price changes are about 5 percentage points bigger after one year than after one month (see Figure XIV). Although the pattern is similar to that in the GL model, the explanation is different. As a price ages, shocks (and trend inflation) accumulate, moving the actual price further away from the desired price on average. In the GL model, selection dulls this effect. Either way, the rising size with age pattern is inconsistent with the flat profile in the data.

The third column of Table VIII reiterates that the Calvo model matches all the facts except the flat size vs. age profile and the fraction of price increases vs. decreases moving a lot with inflation.

IV.H. Taylor vs. the Facts

The timing of price adjustment is deterministic in the Taylor model, as opposed to random in the Calvo model. For any sector in the Taylor model, the distribution of price spell durations is degenerate, putting all weight on the single duration of price changes in that sector. This is sharply at odds with the data, in which spell durations are variable even at the quote-line level (standard deviation of at least four months). Related, the Taylor model predicts zero hazard except at a single duration, in stark contrast to the flat empirical hazards. For other facts, the Taylor model does as well as the Calvo model due to the presence of large idiosyncratic shocks to the firm's price (see Table VIII).

We recap this section by describing those features of sticky price models that help them succeed in matching the highlighted aspects of the microdata. Large idiosyncratic disturbances help replicate large absolute price changes. A range of menu costs, arriving irregularly, helps generate many small price changes and variable spell durations. The Calvo model generates flat hazard rates. The GL model with low goods demand elasticity is able to produce flat hazards vs. age and flat size vs. age. Finally, the evidence on the importance of the intensive margin for the variance of inflation is reproduced in all models with large idiosyncratic shocks to the firm's price. But the SDP models do a better job than the TDP models of matching the rise and fall of price increases vs. decreases with inflation.

V. CONCLUSION

Using the CPI Research Database maintained by the BLS, we documented a number of facts about consumer prices from January 1988 through January 2005. In this data set, price changes are frequent and typically (but far from exclusively) large in absolute value. For a given item, price durations vary, with no correlation between the size or probability of price changes and time since the last price change. The variance of inflation stems mostly from the intensive margin (the size of price changes) rather than the extensive margin (the fraction of prices changing). But this dichotomy is due to offsetting movements in the fraction of price increases and price decreases, which correlate strongly with inflation.

We compare these empirical patterns to several prominent TDP and SDP models. The GL model produces too few small price changes. The DKW model does not generate enough large price changes, and gives too much prominence to the extensive margin. The Taylor model falters in predicting a degenerate duration for a given item. The Calvo and Taylor models wrongly predict bigger price changes for older prices.

In contrast to the empirical shortcomings of first-generation SDP models, some second-generation SDP models enjoy greater success. Midrigan (2006) adds economies of scope to the Golosov and Lucas model to fill in the "missing middle" of small price changes. Gertler and Leahy (2006) propose a model with Poisson arrival of idiosyncratic shocks and small menu costs, so that many price changes are small and the size of price changes is unrelated

to the time since the last change. Dotsey, King, and Wolman (2006) add idiosyncratic shocks to their earlier model, thereby generating large price changes and mitigating the importance of the extensive margin.

APPENDIX I: RELATED STATISTICS PROVIDED TO GOLOSOV AND LUCAS

To calibrate their model, Golosov and Lucas (2007) use five statistics that we calculated from the BLS CPI Research Database. For all five moments, we used the sample of regular price changes (as defined in our Section II) in New York City and the Connecticut suburbs of New York from January 1988 through December 1997. Let p_{ijt} denote the log price of item i in category j in month t , ω_j^{93} the BLS consumption expenditure weight of j , and ω_{ijt} the BLS weight on item i within category j in month t .

The geometric mean inflation rate in month t is

$$\pi_t = \sum_j \omega_j^{93} \sum_i \omega_{ijt} (p_{ijt} - p_{ijt-1}),$$

and the associated log price index is $\bar{p}_t = \sum_{s=0}^t \pi_s$. The first two statistics are the (across time) mean and standard deviation of quarterly inflation, 0.64% and 0.62%, respectively.

The fraction of items changing price in month t is defined in the text. The mean fraction for New York areas is 0.219 (the third statistic).

