FIVE FACTS ABOUT PRICES: A REEVALUATION OF MENU COST MODELS*

EMI NAKAMURA AND JÓN STEINSSON

We establish five facts about prices in the U.S. economy: (1) For consumer prices, the median frequency of nonsale price change is roughly half of what it is including sales (9–12% per month versus 19–20% per month for identical items; 11–13% per month versus 21–22% per month including product substitutions). The median frequency of price change for finished-goods producer prices is comparable to that of consumer prices excluding sales. (2) One-third of nonsale price changes are price decreases. (3) The frequency of price increases covaries strongly with inflation, whereas the frequency of price decreases and the size of price increases and price decreases do not. (4) The frequency of price change is highly seasonal: it is highest in the first quarter and then declines. (5) We find no evidence of upward-sloping hazard functions of price changes for individual products. We show that the first, second, and third facts are consistent with a benchmark menu-cost model, whereas the fourth and fifth facts are not.

I. INTRODUCTION

The nature of price setting has important implications for a range of issues in macroeconomics, including the welfare consequences of business cycles, the behavior of real exchange rates, and optimal monetary policy. We use Bureau of Labor Statistics (BLS) microdata underlying the consumer and producer price indices to document five basic features of price adjustment. We interpret this evidence through the lens of a benchmark menu cost model.

We begin by estimating the frequency of price change. Until recently, the best sources of information on U.S. pricing behavior were studies of price adjustment for particular products (Cecchetti 1986; Kashyap 1995), broader surveys of firm managers (Blinder et al. 1998), and evidence on the dynamics of industrial prices (Carlton 1986). The conventional wisdom from this literature was

*We would like to thank Robert Barro for invaluable advice and encouragement. We would like to thank Daniel Benjamin, David Berger, Leon Berkelmans, Craig Brown, Charles Carlstrom, Gary Chamberlain, Tim Erickson, Mark Gertler, Mike Golosov, Gita Gopinath, Teague Ruder, Oleksiy Kryvtsov, Gregory Kurtzon, Robert McClelland, Greg Mankiw, Ariel Pakes, Ricardo Reis, Roberto Rigobon, John Rogers, Ken Rogoff, Philippa Scott, Aleh Tsyvinsky, Randal Verbrugge, Michael Woodford, seminar participants at various institutions as well as our editor, Larry Katz, and anonymous referees for helpful comments and discussions. We particularly want to thank Mark Bils and Pete Klenow for thoughtful and inspiring conversations. We are grateful to Martin Feldstein for helping us obtain access to the data; without his help this work would not have been possible. We are grateful to the Warburg Fund at Harvard University for financial support.

The Quarterly Journal of Economics, November 2008

that prices adjusted on average once a year. Bils and Klenow (2004) dramatically altered this conventional wisdom by showing that the median frequency of price change for nonshelter consumer prices in 1995–1997 was 21%, implying a median duration of 4.3 months.

We use a substantially more detailed data set than Bils and Klenow (2004) that contains the micro-level price data underlying the nonshelter component of the Consumer Price Index (CPI).[1] This data set has been used by Hosken and Reiffen (2007, 2004) and Klenow and Kryvtsov (2008) to analyze price adjustment behavior. We find that temporary sales play an important role in generating price flexibility for retail prices in categories that account for about 40% of nonshelter consumer expenditures. Whereas the median frequency of price change including sales is 19%–20% per month, we find that the median frequency of nonsale price change for identical items is only 9%–12% per month depending on the time period and how we treat nonsale price changes over the course of sales and stockouts.

Our estimates of the median frequency of price change for identical items may be inverted to obtain estimates of the median duration of regular prices. Excluding product substitutions, these frequency estimates imply uncensored durations of regular prices of between 8 and 11 months. Yet, substitutions often truncate regular price spells. If we include price changes associated with product substitutions, the median frequency of nonsale price change increases by between 1 and 2 percentage points. This implies median durations until either the regular price changes or the product disappears at between 7 and 9 months.

The importance of temporary sales—and to a lesser extent substitutions—in generating price changes in the U.S. data draws attention to the question of whether the relative frequency of different types of price changes is an important determinant of the macroeconomic implications of price rigidity. In other words: "Is a price change just a price change?" An important lesson from the theoretical literature on price adjustment is that different types of price adjustments can have strikingly different macroeconomic

1. Bils and Klenow (2004) used the BLS Commodities and Services Substitution Rate Table for 1995–1997. This data set contains average frequencies of price changes and substitutions by disaggregated product categories over the 1995–1997 period. In contrast, the CPI research database contains the actual data series on prices underlying the Consumer Price Index for the 1988–2005 period. See Section II for a more detailed discussion of the data.

implications. For example, the Calvo (1983) model and the Caplin and Spulber (1987) model have very different macroeconomic implications for the same frequency of price change.

For this reason, an important focus of this paper is to document and contrast the empirical characteristics of the different types of price changes observed in U.S. consumer data. First, we document that sale price changes display markedly different empirical features than regular price changes. Sale price changes are much more transient than regular price changes; and in most cases where a price is observed before and after a sale, the price returns to its original level following the sale.

There are a number of reasons why it may be important to distinguish between sale and nonsale price changes. First, the transience of price adjustment associated with sales implies that a given number of price changes due to sales yield much less aggregate price adjustment than the same number of regular price changes (Kehoe and Midrigan 2007). Second, some types of sales may be orthogonal to macroeconomic conditions. Third, transitory sales are a much more pervasive phenomenon in retail prices than in wholesale prices, implying that temporary sales may be less responsive to shocks at the wholesale than at the retail level of production.

Price changes due to product substitutions are a second class of price changes that we argue is fundamentally different from the regular price changes typically emphasized by macroeconomists. This source of price flexibility is particularly important for durable goods. For example, the spring and fall clothing seasons in apparel and the new model year for cars are associated with a large number of price changes due to the introduction of new products. Many factors other than a firm's desire to change its price influence its decision to introduce a new product. The theoretical literature on price adjustment has shown that price changes that are motivated primarily by a large difference between a firm's current price and its desired price yield much greater aggregate price flexibility than those that are motivated by other factors (Caplin and Spulber 1987; Golosov and Lucas 2007). In state-dependent pricing models, it is therefore crucial to treat product substitutions separately from other types of price changes (Nakamura and Steinsson 2007). In contrast, time-dependent pricing models should arguably be calibrated to the frequency of price change including substitutions because in these models the timing of *all* price changes is exogenous.

We also present the first broad-based evidence on U.S. price dynamics at the producer level. To study this issue, we created a new data set on producer prices from the production files used by the BLS to construct the Producer Price Index (PPI). The median frequency of price change for finished-goods producer prices was 10.8% in 1998–2005; it was 13.3% for intermediate-goods producer prices; and it was 98.9% for crude materials. Moreover, we document a high correlation between the frequency of nonsale consumer price changes and the frequency of producer price changes at a very disaggregated level. The price rigidity in finished-goods producer prices is comparable to the rigidity of consumer prices excluding sales but substantially greater than the rigidity of consumer prices including sales.

There is a tremendous amount of heterogeneity across sectors in both the frequency of price change and the importance of temporary sales. Different summary statistics on price flexibility therefore give very different answers regarding the degree of price flexibility in the U.S. economy. Following Bils and Klenow (2004), we focus on the weighted median frequency of price adjustment across categories. Excluding sales lowers the median frequency of price change of consumer prices by over 50%, while it lowers the mean frequency of price change by only about 20%. This is because sales are concentrated in sectors of the economy, such as food and apparel, that have a frequency of price change close to the median frequency of price change across sectors.

There is no model-free way of selecting what is the appropriate summary statistic to describe the amount of aggregate price flexibility in an economy in which the frequency of price change varies across sectors from over 90% per month to less than 5% per month. In Nakamura and Steinsson (2007), we calibrate a multisector menu cost model to the sectoral distribution of the frequency and absolute size of price changes excluding sales. We use this model to investigate which statistic about price rigidity is most informative about the degree of monetary nonneutrality in the economy. The degree of monetary nonneutrality implied by our multisector model is triple that implied by a single-sector model calibrated to the mean frequency of price change of all firms but similar to that implied by a single-sector model calibrated to the median frequency of price change.[2]

2. Carvalho (2006) studies the effect of heterogeneous price rigidity in time-dependent models. For the Calvo model, he finds that a single-sector model

The second feature of price change that we investigate is the fraction of price changes that are price decreases. We find this fraction to be roughly one-third in both consumer prices excluding sales and finished-goods producer prices. We present a benchmark menu cost model along the lines of Golosov and Lucas (2007) and show that the fraction of price changes that are decreases helps pin down the key parameters of this model. The third feature of price change that we investigate is how the frequency and size of price change covaries with the inflation rate. We find that the frequency of price increases covaries quite strongly with the rate of inflation, whereas the frequency of price decreases and the size of price increases and decreases do not. This fact provides a natural test for our calibrated benchmark menu cost model. The fourth feature of price change that we investigate is the extent of seasonal synchronization. We find that price rigidity is highly seasonal both for consumer and producer prices. Prices are substantially more likely to change in the first quarter than in other quarters.

The fifth and final issue that we investigate is the hazard function of price change. The main empirical challenge in estimating the hazard function of price change is the fact that heterogeneity in the level of the hazard function across products—if not properly accounted for—leads to a downward bias in the slope of the hazard function. We use the empirical model of Meyer (1990) to account for heterogeneity. The hazard function of consumer prices including sales is steeply downward sloping for sectors with frequent sales. In contrast, the estimated hazard function of price change for both consumer prices excluding sales and producer prices is slightly downward sloping for the first few months and then mostly flat. The only substantial deviation from a flat hazard after the first few months is a large spike in the hazard at twelve months for services and producer prices.[3] We show that menu cost models can give rise to a wide

calibrated to the mean duration of price spells in the economy replicates the degree of monetary nonneutrality in a multisector model. We present estimates of the mean duration in Table I.

3. Earlier empirical work on the hazard function of price changes includes Cecchetti (1986), Jonker, Folkertsma, and Blijenberg (2004), Álvarez, Burriel, and Hernando (2005), Baumgartner et al. (2005), Campbell and Eden (2005), Dias, Robalo Marques, and Santo Silva (2005), Fougére, Bihan, and Sevestre (2005), and Goette, Minsch, and Tyran (2005). Empirical support for upward-sloping hazard functions appears to arise mostly in studies in which almost all price changes are increases. Several of these papers use the conditional logit specification to account for unobserved heterogeneity. Unfortunately, this specification yields inconsistent estimates of the shape of the hazard function (Willis 2006).

variety of hazard functions, depending on the relative importance of inflation and idiosyncratic shocks. The hazard function implied by our calibrated benchmark menu cost model is sharply upward sloping for the first few months.

Klenow and Kryvtsov (2008) report related statistics regarding the frequency of price change, the relationship between the size and frequency of price adjustments and the ination rate, and the hazard function of price change. Their frequency of price change estimates are very similar to ours, although their interpretation of these statistics is somewhat different. They estimate the median implied duration of regular prices, including substitutions, to be 7.2 months. Their estimator is similar to the one we use in line 10 of Table I.[4] A time-weighted average of our estimates in line 10 of Table I is 7.5 months. For regular prices excluding substitutions, they report 8.7 months and a time-weighted average of our estimates is also 8.7 months (line 6 of Table I). They report a median implied duration of 9.3 months based on adjacent regular prices. A time-weighted average of our estimates is 9.6 months (line 3 of Table I). The range of numbers we report has a higher upper bound because we split the sample and report results separately for the subsample 1998–2005 for which the rate of inflation was lower.

An important body of work on price adjustment in Europe has been carried out by the Inflation Persistence Network of the European Central Bank. Álvarez et al. (2006) and Dhyne et al. (2006) summarize the conclusions of a number of papers on the frequency of price adjustment in consumer prices for the countries of the Euro area. Vermeulen et al. (2006) summarize analogous studies on producer prices in the Euro area. Fabiani et al. (2004) summarize survey evidence on price adjustment in the Euro area. Our findings regarding the frequency of price change, the relationship between the frequency of price increases and inflation, and the seasonality of price changes find strong support in a number of European countries.[5]

4. The main difference between the estimate in line 10 of Table I and the estimate used by Klenow and Kryvtsov (2008) is that their estimator does not condition on the state (regular price, sale, or stockout), whereas ours does.

5. A number of other recent papers have studied the frequency and size of price changes using disaggregated price data, including Lach and Tsiddon (1992), Baharad and Eden (2004), Konieczny and Skrzypacz (2005), Hobijn, Ravenna, and Tambalotti (2006), Midrigan (2006), Kackmeister (2007), and Gopinath and Rigobon (2008).

The paper is organized as follows. In Section II, we describe the data. In Section III, we present evidence on the frequency of price change, the fraction of price changes that are price increases, the frequency of product turnover, the absolute size of price changes, and temporary sales. In Section IV, we present and calibrate a benchmark menu cost model. In Section V, we present evidence on how the frequency and size of price changes vary with inflation. In Section VI, we present evidence on the seasonality of price changes and sales. In Section VII, we present our estimates of the hazard function of price change. Section VIII concludes.

II. The Data

We use two data sets gathered by the BLS in this paper. The first is the CPI Research Database. This is a confidential data set that contains product-level price data used to construct the CPI. The second is an analogous data set of producer prices that we have created from the production files underlying the PPI. We will refer to this data set as the PPI Research Database. The CPI Research Database has been used by Hosken and Reiffen (2007, 2004) and Klenow and Kryvtsov (2008).[6] The PPI Research Database has not been used before.

II.A. The CPI Research Database

Each month the BLS collects prices of thousands of individual goods and services for the purpose of constructing the CPI. The CPI Research Database contains the nonshelter component of this data set from 1988 to the present. The goods and services included in the CPI Research Database constitute about 70% of consumer expenditures. Prices are sampled in 87 geographical areas across the United States. Prices of all items are collected monthly in the three most populous locations (New York, Los Angeles, and Chicago). Prices of food and energy are collected monthly in all other locations as well. Prices of other items are collected bimonthly. In most of our analysis, we use only monthly observations.[7]

6. Bils and Klenow (2004) used the BLS Commodities and Services Substitution Rate Table for 1995–1997. The Substitution Rate Table contains the average frequency of price change including product substitutions and imputed missing values for all products in the CPI.

