

Five Facts About Prices: A Reevaluation of Menu Cost Models

Emi Nakamura and Jón Steinsson*

Harvard University

May 5, 2007

Abstract

We establish five facts about prices in the U.S. economy: 1) The median frequency of non-sale 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 price change for finished goods producer prices is roughly 11% per month. 2) One-third of regular price changes are price decreases. 3) The frequency of price increases covaries strongly with inflation while 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 1st quarter and lowest in the 4th quarter. 5) The hazard function of price changes for individual consumer and producer goods is downward sloping for the first few months and then flat (except for a large spike at 12 months in consumer services and all producer prices). These facts are based on CPI microdata and a new comprehensive data set of microdata on producer prices that we construct from raw production files underlying the PPI. We show that the 1st, 2nd and 3rd facts are consistent with a benchmark menu-cost model, while the 4th and 5th facts are not.

Keywords: Price Rigidity, Hazard Functions, Menu Cost Models.

JEL Classification: E30

*We would like to thank Robert Barro for invaluable advice and encouragement. We would like to thank Daniel Benjamin, David Berger, Leon Berkemans, Craig Brown, Charles Carlstrom, Gary Chamberlain, Tim Erickson, Mark Gertler, Mike Golosov, Gita Gopinath, 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 and seminar participants at Harvard, the Federal Reserve Board and the Federal Reserve Bank of New York for helpful comments and discussions. We particularly want to thank Mark Bilal 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.

1 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. For this reason, macroeconomists have had a persistent interest in micro-level empirical evidence about the behavior of prices. We use 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 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 non-shelter consumer prices in 1995-1997 was 21%, implying a median duration of 4.3 months.

We use a substantially more detailed dataset than Bils and Klenow (2004) that contains the micro-level price data underlying the non-shelter component of the consumer price index.¹ This dataset has been used by Klenow and Kryvtsov (2005) 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 non-shelter consumer expenditures. While the median frequency of price change including sales is 19-21% per month, we find that the median frequency of non-sale price change is only 9-12% per month depending on the time period and how we treat non-sale price changes over the course of sales and stockouts. These frequency estimates imply an uncensored duration of regular prices of between 8 and 11 months. If we include price changes associated with product substitutions, the frequency of price change increases by between 1 and 2 percentage points.

The importance of temporary sales in generating price flexibility draws attention to the question of how sales should be viewed when thinking about the macroeconomic implications of price

¹Bils 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 2 for a more detailed discussion of the data.

rigidity.² Are the macroeconomic implications of sales the same as those of other price changes? Or are they different? The theoretical literature on price adjustment clearly indicates that different types of price adjustments can have very different macroeconomic 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.

We document several important empirical differences between sale price changes and other types of price adjustments. The most important differences are: 1) Sale price changes appear to be much more transient than other types of price changes. In most cases where a price is observed before and after a sale, the price returns to its original level following the sale. 2) Sales price changes are more than twice as large as other price changes on average. 3) The frequency and size of sales have a very different relationship to aggregate variables than regular price changes. 4) The hazard function of price change including sales is very different from that excluding sales. This last feature is a direct consequence of the transient nature of sales—the hazard function of price change including sales reflects the presence of a greater number of short price spells.

There are a number of reasons why it may be important to distinguish between sale and non-sale 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.

Our frequency of price change measures are for identical items. Product turnover is another source of price flexibility that plays an important role in truncating price spells, particularly in durable goods categories such as automobiles and apparel. The median frequency of product substitution over the period 1998-2005 in transportation goods and apparel was 10.2% and 9.9% per month, respectively. As with sales, it is important to distinguish between price changes due to product turnover and price changes for identical items. Many factors other than a firm's desire to change its price influence its decision to introduce a new product. The theoretical literature suggests that it is crucial to distinguish between price adjustments that are motivated primarily by a large difference between a firm's current price and its desired price and those that are motivated

²A number of recent theoretical studies base their analysis on statistics for data where sales have been excluded (e.g., Golosov and Lucas, 2006; and Midrigan, 2005).

by other factors.³

We also present the first broad-based evidence on U.S. price dynamics at the producer level. Price rigidity at the producer level is potentially important because even if retail prices are perfectly flexible, price rigidity of producer prices could imply that shocks to production costs are not immediately passed through to consumer prices. In order 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. 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. Price rigidity in finished goods producer prices thus seems comparable to the rigidity of consumer prices excluding sales but substantially more 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 due to the fact that 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 degree of monetary non-neutrality in an economy with heterogeneous price rigidity. In Nakamura and Steinsson (2006), we calibrate a multi-sector menu cost model to the sectoral distribution of the frequency and absolute size of price changes excluding sales. The degree of monetary non-neutrality implied by this multi-sector 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. Bils and Klenow (2002) and Carvalho (2006) study the effect of heterogeneous price rigidity in time-dependent models.