The weighted average price increase is

$$\sum_j \omega_j^{93} \left[\frac{\sum_{p_{ijt} > p_{ijt-1}} \omega_{ijt} (p_{ijt} - p_{ijt-1})}{\sum_{p_{ijt} > p_{ijt-1}} \omega_{ijt}} \right].$$

For the New York areas sample, its mean across time is 0.095 (the fourth statistic).

Finally, let z_{ijt} be the log deviation of a price from the price index in the category in month t , $z_{ijt} = p_{ijt} - \bar{p}_{jt}$. For each item we calculated the standard deviation of z_{ijt} across months with price increases. We then took the weighted mean of these standard deviations across items and categories:

$$\sum_j \omega_j^{93} \left[\frac{\sum_i \omega_{ij} \text{std}(z_{ijt})}{\sum_i \omega_{ij}} \right].$$

For the New York areas, the average standard deviation of new prices is 0.087 (the fifth statistic).

APPENDIX II: NUMERICAL SOLUTION METHOD

To solve the firm’s problem, we apply the projection method with Chebyshev collocation discussed by Judd (1998). Specifically, we approximate each of the two value functions $V_A(p,z)$, $V_N(p,z)$ by a sum of Chebyshev polynomials:

$$\hat{V}_A(p, z; a) = \sum_{i=0}^{M_P} \sum_{j=0}^{M_Z} a_{ij}^A T_i \left(2 \frac{p - p_m}{p_M - p_m} - 1 \right) T_j \left(2 \frac{z - z_m}{z_M - z_m} - 1 \right),$$

$$\hat{V}_N(p, z; a) = \sum_{i=0}^{M_P} \sum_{j=0}^{M_Z} a_{ij}^N T_i \left(2 \frac{p - p_m}{p_M - p_m} - 1 \right) T_j \left(2 \frac{z - z_m}{z_M - z_m} - 1 \right),$$

where $T_i(\cdot)$ is the i th-order Chebyshev polynomial; a_{ij}^A , a_{ij}^N for $i = 0, \dots, M_P$, $j = 0, \dots, M_Z$ is a set of coefficients. The $M_P \times M_Z$ -element matrices, \mathbf{a}^A and \mathbf{a}^N , denote these coefficients. We choose values for \mathbf{a}^A and \mathbf{a}^N to get the differences between the right- and left-hand sides of equations (10) and (11) to be close to zero at pre-determined grid points. Let p_j , for $j = 1, \dots, N_P$ denote the values of p satisfying $T_{N_P}(p) = 0$, where $N_P \geq M_P$. Similarly, choose z_j , with $j = 1, \dots, N_Z$. Let \mathbf{A}_P denote the $M_P \times N_P$ matrix with components $A_{P,ij} = T_{i-1}(p_j)$, for $i = 1, \dots, M_P$, $j = 1, \dots, N_P$, and \mathbf{A}_Z be the $M_Z \times N_Z$ matrix with components $A_{Z,ij} = T_{i-1}(z_j)$, for $i = 1, \dots, M_Z$, $j = 1, \dots, N_Z$. Let $\mathbf{R}_A(\mathbf{a})$, $\mathbf{R}_N(\mathbf{a})$ denote the $N_P \times N_Z$ matrices formed by evaluating the residuals of the equations (10) and (11) respectively, using the decision rule $p'(p, z; a)$ at the $N_P \times N_Z$ values of (p, z) . We select the $2M_P M_Z$ elements of \mathbf{a}^A , \mathbf{a}^N so that the $2M_P M_Z$ equations, $\mathbf{A}_P \mathbf{R}_A(\mathbf{a}) \mathbf{A}'_Z$ and $\mathbf{A}_P \mathbf{R}_N(\mathbf{a}) \mathbf{A}'_Z$ are zero.

STANFORD UNIVERSITY AND NBER
BANK OF CANADA

REFERENCES

Bils, Mark, Measuring Growth from Better and Better Goods, unpublished paper, University of Rochester, 2005.
Bils, Mark, and Peter J. Klenow, “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112 (2004), 947–985.