7. As a robustness test, we compared the bimonthly frequency of price change in the portion of our data set that is sampled bimonthly to the bimonthly frequency of price change in the portion of our data set that is sampled monthly. The

The CPI Research Database identifies products at an extremely detailed level. In general, two products are considered different products in the database if they carry different bar codes. In addition, the same product at two different outlets is considered different products in the database. An example of a product in the database is a two-liter bottle of Diet Coke sold at a particular supermarket in New York. The database reports whether a product was "on sale" when its price was sampled in a particular month.[8] We use this sales flag to calculate statistics about the frequency and size of price change excluding sales. Some prices in the database are derived from the price of other products rather than being based on a collected price. We drop all such observations.[9]

The BLS divides products into so called Entry Level Items (ELIs). Examples of ELIs are "Carbonated Drinks," "Washers & Driers," "Woman's Outerwear," and "Funeral Expenses." Before 1998, the BLS divided the data set into roughly 360 ELIs. In 1998, the BLS revised the ELI structure of the data set. Since then, it has divided the data set into roughly 270 ELIs. The revision in the ELI structure of the data set in 1998 implies that in many cases we calculate statistics separately for the periods 1988–1997 and 1998–2005. Most of our results are similar for the two sample periods. For concreteness, we will refer to the estimates for the latter period in the text unless we indicate otherwise. In all of the statistics we present on the frequency and size of price changes, we focus on weighted medians and means across ELIs.

bimonthly frequency of price change is slightly lower in the bimonthly data than in the monthly data.

8. BLS field agents are instructed to mark a price as a sale price if it is considered by the outlet to be temporarily lower than the regular selling price and is available to all consumers. In practice, the BLS sales flag corresponds roughly to whether there is a "sale" sign next to the price when it is collected. If an outlet never sells a product at its "regular" price—that is, the product is always on sale—the BLS field agent is directed not to label it as a sale price. Sales available to customers with savings or discount cards are reported as sales only if the outlet confirms that more than 50% of its customers use these cards. Bonus items may be reported as sales as long as they satisfy the normal criteria for sales described above. Three categories in which the sale flag is never used by design are new and used cars and airfares. The approach that is used to collect price data for these categories is quite different from the procedure used to collect price data for other categories. The price series for new cars combines data on list prices with data on average "deals" obtained by consumers. The used car data are based on an index of used car prices. The data on airline tickets are based on a sample of tickets from the U.S. Department of Transportation data bank. Chapter 10 of the unpublished BLS manual *Price Reporting Rules* contains a more detailed description of the definition of sales used by the BLS.

9. Chapter 17 of the *BLS Handbook of Methods* (U.S. Department of Labor 1997) contains a far more detailed description of the consumer price data collected by the BLS.

The weights we use are CPI expenditure weights from 1990 for our statistics on the period 1988–1997 and weights from 2000 for our statistics on the period 1998–2005. The statistics at the ELI level are unweighted averages within the ELI.

II.B. The PPI Research Database

We construct the PPI Research Database from the production files underlying the U.S. PPI. The earliest prices in the database are from the late 1970s. For the period 1988–2005, which we focus on in most of our analysis, the PPI Research Database contains data for categories that constitute well in excess of 90% of the value weight for the Finished Goods PPI.[10] An important difference between the CPI and the PPI is that the PPI is collected by BLS through a survey of firms. This methodology introduces greater concerns about data quality than in the CPI where BLS agents sample prices of products "on the shelf." Stigler and Kindahl (1970) criticized the methodology used to gather the PPI data because it relied on "list" prices rather than transaction prices. Since then, the BLS has revamped its data collection methodology to focus expressly on collecting actual transaction prices. Specifically, the BLS requests the price of actual shipments transacted within a particular time frame.[11] Note that many of the transactions for which prices are collected as part of the PPI are part of implicit or explicit long-term contracts between firms and their suppliers. The presence of such long-term contracts makes interpreting PPI data more complicated than interpreting CPI data, as we discuss further in Section III.D.

Another difference between the consumer and producer price data is that the definition of a good in the PPI Research Database typically includes information about the buyer of the product as well as a detailed set of product and transaction characteristics. The definition is meant to capture all "price-determining variables." Price-determining variables may include the buyer, the quantity being bought, the method of shipment, the transaction terms, the day of the month on which the transaction takes place, and product characteristics. This implies that if a seller charges a different price to different customers, the BLS will collect prices for a transaction involving the same customer month after month.

10. The weights referred to here are the post-1997 value weights used to construct the Finished Goods PPI.

11. See Chapter 14 of the *BLS Handbook of Methods* (U.S. Department of Labor 1997) for a more detailed description of BLS procedures.

The price data in the PPI are collected in two steps. When a product is first introduced into the data set, the BLS collects "checklist" information by conducting a personal visit to the firm. The checklist contains information on characteristics of the product, buyer, and seller as well as the terms and date of the transaction. The checklist also contains information on various types of addendums to the standard price, for example, whether the price involves a trade or quantity discount or other type of discounts or surcharges. Once the product is initiated, price information is collected using a repricing form. The repricing forms are mailed or faxed. If the form is not returned, a BLS industry analyst calls the firm to collect information over the phone. The checklist information is updated when an industry is resampled every five to seven years.

An important concern with the methods used to collect the PPI data is that the repricing form used to update prices in the PPI first asks whether the price has changed relative to the previous month and then asks the respondent to report a new price if the price did change. This structure of the repricing form may introduce a bias toward no change into the data. To evaluate sensitivity of the price data to the method used to collect prices, we compared the behavior of prices during the anthrax scare of 2001 to the behavior of prices during other time periods. In October and November 2001, all mail to government agencies was rerouted, and PPI collected all prices by a phone survey. Controlling for the relationship between the frequency of price change and inflation, we found no significant differences in the frequency of price change in 2001 versus the same months in other years.[12] Another feature of the data that suggests that the producer price data contain meaningful information is the high correlation between the frequency of price change for manufacturing prices and consumer prices excluding sales documented in Section III.E.

The BLS constructs indexes for three different stages of processing: finished goods, intermediate goods, and crude materials. We focus attention on finished goods, but also report basic results

12. The idea of using the anthrax scare of 2001 for this purpose is due to Gopinath and Rigobon (2008). Our approach is slightly different than theirs. We compare the frequency of price change during the anthrax scare with the frequency of price change in the same months of other years rather than the adjacent months because the frequency of producer prices is highly seasonal. Specifically, we regress the absolute size and frequency of price change in October and November of each year on a dummy for 1998 and the PPI. The coefficient on the "anthrax dummy" in the frequency regression is 0.0057 with a standard error of 0.0084, and in the absolute size regression it is 0.0041 with a standard error of 0.0030. Neither coefficient is statistically significant.

for intermediate goods and crude materials. Our method for calculating statistics at various levels of aggregation in the PPI is somewhat more complicated than in the CPI. The most detailed grouping in the PPI Research Database is the cell code. We do not attempt to construct value weights at this level because there is a substantial amount of churning in the cell codes used in the PPI from year to year. We instead obtain value weights for the PPI at the four-digit commodity code level. We then construct statistics on the frequency of price change at the four-digit commodity code level in the following way. First, we calculate the unweighted average frequency of price change within cell codes. Next, we calculate the unweighted median frequency of price change across cell codes within the four-digit commodity code. Finally, we construct aggregate statistics by taking value-weighted medians over the median price change frequencies at the four-digit commodity code level. This approach is similar to the approach taken by Gopinath and Rigobon (2008) for import and export price data. For the purpose of matching PPI categories with CPI ELIs, we construct unweighted medians within six-digit and eight-digit product categories.

III. How Often and How Much Do Prices Change?

In this section, we present statistics on the frequency and size of price changes in the U.S. economy. While this may seem straightforward, there are a number of important issues to be considered. We therefore discuss the construction of these statistics in some detail. An important lesson from the theoretical literature on price adjustment is that different types of price adjustments have substantially different macroeconomic implications. The menu cost model has the strong prediction that the products "selected" to change their prices in response to an expansionary monetary shock disproportionately have prices that are far below their current optimum level. As a consequence of this selection effect, the price level responds relatively rapidly to the shock, and the effects of the shock on aggregate output are relatively transient (Caplin and Spulber 1987; Golosov and Lucas 2007). In contrast, if the timing of price changes is random, as in the Calvo (1983) model, monetary shocks have significantly more persistent effects on output.

Motivated by this theoretical literature, we distinguish between three broad classes of price changes: (1) regular price changes for identical items, (2) temporary sales, and (3) price changes due to product substitution. We present statistics for these different types of price changes separately. We document that sales have empirical characteristics very different from regular price changes. Price changes associated with sales are highly transient and the price of the product returns to the old regular price after most sales. Kehoe and Midrigan (2007) argue that transitory price changes, such as temporary sales, yield much less aggregate price flexibility than an equal number of permanent price changes.

Our motive for distinguishing between price changes associated with product substitution and price changes for identical items is that new product introduction is motivated by many factors other than a firm's desire to change its price. Product substitutions are by far most common in the apparel and transportation sectors. In these sectors, the introduction of new products is driven by factors such as the fall and spring clothing seasons, and the new model year for automobiles. Although the introduction of the new spring clothing line may be a good opportunity for a firm to adjust its price, this type of new product introduction does not occur *because* of the firm's desire to adjust its price, limiting the strength of the selection effect.

Two important measurement issues arise. First, how do we identify the presence of temporary sales? The BLS gathers data on whether a product was "on sale" when its price was sampled in a particular month. We use this sale flag as our primary measure of the presence of temporary sales.[13] We also consider identifying sales based on a "sale filter" in Section III.H. Second, estimating the frequency of adjustment for regular prices is complicated by times when firms' regular prices are not observed because of sales and stockouts. In the absence of a theory of sales and stockouts, there is no unique way of filling in these gaps in the regular price series.[14] We present four estimates of the frequency of regular

13. This is the approach adopted by Bils and Klenow (2004) and the main approach used by Klenow and Kryvtsov (2008).

14. One can object to the notion that it is meaningful to say that a latent regular price exists during sales and stockouts. After all, the good is not available at the regular price in the case of sales and not available at all in the case of stockouts. However, the fact that the regular price of the product is often the same after sales and stockouts as before suggests that the old regular price should be viewed as a latent state variable in the firm's pricing problem. Also, anecdotial evidence suggests that many sales take the form of a discount relative to the product's most recent regular price, suggesting that the latent regular price influences the sale price. Finally, high-frequency data suggest that most sales and stockouts last for a much shorter period than an entire month. This implies that the products may

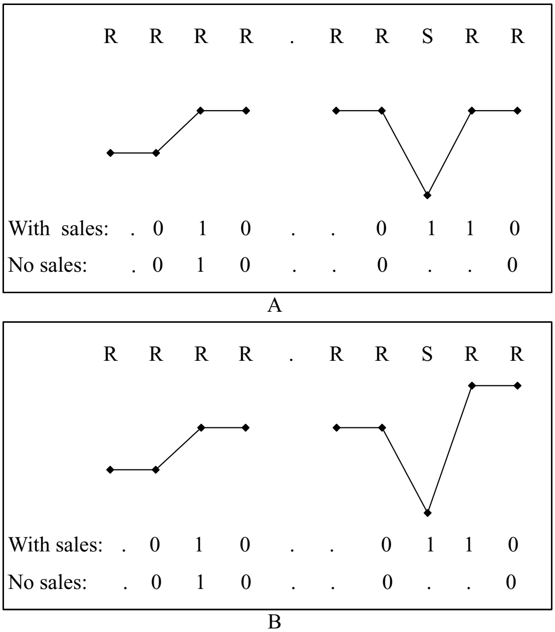

FIGURE I

Construction of Price Change Variables with and without Sales

Notes. Each panel reports the first ten observations for a hypothetical price series. The top row of each panel records the values of the sales flag for the ten observations. The letter "R" denotes "regular price" while the letter "S" denotes "sales." Below the flag is a graph of the evolution of the price of the product. At the bottom of each panel are two indicator variables. The first records price changes, while the second records regular price changes.

price change corresponding to four different treatments of regular prices during sales and stockouts.

The simplest approach is to estimate the frequency of regular price change during periods when the presence or absence of regular price changes is directly observable, that is, when contiguous nonsale observations are made. Figure I graphically illustrates this simple procedure. The two panels in the figure report the first ten observations for two hypothetical products. Each panel contains a graph of the evolution of the price of the product for these ten observations. At the top of each panel in the figure, we record with the letter "R" and the letter "S" whether each observation is a regular price or a sale, respectively. At the bottom of each panel,

in fact be available at some unobserved regular price for a large fraction of the month in question.

there are indicator variables that record price changes including and excluding sales. First, notice that the price change variable is missing for the first observation. This is because the price in the previous month is not observed. Second, notice that the fifth price observation is missing. This yields two missing values in the price change variable. Third, notice that the eighth observation is a sale. The sale yields two price changes in the "raw" price change variable. However, dropping the sale observation from the data set yields two missing observations for the regular price change variable. In this example, our estimate of the frequency of regular price change based on contiguous observations would therefore be $1/5 = 20\%$.

This procedure has the advantage that it does not make any assumption regarding the behavior of the unobserved regular price series over the course of the sale. It provides a direct estimate of the extent of price flexibility that does not arise from sales. Yet, this procedure has the disadvantage that it does not incorporate in any way information in the data set about whether a product's regular price is the same before and after sales and stockouts. If regular prices follow a constant hazard model, then the frequency of regular price change during the periods when regular prices are observed provides a good estimate of the frequency of regular price change during periods when regular prices are not observed.[15] However, the behavior of regular prices may be systematically different over the course of sales and stockouts than during other periods. In this case, the implied durations of regular prices associated with this method are likely to be systematically biased.