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

³In Nakamura and Steinsson (2006), we analyze a menu cost model in which product introduction represents a random opportunity to set a new price. We show that increasing the frequency of product introduction lowers monetary non-neutrality by about 5 times less than an equal increase in the frequency of regular price changes.

sales and finished goods producer prices. We present a benchmark menu cost model along the lines of Golosov and Lucas (2006) and show that the fraction of price changes that are decreases helps pin down the key parameters of this model. Building on the insights in Golosov and Lucas (2006), we find that the combination of the fact that 1/3 of price changes are price decreases and the fact that the average absolute size of price changes is large favors a model in which large but relatively transient idiosyncratic shocks to firms are an important driving force behind most price changes.

The third feature of price change that we investigate is how the frequency and size of price changes covary with variations in the inflation rate. We find that the frequency of price increases covaries quite strongly with the rate of inflation, while 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. We find that the model matches the data well along this dimension. The frequency of price increases covaries much more with inflation than the other three components in the model as in the data.

The fourth feature of price change that we investigate is the extent of seasonal synchronization of price changes. 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 difference is particularly large for producer prices. For consumer prices, we furthermore find a consistent pattern within quarter. The frequency of price change is highest in the first month of each quarter and falls monotonically across months within the quarter. This feature of price change does not arise in our benchmark menu cost model. It could arise in a menu cost model in which firms face seasonal variation in marginal costs or demand; or it may be evidence of a time-dependent element of the pricing decisions of firms.

The fifth and final issue that we investigate is the hazard function of price change. We are primarily interested in the slope of the hazard function. Menu cost models can give rise to a wide variety of hazard functions depending on the specification of marginal costs. The hazard function implied by our calibrated benchmark menu cost model is sharply upward sloping for the first few months. This implies that prices are unlikely to change again in the month immediately following a 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 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 12 months for services and producer prices. We also estimate the hazard of price changes for consumer prices including sales. It is much more sharply downward sloping for categories with frequent sales.⁴

An important body of work on the nature of price adjustment in the European context has been carried out by the Inflation Persistence Network (IPN) of the European Central Bank. Álvarez et al. (2005b) 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) summarizes analogous studies on producer prices in the Euro Area. Fabiani et al. (2004) summarizes the conclusions of a set of papers that analyze survey evidence on price adjustment in the Euro Area. A number of other recent papers have studied the size and frequency of price changes using disaggregated price data, including Lach and Tsiddon (1992), Konieczny and Skrzypacz (2005), Baharad and Eden (2004), Kackmeister (2005), Gopinath and Rigobon (2006), Hobijn et al. (2006) and Midrigan (2005). Hosken and Reiffen (2004) use CPI data to analyze the ability of industrial organization models to explain the sales observed in consumer prices, concluding that none of the existing models are particularly successful.

The paper is organized as follows. In section 2, we describe the data. In section 3, 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 4, we present and calibrate a benchmark menu cost model. In section 5, we present evidence on how the frequency and size of price changes vary with inflation. In section 6, we present evidence on the seasonality of price changes and sales. In section 7, we present our estimates of the hazard function of price change. Section 8 concludes.

⁴Earlier empirical work on the hazard function of price changes includes Cecchetti (1986), Campbell and Eden (2004), Baumgartner et al. (2005), Álvarez et al. (2005a), Jenker et al. (2004), Dias et al. (2005), Fougere et al. (2005), Goette et al. (2005) and Gagnon (2005). The evidence on the shape of the hazard function from these papers is mixed. Most of the papers in this literature do not account for unobserved heterogeneity at the good level. However, several of them use the conditional logit specification to account for unobserved heterogeneity. Unfortunately, this specification yields inconsistent estimates of the shape of the hazard function, as discussed in Willis (2006).