- Broda, Christian, and David E. Weinstein, "Defining Price Stability in Japan: A View from America," NBER Working Paper No. 13255, 2007.
- Burstein, Ariel T., and Christian Hellwig, "Prices and Market Shares in State-Dependent Pricing Models," unpublished paper, University of California at Los Angeles, 2006.
- Caballero, Ricardo, and Eduardo Engel, "Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy," *Review of Economic Studies*, 60 (1993), 95–119.
- , "Price Stickiness in Ss Models: Basic Properties," unpublished paper, Massachusetts Institute of Technology, 2006.
- Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12 (1983), 383–398.
- Caplin, Andrew, and John Leahy, "State-Dependent Pricing and the Dynamics of Money and Output," *Quarterly Journal of Economics*, 106 (1991), 683–708.
- , "Aggregation and Optimization with State-Dependent Pricing," *Econometrica*, 65 (1997), 601–625.
- Caplin, Andrew, and Daniel Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, 102 (1987), 703–725.
- Carvalho, Carlos, "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks," *Frontiers of Macroeconomics* 2 (2006), available at <http://www.bepress.com/bejm/frontiers/vol2/iss1/art1>.
- Cecchetti, Stephen G., "Staggered Contracts and the Frequency of Price Adjustment," *Quarterly Journal of Economics*, 100 (1985), 935–959.
- , "Indexation and Incomes Policy: A Study of Wage Adjustments in Unionized Manufacturing," *Journal of Labor Economics*, 5 (1987), 391–412.
- Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi, "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, 93 (2003), 15–37.
- Dotsey, Michael, Robert King, and Alexander Wolman, "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics*, 114 (1999), 655–690.
- , "Inflation and Real Activity with Firm-Level Productivity Shocks: A Quantitative Framework," unpublished paper, Federal Reserve Bank of Richmond, 2006.
- Gabriel, Peter, and Adam Reiff, "Estimating the Extent of Price Stickiness in Hungary: A Hazard-Based Approach," unpublished paper, Central Bank of Hungary, 2007.
- Gagnon, Etienne, "Price Setting During High and Low Inflation: Evidence from Mexico," International Finance Discussion Paper No. 896, Board of Governors of the Federal Reserve System, 2007.
- Gertler, Mark, and John Leahy, "A Phillips Curve with an Ss Foundation," NBER Working Paper No. 11971, 2006.
- Golosov, Mikhail, and Robert E. Lucas, Jr., "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115 (2007), 171–199.
- Hobijn, Bart, Federico Ravenna, and Andrea Tambalotti, "Menu Costs at Work: Restaurant Prices and the Introduction of the Euro," *Quarterly Journal of Economics*, 121 (2006), 1103–1131.
- Judd, Kenneth L., *Numerical Methods in Economics*, 1st ed. (Cambridge, MA: MIT Press, 1998).
- Kehoe, Patrick J., and Virgiliu Midrigan, "Sales and the Real Effects of Monetary Policy," Federal Reserve Bank of Minneapolis Working Paper No. 652, 2007.
- Klenow, Peter J., and Oleksiy Kryvtsov, "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" NBER Working Paper No. 11043, 2005.
- Klenow, Peter J., and Jonathan Willis, "Sticky Information and Sticky Prices," *Journal of Monetary Economics*, 54 (2007), 79–99.
- Mankiw, Gregory N., and Ricardo Reis, "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117 (2002), 1295–1328.
- Midrigan, Virgiliu, "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," unpublished paper, New York University, 2006.

- Moulton, Brent R., and Karin E. Moses, "Addressing the Quality Change Issue in the Consumer Price Index," *Brookings Papers on Economic Activity* 1 (1997), 305–349.
- Nakamura, Emi, and Jon Steinsson, "Five Facts About Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, forthcoming, 2008.
- Shapiro, Matthew D., and David W. Wilcox, "Mismeasurement in the Consumer Price Index: An Evaluation," in *NBER Macroeconomics Annual*, Ben S. Bernanke and Julio J. Rotemberg, eds. (Cambridge, MA: MIT Press, 1996).
- Taylor, John B., "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88 (1980), 1–23.
- U.S. Department of Labor, *BLS Handbook of Methods*, Bulletin 2490 (U.S. Government Printing Office, 1997).