Our second procedure for calculating the frequency of price change assumes that the latent regular price series is equal to the last observed regular price until a change is observed. This is the procedure used by Klenow and Kryvtsov (2005) and Gopinath and Rigobon (2008). In the context of the menu cost model, this procedure would be appropriate if regular prices were systematically readjusted at the end of (but not during) sales or stockouts. To implement this procedure, we carry forward the last observed regular price through sale and stockout periods and calculate the

15. It may at first seem that the procedure based only on contiguous observations necessarily underestimates the frequency of regular price change because it does not count regular price changes during sales and stockouts. In this regard, it is important to notice that while using only contiguous observations leads one to drop a price change during the sale in Panel B, it also causes one to drop a "no change" during the sale in period A and during the stockout in both panels. If the probability of regular price change is the same during sales and stockouts as it is during other periods (as in a constant hazard model) then dropping the sale price observations has no systematic effect on the estimates.

frequency of price change of the resulting series.[16] This procedure has the appealing feature that it captures the price changes and "no changes" after sales and stockouts in a particularly simple way. However, it assumes that only one regular price change can occur over the course of a sale or stockout. It therefore assumes a maximum amount of rigidity during sales and stockouts.

Our third procedure makes the assumption that the latent regular price series evolves stochastically over the course of a sale. In the context of the menu cost model, this procedure would be appropriate if regular prices were adjusted both during and after sales. The key difference between this procedure and the previous ones is that it allows for more than one price change over the course of the period when regular prices are unobserved. To implement this procedure, we take the following weighted average: $(1-s)f+sf'$, where f is the measure of the frequency of price change based on contiguous nonsale, nonstockout observation, f' is a direct estimate of the frequency of regular price change during one- and two-month sales, and s is the fraction of price change observations corresponding to sales.[17] Our fourth procedure is analogous to the third procedure except that it estimates a separate process for latent regular prices during both sales and stockouts. This is a very similar procedure to the one used by Klenow and Kryvtsov (2008).

16. We carry forward prices if they are followed by another regular price within five months. For longer gaps, we do not fill in the missing observations. It is not obvious how to construct statistics on price flexibility for goods that are not always available. Analogous to stockouts, almost all stores close at night. One could therefore say that prices all rise to infinity at night as well as during stockouts. Consider one economy with 24-hour stores and another with 12-hour stores that both reset all prices on January 1. One measure of price flexibility would be the frequency of price change relative to the total amount of time the good is available. According to this measure the second economy has twice as much price flexibility as the first. Yet prices in the latter economy would not respond more rapidly to aggregate shocks. In contrast, suppose all goods in the economy are available for only one month a year and firms reset their prices at that time. In this economy, prices are perfectly flexible even though the price of each good changes only once a year. The key distinction is whether prices reflect current economic conditions. This distinction motivates our decision to carry prices forward only for four months or less.

17. We calculate $f'=\omega_1 f'_1+(1-\omega_1)f'_2$, where f'_1 and f'_2 are the monthly frequency of regular price change during one- and two-periods sales, respectively. These frequencies are estimated using the method described in Section III.B.; ω_1 is the fraction of sales that are one-period sales. In small samples, this procedure yields an upward-biased estimate of the probability of price change during sale and missing periods due to Jensen's inequality. We have considered other weighted averages of sales spells of different lengths; this choice makes little difference because most sales spells are short. We only make use of cases where a price is observed before and after the event in calculating the probability of price change over the course of sales and stockouts. In particular, clearance sales do not contribute to these statistics.

Following Bils and Klenow (2004) and Dhyne et al. (2006), we have focused on estimating the frequency of price change. An alternative empirical strategy is to record the duration of each price spell and calculate the weighted median duration across all price spells. However, the presence of a large number of censored price spells complicates this approach. To account for right censoring, one must estimate a hazard model. This is a challenging problem because of the presence of heterogeneity. Left censoring is particularly problematic in applications with heterogeneity. The standard practice in the duration literature is to drop left-censored spells. This introduces an initial-conditions problem that biases the estimated duration downward in the presence of heterogeneity (Heckman and Singer 1986). Intuitively, longer spells are more likely to be left censored.

III.A. The Frequency of Price Change: Consumer Prices

Table I reports estimates of the frequency of price change for nonshelter goods and services in the CPI. The first two columns in the table report estimates for the median frequency of price change excluding and including both sales and substitutions.[18] Our four estimates of the median frequency of regular price change for identical items range from 8.7% per month to 11.9%, depending on the sample period and treatment of missing observations. We define the corresponding median implied duration to be $d = -1/\ln(1-f)$, where f is the median frequency.[19] These estimates therefore imply median durations of eight to eleven months. Procedures 3 and 4 yield higher estimates than Procedure 1 because the frequency of regular price change over the course of sales and stockouts is estimated to be on average about 2 percentage points higher than during other periods. Including substitutions raises the frequency of price change by 1–2 percentage points (Panel C).[20]

18. These statistics are estimated by first calculating the mean frequency of price change for each ELI and then taking a weighted median across ELIs.

19. A constant hazard λ of price change implies a monthly probability of a price change equal to $f = 1 - e^{-\lambda}$. This implies $\lambda = -\ln(1-f)$ and $d = 1/\lambda = -1/\ln(1-f)$. In the case of statistics where substitutions are excluded, the implied duration is an estimate of duration where product exit is viewed as a censoring event. In other words, it is a measure of the median uncensored duration.

20. Our procedure in row (10) of Table I is similar to the procedure used by Klenow and Kryvtsov (2008). The main difference is that our procedure allows for a different frequency of price change during sales/stockout than during periods when the regular price is observed.

TABLE I
FREQUENCY OF PRICE CHANGE IN THE CPI

	Median frequency		Median implied duration		Mean frequency		Mean implied duration	
	1988–1997 (%)	1998–2005 (%)	1988–1997 (months)	1998–2005 (months)	1988–1997 (%)	1998–2005 (%)	1988–1997 (months)	1998–2005 (months)
A. Including sales								
Excluding substitutions	20.3	19.4	4.4	4.6	23.9	26.5	8.3	9.0
Including substitutions	21.7	20.5	4.1	4.4	25.2	27.7	7.5	7.7
B. Excluding sales and substitutions								
Contiguous observations	11.1	8.7	8.5	11.0	18.7	21.1	11.6	13.0
Carry regular price forward during sales and stockouts	11.2	9.0	8.4	10.6	18.6	20.9	11.0	12.3
Estimate frequency of price change during sales	11.5	9.6	8.2	9.9	19.0	21.3	11.2	12.5
Estimate frequency of price change during sales and stockouts	11.9	9.9	7.9	9.6	18.9	21.5	10.8	11.7
C. Excluding sales, including substitutions								
Contiguous observations	12.7	10.9	7.4	8.7	20.4	22.8	9.3	9.8
Carry regular price forward during sales and stockouts	12.3	10.6	7.6	8.9	19.7	22.0	9.6	10.4
Estimate frequency of price change during sales	12.8	11.3	7.3	8.3	20.8	22.8	9.2	9.8
Estimate frequency of price change during sales and stockouts	13.0	11.8	7.2	8.0	20.7	23.1	9.0	9.3

Notes. All frequencies are reported in percent per month. Implied durations are reported in months. "Median frequency" denotes the weighted median frequency of price change. It is calculated by first calculating the mean frequency of price change for each ELI and then taking a weighted median across ELIs within the major group using CPI expenditure weights. The "Median implied duration" is equal to $-1/\ln(1-f)$, where f is the median frequency of price change. "Mean frequency" denotes the weighted mean frequency of price change. "Mean implied duration" denotes the weighted implied duration of price change. It is calculated by first calculating the implied duration for each ELI as $-1/\ln(1-f)$, where f is the frequency of price change for a particular ELI, and then taking a weighted mean across ELI's using CPI expenditure weights.

In contrast, the frequency of price change for identical items including sales was 19.4% for 1998–2005 and 20.3% for 1988–1997, implying median durations of 4.6 months and 4.4 months, respectively. The frequency of regular price change is therefore roughly 50% lower than the frequency of price change including sales. Adjusting for sales makes such a large difference not only because sales are common in the data—the expenditure-weighted fraction of price changes due to sales is 21.5%—but also because of the uneven distribution of sales across goods. Table II reports the fraction of price change due to sales by major group. On the one extreme, 87.1% of price changes in apparel and 66.8% of price changes in household furnishings are due to sales. On the other, virtually no price changes in utilities, vehicle fuel, and services—a category that has an expenditure weight of 38.5%—are due to sales.

The sectors that have few sales tend to have either very high (utilities, vehicle fuel and travel) or very low (services) unadjusted frequencies of price change. The sales adjustment is therefore concentrated in sectors that start off with a frequency of price change that is relatively close to the median frequency of price change. This heterogeneity in the prevalence of sales implies that the median frequency of price change drops by roughly 50% when sales are excluded, rather than 21.5%.[21]

To see this point more clearly, consider the three-sector example presented in Table III. Suppose the three sectors in the economy are services, food, and gasoline. Each has an expenditure weight of 1/3. Prices of services change once a year and have no sales. Prices of food change every other month, but 3/4 of these price changes are sales. The price of gasoline changes every month and gasoline never goes on sale. In this example, as in our data, sales are concentrated in the sector that is in the middle of the distribution of price change frequency. Adjusting for sales sector by sector yields a median frequency of regular price change of 1/8 and a median duration of 8 months. However, applying a blanket adjustment of 3/12 to all sectors—the overall fraction of

21. Bils and Klenow (2004) also present a statistic on the frequency of price change adjusted for sales. Because of data limitations, they were not able to adjust for sales at the good level. Instead, they adjusted the median frequency of price change by the fraction of price changes due to sales in the entire data set. This procedure yields an estimate of the sales-adjusted median duration of 5.5 months. It is a valid adjustment for sales under the assumption that sales account for the same fraction of price changes in all sectors. As we discuss, this assumption is dramatically at odds with the data.

TABLE II
FREQUENCY OF PRICE CHANGE BY MAJOR GROUP IN 1998–2005

		Regular prices				Prices				Sales	
		Median				Median					
Major group	Weight	Freq.	Impl. dur.	Mean freq.	Frac. up	Freq.	Impl. dur.	Mean freq.	Frac. up	Frac. price ch.	Frac. obs.
Processed food	8.2	10.5	9.0	10.6	65.4	25.9	3.3	25.5	54.7	57.9	16.6
Unprocessed food	5.9	25.0	3.5	25.4	61.2	37.3	2.1	39.5	53.3	37.9	17.1
Household furnishing	5.0	6.0	16.1	6.5	62.9	19.4	4.6	20.6	49.0	66.8	21.2
Apparel	6.5	3.6	27.3	3.6	57.1	31.0	2.7	30.1	36.1	87.1	34.5
Transportation goods	8.3	31.3	2.7	21.3	45.9	31.3	2.7	22.2	44.0	8.0	2.7
Recreation goods	3.6	6.0	16.3	6.1	62.0	11.9	7.9	13.7	51.3	49.1	10.9
Other goods	5.4	15.0	6.1	13.9	73.7	15.5	5.9	20.6	61.3	32.6	15.3
Utilities	5.3	38.1	2.1	49.4	53.1	38.1	2.1	49.4	53.1	0.0	0.0
Vehicle fuel	5.1	87.6	0.5	87.4	53.5	87.6	0.5	87.5	53.4	0.0	0.3
Travel	5.5	41.7	1.9	43.7	52.8	42.8	1.8	44.4	52.2	1.5	2.1
Services (excl. travel)	38.5	6.1	15.8	8.8	79.0	6.6	14.6	9.1	76.8	3.1	0.5
All sectors	100.0	8.7	11.0	21.1	64.8	19.4	4.6	26.5	57.1	21.5	7.4

Notes. All frequencies are reported in percent per month. Durations are reported in months. Fractions are reported as percentages. Regular prices denote prices excluding sales. "Weight" denotes the CPI expenditure weight of the major group; "median freq." denotes the weighted median frequency of price change. It is calculated by first calculating the mean frequency of price change for each ELI and then taking a weighted median across ELIs within the major group using CPI expenditure weights. The other median statistics in this table are calculated in an analogous manner: "median impl. dur." is equal to $-1/\ln(1-f)$, where f is the median frequency of price change. "Mean freq." denotes the expenditure weighted mean frequency of price change; "frac. up" denotes the median fraction of price changes that are price increases; "frac. price ch." and "frac. obs." denote the expenditure weighted mean fraction of price changes that are due to sales and fraction of observations that are sales. The sector weights add up to 97.4% because used cars are not included in any sector.

TABLE III
SALES ADJUSTMENT WHEN SALES ARE CONCENTRATED IN CERTAIN SECTORS

	Services	Food	Gasoline
Expenditure weight	1/3	1/3	1/3
Frequency of price change	1/12	1/2	1
Implied duration of price spells	12 months	2 months	1 month
Fraction of price changes due to sales	0	3/4	0
Frequency of regular price change	1/12	1/8	1
Implied duration of regular price spells	12 months	8 months	1 month
Assuming a constant fraction of price changes due to sales:			
Frequency of regular price change	1/16	3/8	9/12
Implied duration of regular price spells	16 months	2.66 months	1.33 months

Notes. In this example the expenditure-weighted fraction of price changes due to sales is 3/12. Assuming that the fraction of price changes due to sales is the same across sectors, the frequency of regular price change equals the frequency of price change multiplied by $1 - 3/12 = 9/12$. For simplicity, we assume that only one price change can occur per month in this example.

price changes due to sales in the entire economy—yields a median frequency of price change of 3/8 and a median duration of 2.67 months.

There is a huge amount of heterogeneity in the frequency of regular price change across sectors in the U.S. economy (Table II). Furthermore, the distribution of the frequency of regular price change is very right skewed. Most of the mass of the distribution lies below a frequency of regular price change of 12%, whereas categories such as vehicle fuel have a frequency of price change substantially higher than 50%. As a consequence, the mean frequency of regular price change is almost twice the median frequency of regular price change. Table I reports that the weighted mean frequency of price change in the 1998–2005 period is 26%–28% including sales and 21%–22% excluding sales. These estimates are consistent with the estimates of Klenow and Kryvtsov (2005). Table I also reports the weighted mean implied durations for the various alternative procedures for calculating the frequency of price change. Jensen's inequality implies that the mean implied duration is not the same as the implied duration for a product with the mean frequency of price change. Our estimates of the mean implied duration lie between 9 and 13 months.