2 The Data

We use two data sets gathered by the Bureau of Labor Statistics (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 Consumer Price Index (CPI). The second is an analogous data set of producer prices that we have created from the production files underlying the Producer Price Index (PPI). We will refer to this data set as the PPI Research Database. The CPI Research Database has been used by Klenow and Kryvtsov (2005).⁵ The PPI Research Database has not been used before.

2.1 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 non-shelter 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.

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 are considered different products in the database. An example of a product in the database is a 2 liter bottle of Diet Coke sold at a particular supermarket in New York. The database reports whether or not a product was “on sale” when its price was sampled in a particular month.⁶ We use this sales flag to calculate statistics

⁵Bils 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.

⁶BLS field agents are instructed to mark a price as a sale price if it is considered by the outlet to be lower than the regular selling price, temporarily, 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—i.e. 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 is based on an index of used car

about the frequency and size of price change excluding sales. We also consider identifying sales based on a “sales filter” in section 3.8. 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.⁷

We present results for consumer prices at three levels of aggregation. First, we report statistics that are calculated using the entire cross section of goods. Second, we break the data set into 11 “Major Groups” (see table 3). Third, we report results for so called Entry Level Items (ELIs). Examples of ELIs are “Bread”, “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 must report separate estimates 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 across ELIs. The weights we use are CPI expenditure weights from 1990 for the period 1988-1997 and from 2000 for the period 1998-2005. The statistics at the ELI level are unweighted averages within the ELI.

2.2 The PPI Research Database

The PPI Research Database contains an unbalanced panel of raw data from the productions files used to construct the PPI. The earliest prices in the database are from the late 1970’s. For most categories, however, the sample period begins some time during the early to mid 1980’s. Throughout the sample period a number of categories are discontinued and others appear. To a large extent this “churning” reflects, on the one hand, ongoing modernization of the data set, and on the other hand, the expansion of the data set into new sectors. 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.⁸

prices. The data on airline tickets is 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.

⁷Chapter 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.

⁸The weights referred to here are the post-1997 value weights used to construct the Finished Goods PPI.

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 actually observe 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.⁹ It is important to note that many of the transactions for which prices are collected as part of the PPI are a part of implicit or explicit long-term contracts between firms and their suppliers. The presence of such long-term contracts makes interpreting the PPI data more complicated than interpreting CPI data as we discuss further in section 3.4.

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”. In its Handbook of Methods, the BLS says: “For example, if a company charges more for a red widget than a white one, color is one of the price-determining variables.” Price-determining variables may include the buyer, the quantity being bought, the method of shipment, the transactions terms, the day of the month on which the transaction takes place as well as 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.

The price data in the PPI are collected in two steps. When a product is first introduced into the dataset or when an industry is resampled, 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 may involve 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 to the respondent. If the form is not returned, a BLS Industry Analyst will call the firm

⁹See Chapter 14 of the BLS Handbook of Methods (U.S. Department of Labor, 1997) for a more detailed description of BLS procedures.

or establishment 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. In order 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.¹⁰ 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 other years. 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 manufacturer prices and consumer prices excluding sales documented in section 3.5.

The BLS publishes producer price indexes based on several different classification systems. The most important classification systems are industry classifications and stage of processing classifications. Indexes in all classification systems are based on the same pool of price information. We use the stage of processing classification system. This classification system groups goods according to the class of buyer and the amount of physical processing or assembly the products have undergone. The BLS constructs indexes for three different stages of processing: finished goods, intermediate goods and crude material. We focus attention on finished goods, but also report basic results for intermediate goods and crude materials. As with the consumer price data, we present results at three different levels of aggregation. First, we present results that are based on the entire cross section of goods. Second, we present results for 15 Major Groups. These Major Groups are the two digit stage of processing groupings used by the BLS. Third, we present results based on a matching between more disaggregated stage of processing groupings and CPI ELIs.

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, since there is a substantial

¹⁰We thank Roberto Rigobon for suggesting this robustness check.