One issue that arises in considering the macroeconomic implications of sales is that the quantity sold on sale is likely to be disproportionately large relative to the fraction of time the product is on sale. In the extreme, suppose all of the volume for a particular product is sold on sale. In this case, does the rigidity of the regular price influence real quantities? The answer to this question depends on whether sale prices are set entirely independently from nonsale prices or sale prices are partially set relative to a product's regular price (e.g., 20% discount). In the second case, even if all products are sold on sale, the rigidity of regular prices still influences real quantities through its effect on the sales prices.

III.B. Behavior of Prices after Sales

Sales exhibit empirical features markedly different from regular price changes.[22] Table IV presents statistics on sales for the four major groups for which sales are most important: processed food, unprocessed food, household furnishings, and apparel. First, sales are much shorter than regular price spells. The fraction of sales that last just one period ranges between 35% and 60% in the four major groups in Table IV, and the average length of sales is just 1.8–2.3 months. Longer sales are more prevalent in apparel and household furnishings because clearance sales tend to be longer than other sales and are relatively frequent in these sectors.

Second, the price of a product usually returns to its original regular price following a sale. For the major groups in Table IV, prices return to their original regular price between 60.0% and 86.3% of the time after a one-period sale. Evidently, many sales price changes are highly transient. Clearance sales are not included in these statistics because a new regular price is not observed after the sale. Yet, clearance sales, like other types of sales, yield highly transient price changes, as we discuss in the supplementary material to this paper.[23]

22. Explanations for sales in the industrial organization literature may be grouped into two categories: (1) intertemporal price discrimination (Varian 1980; Sobel 1984) and (2) inventory management (Lazear 1986; Pashigian 1988; Aguirregabiria 1999). Hosken and Reiffen (2004) use CPI data to evaluate the empirical implications of these models.

23. Our evidence regarding the length of sales and the fraction of price changes that return to the original regular price is limited by the fact that our data set samples prices only once a month. Higher frequency data sets suggest that many sales are substantially shorter than one month (Pesendorfer 2002). This suggests that our estimates of the length of sales are upward biased and that our estimates of the fraction of price changes that return to the original price are downward biased. The fractions reported in Table IV imply that the frequency of regular price change during sales is highly correlated with the frequency of regular price change during nonsale periods and only slightly higher on average. The supplementary material is available both on the QJE website and the personal websites of the authors.

TABLE IV
SALES AND PRICES DURING SALES

	Freq. reg. price ch.	Freq. price ch. during one-period sales	Frac. return after one-period sales	Frac. of sales that last one period	Freq. price ch. during one-period sales/miss.	Av. dur. sales
Processed food	10.5	11.4	78.5	64.7	11.1	2.0
Unprocessed food	25.0	22.5	60.0	63.2	22.1	1.8
Household furnishings	6.0	11.6	78.2	43.3	9.4	2.3
Apparel	3.6	7.1	86.3	35.8	5.9	2.1

Notes. The sample period is 1998–2005. “Freq. reg. price ch.” denotes the median frequency of price changes excluding sales. “Freq. price ch. during one-period sales” denotes the median monthly frequency of regular price change during sales that last one month. The monthly frequency is calculated as $1 - (1 - f)^{0.5}$ where f is the fraction of prices that return to their original level after one-period sales. “Frac. return after one-period sales” denotes the median fraction of prices that return to their original level after one-period sales. “Frac. of sales that last one period” denotes the median fraction of sales that last one month. In calculating this statistic we drop left-censored sale spells. Medians are calculated by first calculating an average within each ELI and then calculating an expenditure-weighted median across ELIs within the major group. “Freq. price ch. during one-period sales/miss.” denotes the median monthly frequency of regular price change during sales or missing periods that last one month, calculated in the manner described above for sales. “Av. dur. sales” denotes the weighted average duration of sale periods in months.

TABLE V
FREQUENCY OF SUBSTITUTION AND PRICE CHANGE BY CATEGORY

			Pr. ch. w/subs		Price change	
Major group	Weight	Subs. freq.	Freq. reg.	Freq.	Freq. reg.	Freq.
Processed food	8.2	1.3	10.9	26.1	10.5	25.9
Unprocessed food	5.9	1.2	25.6	37.2	25.0	37.3
Household furnishing	5.0	5.0	9.2	20.6	6.0	19.4
Apparel	6.5	9.9	7.9	32.2	3.6	31.0
Transportation goods	8.3	10.2	36.6	36.6	31.3	31.3
Recreation goods	3.6	6.3	7.3	14.3	6.0	11.9
Other goods	5.4	1.0	15.4	16.2	15.0	15.5
Utilities	5.3	0.6	38.5	38.5	38.1	38.1
Vehicle fuel	5.1	0.2	87.6	87.6	87.6	87.6
Travel	5.5	1.9	42.5	43.5	41.7	42.8
Services (excl. travel)	38.5	0.9	7.2	7.4	6.1	6.6

Notes. The sample period is 1998–2005. "Subs. freq." gives the median monthly frequency of price changes associated with forced item substitutions in the Consumer Price Index as a fraction of all months in which the product is available, as well as intermediate periods of five months or less when the product is unavailable at the time of sampling but subsequently becomes available. Pr. ch. w/subs denotes the median monthly frequency of price change including price changes due to product substitutes. "Price change" indicates the median monthly frequency of price change. The median statistics are calculated by first calculating the mean frequency of price change or substitutions within ELIs and then calculating the expenditure-weighted median across ELIs. "Weight" denotes the expenditure weight of the ELI. The sector weights add up to 97.4% because used cars are not included in any sector.

In the sections that follow, we document three additional characteristics of sales that differ from regular price changes: (1) sale price changes are more than twice as large as other price changes on average; (2) sales have a very different relationship to aggregate variables such as inflation than regular price changes; and (3) the hazard function of price change including sales is much more downward sloping than the hazard function of price change for regular prices.

III.C. Product Substitutions

The literature on price rigidity has focused primarily on modeling and measuring the frequency of price change for identical items. However, in many durable goods sectors of the economy, the primary mode of price adjustment is not price changes for identical items; it is product turnover. Table V reports information on product substitutions for consumer products. Because product introductions involve pricing decisions, the frequency of product introduction would be the ideal measure of product turnover for the purpose of measuring price flexibility. The CPI Research Database provides an imperfect measure of product introduction by providing an indicator for whether a product undergoes a "forced substitution." A forced substitution occurs if the BLS is forced to stop sampling a product because it becomes permanently unavailable.[24]

The main complication that arises in relating the frequency of substitutions to the frequency of product introduction is that our data set does not follow products over their entire lifetime. Following a substitution, the BLS procedure for choosing a new product to sample tends to lead to the selection of products that have existed for some time.[25] If older products are more likely to become permanently unavailable than new ones, then the average frequency of forced product substitution is an upward-biased measure of the average frequency of product introduction. Despite this caveat, the frequency of substitutions provides useful information on the frequency of product turnover. We measure the frequency of substitutions as a fraction of the total product lifetime.[26]

Substitutions are most common in durable goods categories, particularly apparel and transportation goods. In apparel, we estimate the frequency of substitutions to be 9.9%. Many clothes categories undergo substitutions twice a year at the beginning of the spring and fall seasons. For some clothes, such as women's dresses, substitutions are even more common. Substitutions are also common in transportation goods. In this category, the monthly rate of substitutions is 10.2%. This high rate of substitutions is driven by the introduction of the new model year in cars each fall. Household furnishings and recreation goods also have high rates of substitution, 5.0% and 6.3%, respectively. Other product categories have a rate of substitutions close to 1%.[27]

24. Moulton and Moses (1997) show that price changes that are concurrent with product substitutions play a disproportionate role in explaining steady-state aggregate inflation. This effect is particularly strong in apparel where "clearance sales" are common just before product substitutions.

25. When a product in the data set becomes unavailable, BLS pricing agents are instructed to substitute the most similar available product. In sectors where fashion is important, this is likely to be an older product.

26. We define a product's lifetime as the total time the product is priced and available, where we also include periods when the product is temporarily unavailable for five months or less. This definition is meant to capture the idea that permanent product exits are likely to be followed by new product introductions; but a new product introduction is less likely to occur when the product is only temporarily absent. This measure differs from the measure used by Bils and Klenow (2004). They define the frequency of substitutions as a fraction of the total number of prices collected. We exclude product substitutions that do not lead to a price change because these substitutions do not yield additional price flexibility.

27. See the supplementary material to this paper for a detailed analysis of the timing and frequency of product substitutions in different sectors of the U.S. economy.

In categories such as apparel and transportation goods, the timing of product substitution is primarily motivated by factors such as seasonal demand variation, fashion, and product cycles rather than a firm's desire to change its price. Price changes occur when new products are introduced. But new products are not introduced because the old products were mispriced. This implies that the selection effect associated with price changes due to product substitution may be weaker than for price changes for identical items (Nakamura and Steinsson 2007). The degree of aggregate price flexibility induced by price changes due to product substitution is therefore likely to be less than that induced by the same number of price changes for identical items.

III.D. Frequency of Price Change: Producer Prices

Panel A of Table VI presents statistics on the median frequency of price change for producer prices at three different stages of processing: finished goods, intermediate goods, and crude materials. The median frequency of price change of finished producer goods in 1998–2005 was 10.8%. The corresponding median implied duration is 8.7 months. The median frequency of price change of intermediate goods in 1998–2005 was 13.3%, and the corresponding median implied duration is 7.0 months. In contrast to finished goods and intermediate goods, crude materials have almost completely flexible prices. The median frequency of price change of crude materials in 1998–2005 was 98.9%, and the corresponding median implied duration is 0.2 months. Sales do not appear to be common in our producer price data set.[28] We therefore make no adjustment for sales when analyzing producer prices.

In the PPI, a relatively small (value-weighted) fraction of the categories have a frequency of price change close to the median. Most of the categories with frequencies of price change above the median, have frequencies of price change substantially higher than 10%. As a consequence, the 55th percentile is 18.7% for 1998–2005, while the median is 10.8%. In contrast, for the CPI the 55th percentile is 10.1% for 1998–2005, while the median is 8.7%.

Panel B of Table VI reports results on the frequency of price change of producer prices by two-digit major groups. As in the case of consumer prices, there is a large amount of heterogeneity

28. The PPI database does not include a sales flag. We used the sales filters described in Section III.H to assess the importance of sales in the producer price data. These sales filters identified very few sales.

TABLE VI
FREQUENCY OF PRICE CHANGE FOR PRODUCER PRICES

Category name	Weight	Med. freq. price ch.	Med. freq. substitions	Med. freq. + subs.	Frac. up
A. Stages of processing					
Finished goods	100.0	10.8	1.9	12.1	60.6
Intermediate goods	100.0	13.3	1.2	14.9	58.4
Crude materials	100.0	98.9	4.1	98.9	56.1
B. Major groups (finished-goods weights)					
Farm products	1.6	87.5	0.0	87.5	48.6
Processed foods and feeds	22.4	26.3	2.7	26.6	57.8
Textile products and apparel	3.6	2.3	3.3	3.7	49.7
Hides, skins, leather, and related products	0.3	3.8	1.2	6.4	80.0
Fuels and related products and power	20.8	48.7	0.5	48.7	54.1
Chemicals and allied products	2.8	6.1	6.5	11.3	61.6
Rubber and plastic products	1.8	3.2	1.1	4.0	83.8
Lumber and wood products	0.1	1.3	2.9	4.4	86.6
Pulp, paper, and allied products	3.0	4.4	3.2	9.4	74.9
Metals and metal products	1.1	3.8	3.0	4.6	72.2
Machinery and equipment	13.0	3.7	4.0	4.9	71.0
Furniture and household durables	5.6	5.1	1.1	5.7	78.6
Nonmetallic mineral products	0.1	4.1	1.0	6.1	67.0
Transportation equipment	16.8	27.3	16.6	45.2	53.7
Miscellaneous products	6.9	16.5	0.0	16.5	81.3

Notes. The sample period is 1998–2005. Frequencies are reported in percent per month. Fractions are reported in percentages. "Weight" denotes the post-1997 final goods value weight of the major groups. "Med. freq. price ch." denotes the median frequency of price change. It is calculated by first calculating the mean frequency of price change for each cell code, then taking an unweighted median within the four-digit commodity code, and then taking a value-weighted median across four-digit commodity codes. "Frac. up" denotes the median fraction of price increases. It is calculated in a manner analogous to the median frequency of price change.

across sectors. Table VI also reports the frequency of product substitution for these two-digit major groups. The frequency of product substitution varies across the major groups from 0% in farm products to 16.6% in transportation goods.

The finding that finished-goods producer prices exhibit a substantial degree of rigidity confirms for a broader set of products the results of a number of previous studies (e.g., Blinder et al. [1998]; Carlton [1986]). Interpreting this evidence is, however, more complicated than interpreting evidence on consumer prices. Buyers and sellers often enter into long-term relationships in wholesale markets. It is therefore possible that buyers and sellers enter

TABLE VII
FREQUENCY OF PRICE CHANGE: COMPARISON OF CPI AND PPI CATEGORIES

Category	Number of matches	Frequency CPI w/sales	Frequency CPI nonsale	Frequency PPI	Implied duration CPI w/sales	Implied duration CPI nonsale	Implied duration PPI
Processed food	32	26.1	10.5	7.2	3.3	9.0	13.4
Unprocessed food	24	37.3	25.9	67.9	2.1	3.3	0.9
Household furnishings	27	23.0	6.5	5.6	3.8	14.9	17.3
Apparel	32	31.0	3.6	2.7	2.7	27.3	36.3
Recreation goods	16	14.5	6.8	6.1	6.4	14.2	15.9
Other goods	13	33.6	23.2	17.1	2.4	3.8	5.3

Notes. "Number of matches" denotes the number of ELIs matched to four-, six-, or eight-digit commodity codes within the PPI in the major group. "Frequency" denotes the median frequency of price change. "Implied duration" denotes $-1/\ln(1-f)$, where f is the median frequency of price change. Medians for the consumer price data are calculated by first calculating an average within each ELI and then calculating an expenditure-weighted median across ELIs within the major group. Medians for the producer price data are calculated by first calculating the mean frequency of price change for each cell code, then taking an unweighted median within a four-digit commodity code, and then taking a value-weighted median across four-digit commodity codes. All statistics are for the period 1998–2005.

into long-term "implicit contracts" in which observed transaction prices are essentially installments on a "running tab" that the buyer has with the seller (Barro 1977). In such cases, the buyer would perceive a marginal cost equal to the shadow effect of purchasing the product on the total amount he would eventually pay the seller. But this shadow price would be unobserved. Of course, it is not clear why buyers or sellers would choose to enter into such implicit contracts, or how and why they would then choose to subsequently uphold these contracts. In this type of situation, retail prices might react to changes in the shadow marginal cost even if wholesale prices did not change. Another complication in wholesale markets is that sellers may choose to vary quality margins, such as delivery lags, rather than varying the price (Carlton 1979).

III.E. Frequency of Price Change: CPI vs. PPI

To compare price flexibility at the consumer and producer levels, we matched 153 ELIs from the CPI with product codes from the PPI.[29] Table VII presents comparisons between the frequency of price change at the consumer and producer levels for the major groups in which a substantial number of matches

29. Forty-two ELIs were matched to PPI categories at the eight-digit product-code level, 71 ELIs were matched to PPI categories at the six-digit product-code level, and 40 ELIs were matched to PPI categories at the four-digit product-code level.

were found. In all the major groups except unprocessed food, the median frequency of price change for producer prices is similar to that for consumer prices excluding sales, but substantially lower than the median frequency of price change of consumer prices including sales. For example, for processed food, we find that the median frequency of price change is 7.2% for producer prices, 10.5% for regular consumer prices, and 26.1% for consumer prices including sales. Similarly, for household furnishings, we find that the median frequency of price change is 5.6% for producer prices, 6.5% for regular consumer prices, but 23.0% for consumer prices including sales. For all 153 matches, the correlation between the frequency of price change for producer prices and regular consumer prices is 0.83, while the correlation for producer prices and raw consumer prices is 0.64.

III.F. Relative Frequency of Price Increases and Price Decreases

Most models of price rigidity make the simplifying assumption that price changes occur only in response to aggregate shocks.[30] With even a modest amount of inflation, these models imply that almost all price changes are price increases. Table II shows that this assumption is far from being realistic. The weighted median fraction of regular price changes in consumer prices that are price increases is 64.8%, while the weighted median fraction of price changes including sales that are increases is 57.1%.[31] Table VI shows that the same pattern emerges for producer prices. The fraction of price changes in producer prices that are increases is 60.6%. This result has important implications for the calibration of models of price rigidity. Along with the large average size of price changes, emphasized by Golosov and Lucas (2007), it provides strong evidence for the hypothesis that idiosyncratic shocks are an important driving force for price changes.

III.G. Size of Price Changes

Price adjustment is lumpy not only because prices often remain unchanged for substantial periods of time but also because prices change by large amounts when they do change. Table VIII reports the median absolute size of log changes in consumer prices.

30. Examples include Taylor (1980), Calvo (1983), Caplin and Spulber (1987), Dotsey, King, and Wolman (1999), and Mankiw and Reis (2002). A notable exception is Golosov and Lucas (2007).

31. These statistics are calculated as follows. First, we calculate the fraction of price changes that are increases by ELI. Then, we calculate the weighted median of these statistics across ELI.

TABLE VIII
ABSOLUTE SIZE OF PRICE CHANGES

		Regular prices			Sales			All prices
Major group	Weight	Median change	Median increase	Median decrease	Median change	Median ratio	Frac. price ch.	Median change
Processed food	8.2	13.2	11.5	17.6	33.1	2.6	57.9	26.5
Unprocessed food	5.9	14.2	13.9	15.0	35.1	2.5	37.9	27.1
Household furnishings	5.0	8.7	8.0	9.8	28.0	2.8	66.8	20.8
Apparel	6.5	11.5	10.0	13.3	37.1	3.1	87.1	30.2
Transportation goods	8.3	6.1	5.9	6.2	14.1	0.9	8.0	6.1
Recreation goods	3.6	10.1	8.7	12.0	32.9	3.1	49.1	18.9
Other goods	5.4	7.3	7.2	9.2	26.5	2.9	32.6	10.0
Utilities	5.3	6.3	6.2	6.4	12.6	1.6	0.0	6.3
Vehicle fuel	5.1	6.4	6.8	5.9	11.7	1.8	0.0	6.4
Travel	5.5	21.6	20.9	22.4	29.3	1.4	1.5	21.9
Services (excl. travel)	38.5	7.1	6.5	9.5	29.5	2.9	3.1	7.3
All sectors	100.0	8.5	7.3	10.5	29.5	2.6	21.5	10.7

Notes. The sample period is 1998–2005. “Regular prices” denotes prices excluding sales. “Weight” denotes the CPI expenditure weight of the major group. “Median change,” “Median increase,” and “Median decrease” refer to the weighted median absolute size of log price changes, increases, and decreases, respectively. The median absolute size of log price changes is calculated by first calculating the mean absolute size of log price changes for each ELI and then taking a weighted median across ELIs using CPI expenditure weights. Other median statistics are calculated in an analogous manner. “Median ratio” denotes the weighted median ratio of the mean absolute size of log price changes due to sales to the absolute size of log regular price changes within ELIs. For each ELI the mean size of sales is calculated for all price changes at the beginning and end of sales. “Frac. price ch.” denotes the mean fraction of price changes that are due to sales. The sector weights add up to 97.4% because used cars are not included in any sector.

For consumer prices excluding sales, the median absolute size of price changes is 8.5%.[32] This table also reports the absolute size of price change by major group. The median absolute size of price changes due to sales is 29.5%, more than three times the size of regular prices. The results are similar for finished-goods producer prices. The median absolute size of log changes for finished-goods producer prices is 7.7%. Another result that emerges from Table VIII is that the median size of price decreases is larger than the median size of price increases. For consumer goods, this difference is 3.2 percentage points. For finished-goods producer prices, it is 1 percentage point.

III.H. Alternative Measures of Sales

Up until now we have used the BLS sale flag to identify sales. An alternative approach is to look for V-shaped patterns in the data and identify these patterns as sales. An important conceptual difference between this "sale filter" approach and our previous approach is that clearance sales are not defined as "sales" according to this approach. Sales filters therefore tend to generate somewhat higher estimates of the frequency of price change.

There are two main empirical drawbacks of the sale filter approach to identifying V-shaped sales. First, because prices are observed at a monthly frequency, a simple sale filter that excludes only V-shaped sales would not be able to identify V-shaped sales that are followed by a regular price change within the same month. For example, consider a good that goes on sale for one week, reverts to the original price following the sale, but subsequently experiences a regular price change before the BLS price collector returns to the store. The simple sale filter would not identify this price pattern as a sale, even though the true pattern of prices (unobserved in monthly data) exhibited a V-shaped pattern. Second, in some categories with highly volatile prices, such as gasoline, sale filters may identify sales even when there are none. In these categories, sale filters may identify V-shaped price patterns simply because prices tend to change by discrete amounts, for example, from $2.49 to $2.59. For this reason, sale filters will indicate that gasoline is on sale a significant fraction of the time, whereas the BLS sale flag indicates that there are virtually no sales in the gasoline category.

32. This statistic is calculated by finding the average log change in price by ELI and then taking the weighted median across ELIs.

TABLE IX
FREQUENCY OF PRICE CHANGE FOR SALE FILTERS, 1998–2005

	No subs.	With subs
Sale filter B, 1-month window	15.3	16.4
Sale filter A, 1-month window	13.3	14.7
Sale filter A, 3-month window	11.9	14.1
Sale filter A, 5-month window	11.4	13.3
Price changes	19.4	20.5
Reg. price changes (sale flag)	8.7	10.9
Reg. price ch. + clear	10.7	13.0

Notes. This table gives the weighted median frequency of price change for alternative procedures for filtering out V-shaped sales. Frequencies are reported in percent per month. The median frequency is calculated by first calculating the mean frequency of price change for each ELI and then taking an expenditure-weighted median across ELIs using CPI expenditure weights. In all cases, clearance sales are not removed. Sale filter B removes only symmetric V-shaped sales while Sale filter A also allows for regular price changes immediately preceding or following sales or asymmetric V's. We consider sale filters with a "window" for return to the original price of between one and five months. See supplemental material for this paper for a detailed description of the sale filter algorithm.

The sale filter approach nevertheless provides useful information about both the nature of price adjustment as well as the definition of the "sale flag" variable. Table IX reports results for two types of sale filters, which we refer to as sale filters A and B. Sale filter B removes price patterns in which the price returns to the original price within a set number of months without going above the original price. Sale filter A is designed to also remove price patterns in which a sale is followed by a change in the regular price, that is, asymmetric V's. These procedures are described in detail in the supplementary material for this paper. For each type of filter, we consider different windows between one and five months. For example, for the two-month case, we require that the price return to a regular price in the first two months after the price decline occurs.

The median frequency of price change based on the sale filter A with a window of five months is 11.4% for the 1998–2005 period. This statistic is similar to the weighted median frequency of price change that uses the sale flag to exclude all sales except clearance sales.[33] However, depending on how one parameterizes the sale filter, and depending on whether product substitutions are included as price changes, one can get substantially different answers for the median frequency of price change. In particular,

33. See supplementary material for a discussion of how we identify clearance sales.

if one assumes a window of one month, counts only symmetric V's and includes substitutions as price changes, the frequency of price change rises to 16.4%. For alternative choices of the window and the decision of whether to include substitutions, one can obtain a variety of intermediate values for the median frequency of price change between 11.4% and 16.4%, implying median durations between 5.6 and 8.3 months.[34]

IV. A Benchmark Menu Cost Model

The facts we have established can help distinguish between different models of price-setting behavior. We focus on a benchmark version of the menu cost model developed by Barro (1972), Sheshinski and Weiss (1977), and Golosov and Lucas (2007).

Consider the pricing decision of a single firm. This firm produces a good using a linear technology

$$y_t(z) = A_t(z)L_t(z), \tag{1}$$

where $y_t(z)$ denotes the output of the firm in period t, $A_t(z)$ denotes the productivity of the firm's labor force in period t, and $L_t(z)$ denotes the quantity of labor hired by the firm for production purposes in period t. Assume that demand for the firm's good is

$$c_t(z) = C\left(\frac{p_t(z)}{P_t}\right)^{-\theta}, \tag{2}$$

where $c_t(z)$ denotes the quantity demanded of the firm's good in period t, $p_t(z)$ denotes the nominal price the firm charges in period t, P_t denotes the price level in period t, and C is a constant that determines the size of the market for the firm's good. To generate price rigidity, we assume that the firm must hire an extra K units of labor in order to change its price.

For simplicity, we assume that the real wage rate in the economy is constant and equal to

$$\frac{W_t}{P_t} = \frac{\theta - 1}{\theta}, \tag{3}$$

where W_t denotes nominal wage rate in the economy at time t.[35]

34. Klenow and Kryvtsov (2008) consider a sale filter similar to the one we report in the top right corner of Table IX.

35. In a general equilibrium model with linear disutility of labor and constant aggregate consumption, the real wage would be equal to $W_t/P_t = \alpha U_C(C)$, where α

Using equations (1), (2), and (3) and the fact that markets clear, we can write real profits as

$$\Pi_t(z) = C\left(\frac{p_t(z)}{P_t}\right)^{-\theta}\left(\frac{p_t(z)}{P_t} - \frac{\theta-1}{\theta}\frac{1}{A_t(z)}\right) - \frac{\theta-1}{\theta}KI_t(z). \tag{4}$$

Assume that the logarithm of productivity of the firm's labor force follows an AR(1) process:

$$\log(A_t(z)) = \rho\log(A_{t-1}(z)) + \epsilon_t(z), \tag{5}$$

where $\epsilon_t(z) \sim N(0, \sigma_\epsilon^2)$ is an idiosyncratic productivity shock.

Assume that the logarithm of the price level fluctuates around a trend:

$$\log P_t = \mu + \log P_{t-1} + \eta_t, \tag{6}$$

where $\eta_t \sim N(0, \sigma_\eta^2)$.

The firm maximizes profits discounted at a constant rate $\boldsymbol{\beta}$. The value function of the firm is given by the solution to

$$V(p_{t-1}(z)/P_t, A_t(z)) = \max_{p_t(z)}\left[\Pi_t(z) + \boldsymbol{\beta}E_t V(p_t(z)/P_{t+1}, A_{t+1}(z))\right],$$

where E_t denotes the expectations operator conditional on information known at time t. We solve the firm's problem by value function iteration on a grid. We approximate the processes for $A_t(z)$ and P_t using the method proposed by Tauchen (1986).

The solution to the firm's problem depends on the parameters of the model: $\boldsymbol{\beta}$, θ, K/C, μ, ρ, σ_ϵ, and σ_η. We set the monthly discount factor equal to $\boldsymbol{\beta} = 0.96^{1/12}$. We choose $\theta = 4$ to roughly match estimates from the industrial organization literature on markups of prices over marginal costs.[36] We estimate $\mu = 0.0021$ and $\sigma_\eta = 0.0032$ from data on the CPI from the period 1998–2005. We choose the remaining three parameters to match our estimates of the frequency of regular price change, the fraction of regular price changes that are price increases, and the size of regular

is the marginal disutility of labor. Under the additional assumption that prices are flexible, $W_t/P_t = (\theta-1)/\theta$. More generally, if the degree of monetary nonneutrality is small, variation in C_t will be small and the real wage will be approximately constant.