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 4-digit commodity code level. We then construct statistics on the frequency of price change at the 4-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 4-digit commodity code. Finally, we construct aggregate statistics by taking value weighted medians over the median price change frequencies at the 4-digit commodity code level. For the purpose of matching PPI categories with CPI ELIs, we also construct statistics at the 6-digit and 8-digit level. These statistics are unweighted medians analogous to the statistics we calculate at the 4-digit level.

3 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. An important lesson from the theoretical literature on price adjustment is that different types of price adjustments can have very 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, 2006). In contrast, if the timing of most 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 and present statistics on these different types of price changes separately. We distinguish between price changes for identical items and price changes that occur due to product turnover. We also distinguish between changes in the regular price of a product and price changes associated with temporary sales.

The literature on price rigidity has focused primarily on modeling and measuring the frequency of price change for identical items. However, in sectors such as apparel, the primary mode of price adjustment is not price changes for identical products; it is product turnover. While the notion that prices adjust when they are far from the firm’s desired price is intuitive in the case of price changes for identical items, the timing of product turnover in sectors such as apparel is primarily

motivated by other factors such as seasonal demand variation and fashion. This implies that the selection effect associated with price changes due to product turnover may be weaker than for price changes for identical items. Price changes due to product turnover would then imply less aggregate price flexibility than the same number of price changes for identical items.

Most macro models of price adjustment abstract from temporary sales. Yet, temporary sales represent a large fraction of price changes in certain categories of retail products, such as grocery products and apparel. Like price changes due to product introduction, sales have quite distinct empirical characteristics from regular price changes. In particular, most sales are highly transient. When a product goes on sale, its price drops by a large amount. However, when the sale ends, the price typically returns to its original level. This is very different from regular price changes, which typically represent a quite persistent change in the price of a product. Kehoe and Midrigan (2007) analyze a model in which firms can make permanent and transitory price changes. They argue that temporary sales imply much less aggregate price adjustment than the same number of regular price changes.

3.1 The Frequency of Price Change: Consumer Prices

Table 1 reports estimates of the median frequency of price change for non-shelter goods and services in the CPI. These statistics are estimated by first calculating the mean frequency of price change for each ELI and then taking a weighted median across ELIs. In principle, calculating the frequency of price change is straightforward. It simply involves creating an indicator variable of when prices change and calculating the average of this variable. This procedure is complicated by three features of our retail price data. First, the data contain missing values as a consequence of stockouts. Second, products in the CPI research database are sometimes substituted out of the database and new products introduced in their stead. Third, we would like to construct statistics for price changes both including and excluding retail sales.¹¹

We present several measures of the frequency of price change based on different procedures for addressing these complications. The simplest procedure we use is to calculate the frequency of price change based only on contiguous price observations (Bils and Klenow, 2004). According to this procedure, if a product's price is observed in two consecutive months, and the price differs between

¹¹In order to identify temporary sales in the CPI data, we make use of the BLS sale flag. See section 2 for a discussion of this variable.

the two months, we define this as a price change; if the price is the same in the two months, we define this as no change. However, if either the current price or the price in the previous month are missing, we record a missing value in our price change indicator variable. In order to calculate the frequency of non-sale price change—which we refer to as the frequency of regular price change—we treat sales as missing observations and calculate the frequency of price change on contiguous non-sale price observations.¹²

Figure 1 graphically illustrates this simple procedure. The two panels in the figure report the first 10 observations for two hypothetical products. At the top of each panel, we record whether each observation is a sale or a regular price. The letter “R” denotes “regular price” while the letter “S” denotes “sale”. Below the sale flag is a graph of the evolution of the price of the product for these 10 observations. At the bottom of each panel, are two indicator variables that record price changes and regular price changes, respectively. First, notice that the price change variable and the regular price change variable are 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 variables. Third, notice that for the 8th observation the sale flag indicates that this price observation is a sale. In both panels, 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.

The key difference between the two panels is that in panel B a regular price change occurs during the sale. 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. In this regard, it is important to notice that while using only contiguous observations leads one to drop a price change in the situation depicted in panel B, it also causes one to drop a no change in the situation depicted in panel A. If sales occur randomly, we are simply randomly dropping a small fraction of the observations in the data set.