36. The value of θ we choose implies a markup similar to the mean markup estimated by Berry, Levinsohn, and Pakes (1995) but slightly below the median markup found by Nevo (2001). Broda and Weinstein (2006) report a median elasticity of demand below 3 using trade data. Midrigan (2006) uses $\theta = 3$ while Golosov and Lucas (2007) use $\theta = 7$. Were we to assume $\theta = 10$, our estimate of K/C would rise to 0.07. All other results would be essentially unaffected.

price changes in 1998–2005. The parameter values that imply that the model matches the data along these three dimensions are $K/C = 0.0245$, $\rho = 0.660$, $\sigma_\epsilon = 0.0428$. The model does not generate sale-like behavior for prices. We calibrate the model to match statistics for regular price changes and investigate whether it provides a good positive model of regular price adjustments. The simultaneous existence of rigid regular prices and frequent sales is an important challenge for the theoretical literature on monetary nonneutrality.

We can now test the model calibrated in this way by seeing how well it can account for other empirical features of price change. In the next three sections, we present several new empirical facts about price change and consider how well they line up with the implications of the model presented above.

V. Inflation and the Frequency of Price Change

The frequency of price change is not constant over time. As the rate of inflation varied over the period 1988–2005, the frequency of price change varied systematically along with it. We analyze the evolution of four components of aggregate inflation: the median frequency of price increases, the median frequency of price decreases, the median absolute size of price increases, and the median absolute size of price decreases.[37] Figure II plots the annual evolution of the frequency of price increases and price decreases for consumer prices along with the evolution of CPI inflation. An analogous plot for the size of price increases and decreases is presented in the supplementary material.[38] Of these four components of aggregate inflation, only the frequency of price increases displays a strong relationship with inflation. In contrast, the frequency of price decreases and the size of price increases and price decreases covary much less with inflation.[39]

Table X conveys through regressions what Figure II conveys graphically. We regress the four components at the ELI level on

37. Gagnon (2007) emphasizes the importance of distinguishing between price increases and decreases in this context.

38. As in Section III, these statistics are calculated by first calculating the mean frequency within each ELI and then finding the weighted median across ELIs.

39. This same result has been documented for the Euro area (Vilmunen and Laakkonen 2004; Dhyne et al. 2005). Also, Cecchetti (1986), Lach and Tsiddon (1992), Kashyap (1995), and Goette, Minsch, and Tyran (2005) all find that inflation has a substantial effect on the frequency of price change, but a much weaker effect on the absolute size of price changes.

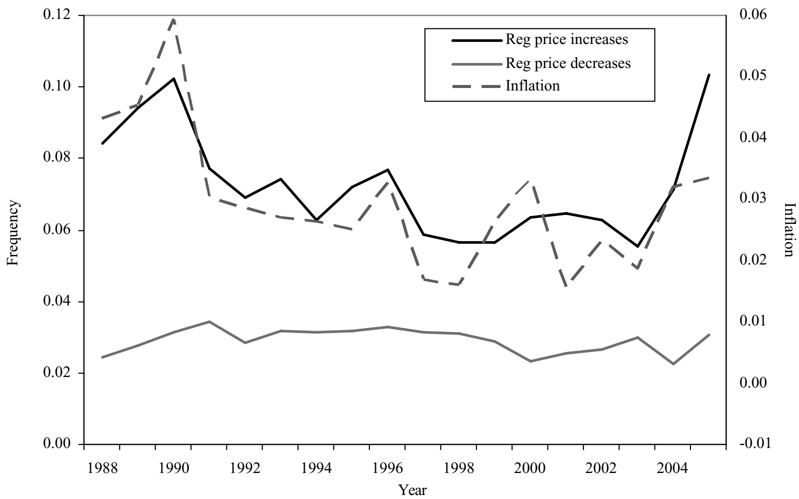

FIGURE II
Inflation and the Frequency of Regular Price Change for Consumer Prices
Note. The figure plots the annual evolution of the weighted median frequency of regular price increases and decreases along with the CPI inflation rate.

the aggregate CPI inflation rate. The regressions include ELI fixed effects and a time trend. We run such regressions both including and excluding sales and separately for 1988–1997 and 1998–2005. The regression coefficient on the frequency of price increases is always positive and statistically significant. The coefficient on price decreases is always negative and statistically significant for regular price decreases. In contrast, the coefficients on the absolute size of price increases and decreases are inconsistent and never significantly different from zero. It is important to note that these results should be interpreted with caution given the small amount of inflation variability over the period we consider.[40]

Figure III compares the evolution of the frequency of price change in the model to its evolution in the data. We simulated the model 100,000 times for the actual evolution of the CPI over 1988–2005 and calculated the average frequency of price increases and decreases by year. Just as in the data, the frequency of price

40. The year 1990 is an outlier in terms of both the frequency of price change and the inflation rate and therefore contributes disproportionately to the statistical significance and magnitude of the regression coefficients for the period 1988–1997. If a dummy for 1990 is included in the regression for the 1988–1997 period, the coefficient falls to 0.68 (0.36) for the frequency of price change and is virtually unchanged at 0.97 (0.18) for the frequency of price increases.

TABLE X
REGRESSIONS OF FREQUENCY AND SIZE OF CONSUMER PRICE CHANGES ON INFLATION

	Regular prices		Prices	
Dependent variable	1988–1997	1998–2005	1988–1997	1998–2005
Consumer price ELI level:				
Frequency of price increase	0.96*	0.56*	0.77*	0.70*
	(0.09)	(0.26)	(0.10)	(0.22)
Frequency of price decrease	−0.22*	−0.36*	−0.22	−0.41
	(0.10)	(0.08)	(0.13)	(0.13)
Size of price increase	0.17	−0.48	−0.06	−0.58
	(0.18)	(0.45)	(0.09)	(0.40)
Size of price decrease	−0.11	−0.43	0.08	0.24
	(0.37)	(0.24)	(0.24)	(0.14)
Frequency of price change	0.74*	0.37	0.56*	0.41
	(0.18)	(0.43)	(0.21)	(0.34)
Size of price change	0.52*	0.49	0.17	0.59
	(0.12)	(0.35)	(0.10)	(0.56)

Notes. The table reports the results of regressions of the mean frequency and absolute size of log price increases and decreases at the ELI level on the aggregate CPI inflation rate (log change over twelve months). For example, the number in the table in the first row of numbers and first column of numbers (i.e., 0.96) refers to the regression coefficient on CPI inflation in a regression where the dependent variable is the frequency of regular price increases in 1988–1997. Each observation is for a particular ELI in a particular year. All regressions include ELI-level fixed effects and ELI-level time trends. Standard errors are in parentheses. The standard errors are cluster-robust standard errors calculated according to the method described by Arellano (1987), where the standard errors are clustered by year. *Significant at 5% level.

increases in the model covaries much more strongly with inflation than the frequency of price decreases and the size of price increases and price decreases. For robustness, we also carry out this exercise in the general equilibrium model presented in Nakamura and Steinsson (2007) and get virtually identical results.

The greater covariance of the frequency of price increases than the frequency of price decreases is a consequence of the fact that the price level is drifting upward. Positive inflation implies that the distribution of relative prices is asymmetric with many more prices bunched toward the lower sS bound than the upper sS bound. The bunching toward the lower sS bound implies that the frequency of price increases covaries more than the frequency of price decreases with shocks to the price level.

The model also matches the fact that the median size of price decreases is larger than that of price increases. Ellingsen, Friberg, and Hassler (2006) show that this asymmetry can arise because the firm's profit function is asymmetric when the elasticity of demand for its product is constant. An alternative explanation for

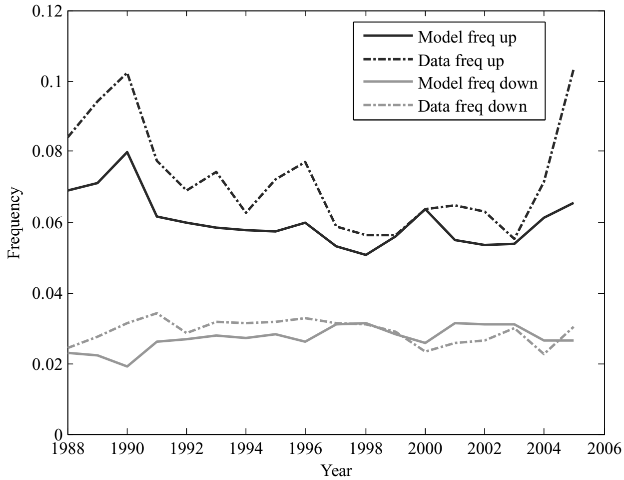

FIGURE III
Frequency of Regular Price Increases and Decreases in the Data and Model

the fact that price decreases are larger than price increases in the data is that we may have failed to filter out all sales.

If new technologies cause the fixed costs of changing prices to fall, the frequency of price change should be increasing over time, other things equal. Figure III shows that for the economy as a whole we do not find evidence of this phenomenon.[41] To the contrary, our menu cost model with a constant menu cost is able to roughly match the evolution of the frequency of regular price change over the period 1988–2005 when we take into account the evolution of inflation.

Klenow and Kryvtsov (2008) find that most of the variation of aggregate inflation stems from variation in the average size of price changes. The average size of price change may be decomposed as $s_{\text{all}} = f_u s_u - f_d s_d$, where f_u and f_d denote the frequency of price increases and price decreases, respectively, and s_u and s_d denote the size of price increases and price decreases, respectively. We find that the frequency of price increases f_u is an important driving force behind variation in the average size of

41. There are two sectors that do not follow this general pattern. These are vehicle fuel and travel services. The frequency of price change for vehicle fuel rose essentially monotonically from approximately 60% in 1988 to approximately 95% in 2005, whereas the frequency of price change for travel services rose again monotonically from approximately 20% in 1988 to 50% in 2005.

price changes. Klenow and Kryvtsov (2008) also find that there is less asymmetry in the relationship between the frequency of price increases versus price decreases and the inflation rate if one looks at the mean frequency of price change across sectors rather than the median frequency of price change. The asymmetry between price increases and decreases is present in virtually all sectors of the U.S. economy in which there is a substantial amount of price rigidity. The difference between means and medians arises because travel and vehicle fuel both have a strong upward trend in the frequency of price change.[42]

The response of producer prices to variation in inflation is similar to the response of consumer prices excluding sales. We regress the frequency of price increases and decreases and the size of price increases and decreases for producer prices on CPI and PPI inflation separately at the four-digit level for the period 1988–2005. The regressions include product fixed effects and a time trend. The frequency of price increases is highly correlated with both inflation rates. The size of price increases is also significantly correlated with both inflation rates. However, the frequency and size of price decreases are not related to inflation in a statistically significant way.

The evolution of sales in consumer prices over the past two decades has been entirely different from the variation in the frequency of regular price changes. Figure IV shows the annual evolution over the period 1988–2005 of the median fraction of price quotes that are sales for the four major groups for which sales are most important. There has been a remarkable increase in the frequency of sales over this period. The frequency of sales increases substantially in all four categories, doubling in both processed food and apparel. The average size of sales has also increased substantially over the sample period in all of the categories except for household furnishings.[43] The increase is most dramatic in processed food, where the size of sales has nearly doubled from about

42. Inflation has fallen over the time period we consider. The upward trend in the frequency of price change in these sectors over a period when inflation has been falling lowers the correlation between the frequency of price increases and inflation, and raises the correlation between the frequency of price decreases and inflation. The Supplementary Material for this paper reports the evolution of the mean and median frequency of price increases and decreases for all sectors of the economy.

43. The size of a sale is measured as the absolute change in prices at the start of a sale (when the sale flag switches from R to S) or at the end of a sale (when the sale flag switches from S to R). Only sales in which prices before or after the sale are observed are included in this calculation. We found no significant

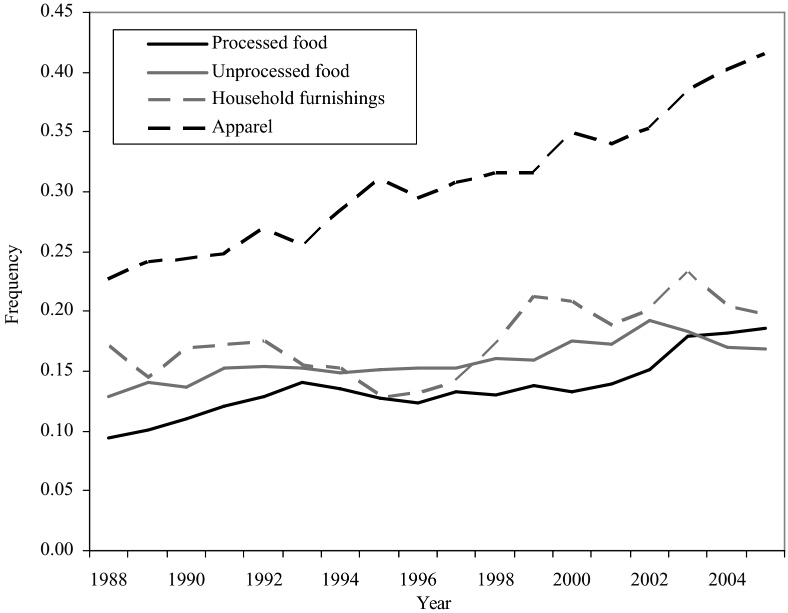

FIGURE IV
Evolution of the Frequency of Sales

Notes. The figure plots the annual evolution of the weighted median across ELIs of the fraction of observations that are sales for the four major groups for which sales are most important.

20% to almost 40%. These facts extend the results of Pashigian (1988), who documents a trend in the frequency and size of sales beginning in the 1960s.

Regressions of the frequency and size of sales on CPI inflation, ELI fixed effects, and a time trend do not find robust evidence of a relationship between either the size or frequency of sales and aggregate variables. For the frequency of sales in 1998–2005, the coefficient on CPI inflation is −0.24 with standard error of 0.20. For the size of sales, the coefficient on CPI inflation is 0.45 with standard error of 0.43. This suggests that a relation may exist between the frequency of sales and inflation, but greater variation in desired prices than is generated by the variation in aggregate inflation over our sample period may be necessary to identify it.

difference between the size of the price decrease at the beginning of sales and the size of the price increase at the end of sales.

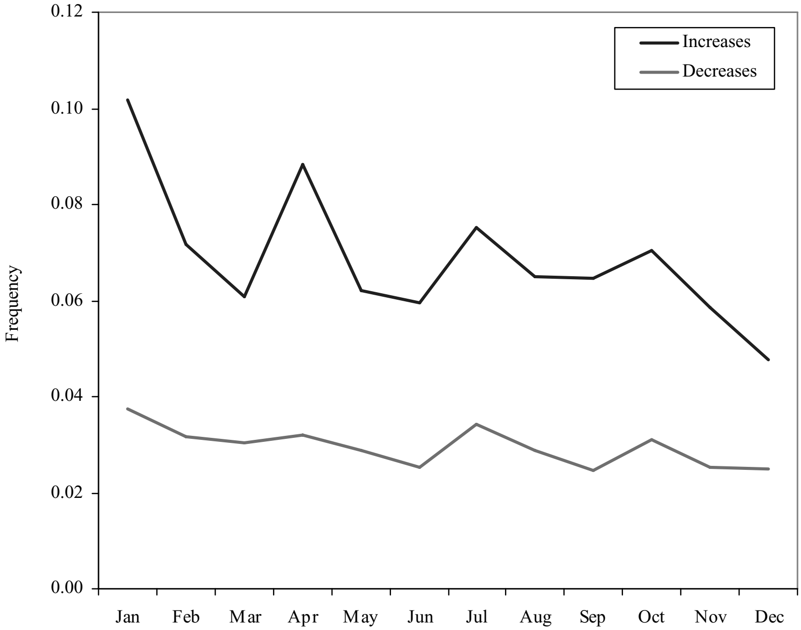

FIGURE V
Frequency of Regular Price Increases and Decreases by Month for Consumer Prices

Note. The figure plots the weighted median frequency of regular price increase and decrease by month.

VI. Seasonality of Price Changes

The synchronization or staggering of price change is an important determinant of the size and persistence of business cycles in models with price rigidity. One form of synchronization of price change is seasonality. We find a substantial seasonal component of price changes for the U.S. economy, for both consumer and producer goods.

Figure V presents the weighted median frequency of price increases and decreases by month for consumer prices excluding sales over the period 1988–2005. Three results emerge. First, the frequency of regular price change declines monotonically over the four quarters. It is 11.1% in the first quarter, 10.0% in the second quarter, 9.8% in the third quarter, and only 8.4% in the fourth quarter. Second, in all four quarters, the frequency of price change is largest in the first month of the quarter and declines monotonically within the quarter. This gives rise to the pattern of local peaks in the frequency of price change in January, April, July, and October. Third, price increases play a disproportionate role in generating seasonality in price changes.[44]

The quarterly seasonal pattern in producer prices mirrors the seasonal patterns in consumer prices qualitatively, but is substantially larger. For producer prices, the frequency of price change is 15.9% in the first quarter, 9.4% in the second quarter, 8.9% in the third quarter, and only 8.2% in the fourth quarter. Most of the seasonality in the frequency of price change in producer prices is due to the fact that producer prices are more than twice as likely to change in January than on average in other months of the year. As in consumer prices, most of the seasonality in the frequency price change comes from the frequency of price increases.

Olivei and Tenreyro (2007) show that the real effects of monetary policy shocks differ depending on the quarter of the year in which the shock hits. They argue that seasonality in the flexibility of wages can explain their empirical findings. Our finding that a disproportionate number of price changes are recorded in January provides an alternative potential explanation for their findings. Of course, seasonality in price setting may simply be evidence that seasonality in wage setting has true allocative effects. Alternatively, seasonality in price setting may suggest some time dependence of price changes.

The seasonal pattern in sales is very different from the seasonal pattern in regular price changes. Figure VI plots the fraction of price quotes that are sales by month for the four major groups for which sales are most important. The major group with by far the most seasonal variation in sales is apparel. The frequency of sales is about 10 percentage points higher in apparel in December, January, and June than in the months with the least sales. Yet, these summer and winter sales are clearly not the only sales in apparel because in the other months, more than 25% of price quotes are sales. We find much less seasonality in sales in other major groups. These patterns have remained roughly unchanged between 1988–1997 and 1998–2005 while the overall level of sales has increased dramatically.

44. Álvarez et al. (2006) find that prices are significantly more likely to change in January in the Euro area.

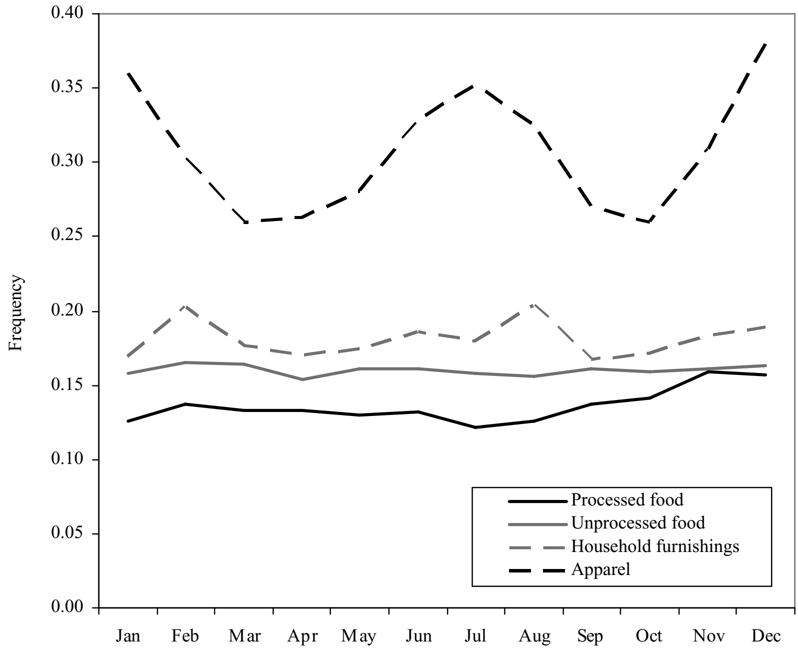

FIGURE VI
Seasonality of the Frequency of Sales

Note. The figure plots the weighted median fraction of observations that are sales by quarter for the four major groups for which sales are most prevalent.

VII. THE HAZARD OF PRICE CHANGE

Are prices that have recently changed more likely than others to change again? Or is it the case that prices become more likely to change the longer they have remained unchanged? These questions are essentially questions about the shape of the hazard function of price change. Let T be a random variable that denotes the duration of a generic price spell. In discrete time, the hazard function is defined as $\lambda(t) = P(T = t \mid T \geq t)$. In other words, the hazard of a price change at time t is the probability that the price will change after t periods given that it has survived for t periods. If prices become more likely to change the longer they have remained unchanged, the hazard function of price change is upward sloping.

Menu cost models can give rise to a multitude of different shapes for the hazard function of price change. If marginal costs follow a random walk, the hazard function will be upward sloping.

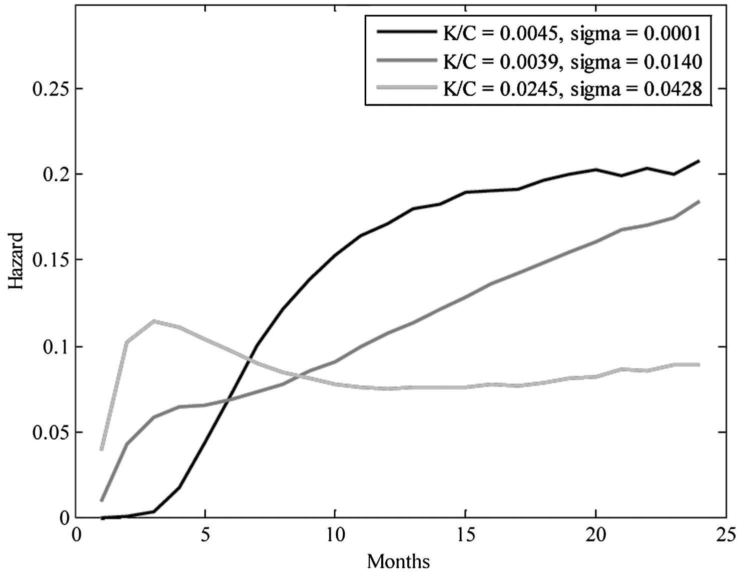

FIGURE VII
Hazard Function in the Menu Cost Model

Notes. Hazard functions with different levels of volatility of the idiosyncratic shock. In all cases $\rho = 0.66$ and the frequency of price change is 8.7%.

More generally, the shape of the hazard function is influenced by the relative size of transient and permanent shocks to marginal costs. Nonstationarity in marginal costs—for example, due to inflation—tends to yield an upward-sloping hazard function, while transient shocks tend to flatten the hazard function and can even yield a downward-sloping hazard. Figure VII illustrates how the shape of the hazard function in our benchmark menu cost model is affected by idiosyncratic shocks to marginal costs. As the variance of idiosyncratic shocks rises relative to the rate of inflation, the hazard function flattens out at longer durations but remains steeply upward sloping in the first few months.[45] In contrast, the Calvo model assumes a flat hazard function of price change.

45. The reason why idiosyncratic shocks flatten the hazard function is that they give rise to temporary price changes that are quickly reversed. Such price changes occur when the idiosyncratic shock is large enough that it is worthwhile for the firm to change its price temporarily to an "abnormal" level even though it realizes that it will soon have to change it back. For calibrations of the model with very large idiosyncratic shocks—much too large to be realistic—the model even generates a downward-sloping hazard in the first few months.

We estimate the hazard function of price change for consumer and producer prices and investigate how it lines up with the implications of our calibrated menu cost model. The main empirical challenge we face in doing this is to account for heterogeneity across products. It is well known in the literature on duration models that estimates of hazard functions based on pooled data from many heterogeneous products lead to a downward bias in the estimated slope of the hazard function (e.g., Kiefer [1988]). We account for heterogeneity in two ways. First, we divide the products in our data set into groups and estimate hazard functions separately for each group. Second, within each group we estimate the empirical model proposed by Lancaster (1979) and analyzed in detail by Meyer (1986, 1990). This model allows for multiplicative unobserved heterogeneity in the level of the hazard function at the product level, while estimating the slope of the hazard function nonparametrically.[46] Specifically, we assume that the hazard function is

$$\lambda_i(t \mid \boldsymbol{x_{i,j}}) = \nu_i \lambda_0(t) \exp(\boldsymbol{x_{i,j}\beta}), \tag{7}$$

where i indexes products, j indexes observations, ν_i is a product-specific random variable that reflects unobserved heterogeneity in the level of the hazard, $\lambda_0(t)$ is a nonparametric baseline hazard function with dummies for each month, $\boldsymbol{x_{i,j}}$ is a vector of covariates for the jth observation of products i, and $\boldsymbol{\beta}$ is a vector of parameters. We assume that $\nu_i \sim \text{Gamma}(1, \sigma_\nu^2)$.[47] An important advantage of our data is that we observe multiple price spells for the same product. This fact substantially enhances our ability to identify the distribution of ν_i. We estimate the model by maximum likelihood. We truncate the price spells at eighteen months and drop left-censored spells.[48]

We divide the data set into groups at the level of major groups. Figure VIII plots the baseline hazard function from the model described by equation (7) for processed food and services. Each panel

46. An example of a "product" is 16-oz Kraft Singles sold at a particular supermarket in New York.

47. We have estimated the model with $\nu_i \sim N(1, \sigma_\nu^2)$. The results are virtually identical.

48. In the presence of heterogeneity, discarding left-censored spells leads us to disproportionately drop price spells arising from subjects with low values of ν_i, because long spells are disproportionately censored (Heckman and Singer 1986). This does not bias our results about the shape of the hazard function under the proportional hazards assumption, though it does affect the estimated level of the hazard function.