The median frequency of price change including sales in 1998-2005 measured based on contiguous observations was 19.4%, while it was 20.3% in 1988-1997. We define the corresponding median

¹²In all of these estimates we completely discard all the bimonthly observations in the data set. As a robustness test, we have compared the bimonthly frequency of price change in the portion of our dataset that is sampled bimonthly to the bimonthly frequency of price change in the portion of our dataset that is sampled monthly. The bimonthly frequency of price change is slightly lower in the bimonthly data than the monthly data.

implied duration to be $d = -1/\ln(1 - f)$, where f is the median frequency.¹³ Measured in this way, the median implied duration of raw prices in 1998-2005 was 4.6 months, while it was 4.4 months in 1988-1997. Including substitutions raises the median frequency by about 1% in each period. This estimator for the case including sales and substitutions is very similar to the procedure used by Bils and Klenow (2004).

We present four procedures for excluding sales. The first procedure estimates the frequency of regular price change off of contiguous non-sale observations as described above. This method yields a median frequency of price change excluding sales in 1998-2005 equal to 8.7%. The corresponding median implied duration is 11 months. The second procedure, calculates the frequency of regular price change by carrying forward the last observed regular price through sale and stockout periods that are followed by another regular price within 5 months.¹⁴ This method yields a median frequency of regular price change in 1998-2005 equal to 9.0%, which implies a median duration of 10.6 months.

The third and fourth procedures make use of direct measures of the frequency of regular price change during sale and stockout periods. Specifically, the third procedure is based on the following weighted average: $(1 - s)f + sf'$, where f is the measure of the frequency of price change based on contiguous non-sale, non-stockout 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. The fourth procedure is analogous to the third procedure except that it allows the frequency of regular price change to differ during stockouts as well as sales.¹⁵

The third and fourth procedures yield estimates of the median frequency of regular price change for 1998-2005 of 9.6% and 9.9%, respectively. The corresponding median implied durations are 9.9

¹³A constant hazard λ of price changes 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 the duration for the case where product exit is viewed as a censoring event. In other words, it is a measure of the median uncensored duration.

¹⁴This procedure is analogous to the procedure used by Klenow and Kryvtsov (2005). In practice, the estimates based on this procedure are not sensitive to the duration of sale or missing spells that we fill in since the majority of these spells last for one or two periods.

¹⁵We 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 period sales and two periods sales respectively. They are estimated using the method described in section 3.3. ω_1 is the fraction of sales that are one period sales. This procedure yields an upward biased estimate of the probability of price change during sale and missing periods due to Jensen's inequality. The choice of how to weight the probability that a price returns to the original price for spells of different lengths make little difference in practice. Using the empirical frequency of price change during sales of all different lengths yields virtually identical results. 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.

and 9.6 months. If the frequency of price change was the same during sale and stockout periods as during non-sale, non-stockout periods, these procedures would yield the same estimate as the procedure based only on contiguous non-sale, non-stockout observations. The fact that they yield higher estimates is because the frequency of regular price changes over the course of sales and stockouts is on average about 2 percentage points higher than during other periods.

Our results on the median frequency of regular price change are roughly in line with recent evidence on the frequency of price change in Europe based on CPI micro-data (Dhyne et al., 2006). The frequency of price change estimates in Dhyne et al. (2006) vary in their treatment of sales. However, Dhyne et al. report that in cases where sales could be identified they had little impact on the frequency of price change suggesting that sales are less important in Europe than in the U.S.

In summary, the median frequency of regular price change for 1998-2005 based on these four different procedures was 9-12% while the median frequency of price change including sales was 20%. Adjusting for sales makes such a large difference not only because sales are common in the data—21.5% of price changes are due to sales (table 2)—but also because of the uneven distribution of sales across goods. Table 3 reports the fraction of price change due to sales by Major Group. There is a huge amount of heterogeneity across Major Groups regarding the prevalence of sales. 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 and Vehicle Fuel are due to sales and only 3.1% of price changes in Services—a category that has an expenditure weight of 38.5%—are due to sales.

The sectors that have relatively few sales tend to be the sectors with 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 55% when sales are excluded, rather than 21.5%.