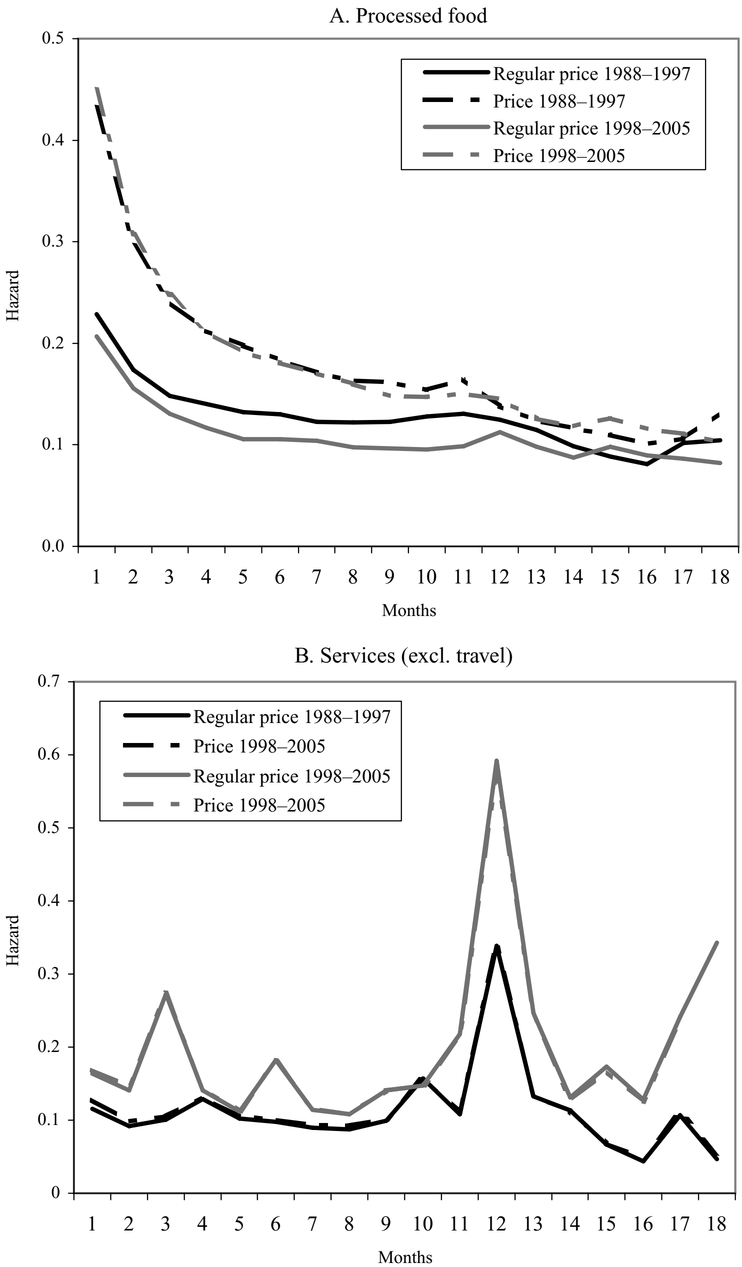

FIGURE VIII
Hazard Function for Consumer Prices

plots the hazard function separately for prices with and without sales and separately for 1988–1997 and 1998–2005. The shape of the hazard function for processed food is representative of the shape of the hazard function for many of the major groups. The hazard function of regular prices is somewhat downward sloping for the first few months and then mostly flat after that. We do not find any evidence of upward-sloping hazard functions.[49] For the major groups in which sales occur frequently, the hazard function including sales is much more steeply downward sloping than the hazard function of regular prices. For services, we estimate a large spike in the hazard function at twelve months. This spike is perhaps most naturally interpreted as an element of time dependence in firms' pricing decisions but may alternatively arise because of seasonality in costs or demand. Interestingly, such a twelve-month spike is completely absent in most other major groups.[50]

For producer prices, we estimate the model described by equation (7) separately for the 15 two-digit major groups. The main stylized facts about the shape of the hazard function for producer prices are similar to those for consumer prices. The hazard functions are downward sloping for the first few months, then mostly flat except for a large twelve-month spike in all major groups. Accounting for heterogeneity leads to a substantial flattening of the hazard functions and a large increase in the size of the spike at twelve months. Interestingly, the twelve-month spike in the hazard function is a much more pervasive phenomenon in producer prices than in consumer prices.

The main difference between the hazard function generated by our benchmark menu cost model and the hazard functions we estimate from the data is the behavior of the hazard in the first few months. In the data the hazard is large and falling whereas

49. The supplementary material for this paper reports plots of the hazard function of eight major groups for consumer prices and another eight major groups for producer prices. The main qualitative features of our results hold even when we estimate our hazard model separately at the ELI level or when we sort products in each major group by their frequency of price change into eight subgroups. We do not report the standard errors of our estimates in Figure VIII because the standard errors are very small.

50. Klenow and Kryvtsov (2008) present hazard function estimates for this same data set. Their estimates are based on a linear probability model with fixed effects. Because of the incidental parameters problem, this estimator yields biased estimates of the shape of the hazard function. Because only a handful of price spells are observed for each product, this bias is potentially quite large. Also, Klenow and Kryvtsov (2008) assume that the shape of the hazard function is the same for all products in the economy. We estimate separate hazard functions for each major group and find large differences across groups.

in the model it is small and rising sharply. We have considered an extension of our benchmark model with heteroscedastic shocks to marginal costs. This model can generate a downward-sloping hazard function in the first few months.

VIII. Conclusion

In this paper, we present new evidence on price adjustment in the U.S. economy. Using BLS microdata we document that the median frequency of nonsale price change is 9%–12% per month, roughly half of what it is including sales. This implies an uncensored median duration of regular prices of 8–11 months. Product turnover plays an important role in truncating price spells in durable goods. The median frequency of nonsale price change including product substitutions is 11%–13%, implying a median duration of 7–9 months. The median frequency of price change for finished-goods producer prices is roughly 11% per month. The frequency of price increases covaries strongly with inflation, whereas the frequency of price decreases and the size of price increases and price decreases do not. We find that the frequency of price change is highly seasonal. Finally, we estimate the hazard function of price changes to be somewhat downward sloping for the first few months and then flat. We argue that the empirical differences between regular price changes, temporary sales, and price changes due to product substitution make it crucial to distinguish between these different classes of price adjustments in macroeconomic models.

Columbia University, NBER, CEPR
Columbia University, NBER, CEPR

References

Aguirregabiria, Victor, "The Dynamics of Markups and Inventories in Retail Firms," *Review of Economic Studies*, 66 (1999), 275–308.

Álvarez, Luis J., Pablo Burriel, and Ignacio Hernando, "Do Decreasing Hazard Functions for Price Changes Make Sense?" Working Paper No. 461, European Central Bank, 2005.

Álvarez, Luis J., Emmanuel Dhyne, Marco M. Hoeberichts, Claudia Kwapil, Harvé Le Bihan, Patrick Lunnemann, Fernando Martins, Roberto Sabbatini, Harald Stahl, Philip Vermeulen, and Jouko Vilmunen, "Sticky Prices in the Euro Area: A Summary of New Micro Evidence," *Journal of the European Economic Association*, 4 (2006), 575–584.

Arellano, Manvel, "Computing Robust Standard Errors for Within-Groups Estimators," *Oxford Bulletin of Economics and Statistics*, 49 (1987), 431–434.

Baharad, Eyal, and Benjamin Eden, "Price Rigidity and Price Dispersion: Evidence from Micro Data," *Review of Economic Dynamics*, 7 (2004), 613–641.
Barro, Robert J., "A Theory of Monopolistic Price Adjustment," *Review of Economic Studies*, 39 (1972), 17–26.
——, "Long Term Contracting, Sticky Prices and Monetary Policy," *Journal of Monetary Economics*, 3 (1977), 305–316.
Baumgartner, Josef, Ernst Glatzer, Fabio Rumler, and Alfred Stiglbauer, "How Frequently Do Consumer Prices Change in Austria?" Working Paper No. 523, European Central Bank, 2005.
Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 63 (1995), 841–890.
Bils, Mark, and Peter J. Klenow, "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112 (2004), 947–985.
Blinder, Alan S., Elie R. D. Canetti, David E. Lebow, and Jeremy B. Rudd, *Asking about Prices* (New York, NY: Russell Sage Foundation, 1998).
Broda, Christian, and David E. Weinstein, "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121 (2006), 541–585.
Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12 (1983), 383–398.
Campbell, Jeffrey R., and Benjamin Eden, "Rigid Prices: Evidence from U.S. Scanner Data," Working Paper, Federal Reserve Bank of Chicago, 2005.
Caplin, Andrew, and Daniel Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, 102 (1987), 703–725.
Carlton, Dennis W., "Contracts, Price Rigidity and Market Equilibrium," *Journal of Political Economy*, 87 (1979), 1034–1062.
——, "The Rigidity of Prices," *American Economic Review*, 76 (1986), 637–658.
Carvalho, Carlos, "Heterogeneity in Price Stickiness and the New Keynesian Phillips Curve," *B.E. Journals in Macroeconomics: Frontiers of Macroeconomics*, 2 (2006), 1–56.
Cecchetti, Stephen G., "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines," *Journal of Econometrics*, 31 (1986), 255–274.
Dhyne, Emmanuel, Luis J. Álvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lunnemann, Fabio Rumler, and Jouko Vilmunen, "Price Setting in the Euro Area: Some Stylized Facts from Individual Consumer Price Data," ECB Working Paper No. 524, 2005.
——, "Price Setting in the Euro Area and the United States: Some Facts from Individual Consumer Price Data," *Journal of Economic Perspectives*, 20 (2006), 171–192.
Dias, Daniel A., Carlos Robalo Marques, and Joao M. Santo Silva, "Time or State Dependent Price Setting Rules? Evidence from Micro Data," *European Economic Review*, 51 (2007), 1589–1613.
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.
Ellingsen, Tore, Richard Friberg, and John Hassler, "Menu Costs and Asymmetric Price Adjustment," CEPR Discussion Paper No. 5749, 2006.
Fabiani, Silvia, Martine Druant, Ignacio Hernando, Claudia Kwapil, Benitta Landau, Claire Loupias, Fernando Martins, Thomas Matha, Roberto Sabbatini, Harald Stahl, and Ad Stokman, "The Pricing Behavior of Firms in the Euro Area: New Survey Evidence," Paper Presented at Conference on Inflation Persistence in the Euro Area at the European Central Bank, 2004.
Fougére, Denis, Hervé Le Bihan, and Patric Sevestre, "Heterogeneity in Consumer Price Stickiness: A Microeconometric Investigation," *Journal of Business and Economic Statistics*, 25 (2005), 247–264.
Gagnon, Etienne, "Price Setting under Low and High Inflation: Evidence from Mexico," International Finance Division Paper No. 896, Federal Reserve Board, 2007.
Goette, Lorenz, Rudolf Minsch, and Jean-Robert Tyran, "Micro Evidence on the Adjustment of Sticky-Price Goods: It's How Often, Not How Much," Discussion Paper, University of Copenhagen, 2005.

Golosov, Mikhail, and Robert E. Lucas, "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115 (2007), 171–199.
Gopinath, Gita, and Roberto Rigobon, "Sticky Borders," *Quarterly Journal of Economics*, 123 (2008), 531–575.
Heckman, James J., and Burton Singer, "Econometric Analysis of Longitudinal Data," in *Handbook of Econometrics*, Vol. III, Z. Grilliches and M. D. Intrilligator, eds. (Amsterdam: Elsevier Science Publishers, 1986, pp. 1689–1763).
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.
Hosken, Daniel, and David Reiffen, "Patterns of Retail Price Variation," *RAND Journal of Economics*, 35 (2004), 128–146.
——, "Pricing Behavior of Multiproduct Retailers," *B.E. Journal of Theoretical Economics: Topics in Theoretical Economics*, 7 (2007), Article 39.
Jonker, Nicole, Carsten Folkertsma, and Harry Blijenberg, "An Empirical Analysis of Price Setting Behavior in the Netherlands in the Period 1998–2003 Using Micro Data," Working Paper No. 413, European Central Bank, 2004.
Kackmeister, Alan, "Yesterday's Bad Times Are Today's Good Old Times: Retail Price Changes Are More Frequent Today Than in 1890s," *Journal of Money, Credit and Banking*, 39 (2007), 1987–2020.
Kashyap, Anil K., "Sticky Prices: New Evidence from Retail Catalogs," *Quarterly Journal of Economics*, 110 (1995), 245–274.
Kehoe, Patrick, and Virgiliu Midrigan, "Sales, Clustering of Price Changes, and the Real Effects of Monetary Policy," Working Paper, University of Minnesota, 2007.
Kiefer, Nicholas M., "Economic Duration Data and Hazard Functions," *Journal of Economic Literature*, 26 (1988), 646–679.
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.
——, "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics*, 123 (2008), 863–904.
Konieczny, Jerzy D., and Andrzej Skrzypacz, "Inflation and Price Setting in a Natural Experiment," *Journal of Monetary Economics*, 52 (2005), 621–632.
Lach, Saul, and Daniel Tsiddon, "The Behavior of Prices and Inflation: An Empirical Analysis of Disaggregated Price Data," *Journal of Political Economy*, 100 (1992), 349–389.
Lancaster, Tony, "Econometric Methods for the Duration of Unemployment," *Econometrica*, 47 (1979), 939–956.
Lazear, Edward P., "Retail Pricing and Clearance Sales," *American Economic Review*, 76 (1986), 14–32.
Mankiw, N. Gregory, 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.
Meyer, Bruce D., "Semiparametric Estimates of Hazard Models," Mimeo, MIT, 1986.
——, "Unemployment Insurance and Unemployment Spells," *Econometrica*, 58 (1990), 757–782.
Midrigan, Virgiliu, "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations," Working Paper, Ohio State University, 2006.
Moulton, Brent R., and Karin E. Moses, "Addressing the Quality Change Issue in the Consumer Price Index," *Brookings Papers on Economic Activity*, (1997), 305–349.
Nakamura, Emi, and Jón Steinsson, "Monetary-Non-Neutrality in a Multi-Sector Menu Cost Model," Working Paper, Columbia University, 2007.
Nevo, Aviv, "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69 (2001), 307–342.
Olivei, Giovanni, and Silvana Tenreyro, "The Timing of Monetary Policy Shocks," *American Economic Review*, 97 (2007), 636–663.
Pashigian, B. Peter, "Demand Uncertainty and Sales: A Study of Fashion and Markdown Pricing," *American Economic Review*, 78 (1988), 936–953.

Pesendorfer, Martin, "Retail Sales: A Study of Pricing Behavior in Supermarkets," *Journal of Business*, 75 (2002), 33–66.
Sheshinski, Eytan, and Yoram Weiss, "Inflation and Costs of Price Adjustment," *Review of Economic Studies*, 44 (1977), 287–303.
Sobel, Joel, "The Timing of Sales," *Review of Economic Studies*, 51 (1984), 353–368.
Stigler, George J., and James K. Kindahl, *The Behavior of Industrial Prices* (New York, NY: Columbia University Press, 1970).
Tauchen, George, "Finite State Markov-Chain Approximation to Univariate and Vector Autoregressions," *Economics Letters*, 20 (1986), 177–181.
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* (Washington, DC: U.S. Government Printing Office, 1997).
Varian, Hal R., "A Model of Sales," *American Economic Review*, 70 (1980), 651–659.
Vermeulen, Philip, Daniel Dias, Maarten Dossche, Erwan Gautier, Ignacio Hernando, Roberto Sabbatini, and Harald Stahl, "Price Setting in the Euro Area: Some Stylised Facts from Individual Producer Price Data and Producer Surveys," ECB Working Paper, 2006.
Vilmunen, Jouko, and Helina Laakkonen, "How Often Do Prices Change in Finland? Micro-Level Evidence from the CPI," Working Paper, Bank of Finland, 2004.
Willis, Jonathan L., "Magazine Prices Revisited," *Journal of Applied Econometrics*, 21 (2006), 337–344.