To see more clearly how heterogeneity in the prevalence of sales across sectors can lead to a large adjustment in the median frequency of price change, consider the three sector example presented in table 4. 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.¹⁶ However, a researcher that only knew that the overall fraction of price changes due to sales in the entire economy is 3/12 and adjusted the frequency of price change in all sectors using this number would conclude that the median frequency of price change is 3/8 and the median duration is 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 3). 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%, while 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 5 shows 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).

The large difference between the mean and median frequency of regular price change raises the question of which statistic is most informative about the degree of monetary non-neutrality in the economy. This question does not have a model free answer. In Nakamura and Steinsson (2006), we show that a multi-sector menu cost model calibrated to the cross-sectional distribution of the frequency of price change excluding sales yields a degree of monetary non-neutrality that is similar to that of a single sector model calibrated to the median frequency of regular price change but much greater than that of a single sector model calibrated to the mean frequency of regular price change.¹⁷ Carvalho (2006) argues that in a multi-sector Calvo model the mean duration provides a good approximation to the overall degree of monetary non-neutrality in the economy. Table 5 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

¹⁶For simplicity, we assume that only one price change can occur per month in this example.

¹⁷In Nakamura and Steinsson (2006), we also calibrate the multi-sector model to the sectoral distribution of the frequency of price change including sales, i.e., we treat sales as if they are the same as regular price changes. The degree of monetary non-neutrality implied by the calibration including and excluding sales differ by a factor between 1.5 and 2.

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.

Following Bils and Klenow (2004) and Dhyne et al. (2006), we have adopted a frequency based approach to estimating the median duration of price changes. A more direct approach would be to record the duration of each price spell and then find 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 complicated by 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.

Klenow and Kryvtsov (2005) present estimates of the median length of price spells based on duration data in footnote 3 of their paper. They do not present an estimate of the median duration of regular prices that adjusts for right and left censoring. The estimates in this footnote are therefore not comparable to our estimates. Klenow and Kryvtsov (2005) also present frequency based estimates of the mean frequency of price change that are similar to the estimates we present in table 5.¹⁸

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 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 non-sale prices or sales prices are partially set relative to a product's regular price. In the second case, even if all products are sold on sale the rigidity of the regular price still influences real quantities through the sales prices.

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

¹⁸Klenow and Kryvtsov's estimate of the mean frequency of price change is slightly higher than our estimate because they treat price changes over multiple periods for seasonal items the same as price changes over a single period. Our estimate of the mean frequency of price change is virtually identical to the corresponding estimate in Klenow and Willis (2006), where price changes over multiple periods for seasonal items are not included.

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 above, this assumption is dramatically at odds with the data.

3.2 Product Turnover

Table 6 reports information on product turnover for consumer products. Since 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.¹⁹

The main complication that arises in trying to relate the frequency of substitutions to the frequency of product introduction is that the CPI research database does not follow products over their entire lifetime. Following a substitution, BLS procedure for choosing a new product to sample tends to lead to the selection of products that have existed for some time.²⁰ Despite this caveat, the frequency of substitutions provides useful information on the frequency product turnover. We measure the frequency of substitutions as a fraction of the total product lifetime.²¹

The frequency of substitutions varies a great deal across difference Major Groups. Substitutions are much more common in durable goods categories than they are in other categories. In apparel, we estimate the frequency of substitutions to be 9.9%. Many clothes categories undergo substitutions

¹⁹Moulton 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.

²⁰Specifically, when a product in the dataset becomes unavailable, BLS pricing agents are instructed to substitute to the most similar available product. In sectors where fashion is important, this is likely to be an older product. 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. For example, Lancaster (1990) shows that if all products have a fixed lifetime, then the average time until a product exits is only half of its expected lifetime of the product.

²¹We define a product’s lifetime as the total time the product is priced and available, where we also include periods where the product is temporarily unavailable for 5 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. In this definition, we do not include any time periods after the last period when the product is priced and available. This measure differs from the measure used in Bils and Klenow (2004). They define the frequency of substitutions as a fraction of the total number of prices collected. We do not include product substitutions for which the price of the new product is observed in the months immediately after the substitution and is the same as the price of the old product.

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%.²²

3.3 The Behavior of Prices During and After Sales

Most of the existing literature on menu cost models does not attempt to fit the behavior of retail sales. Rather, this literature seeks to fit the behavior of prices excluding sales (see, e.g., Golosov and Lucas, 2006; and Midrigan, 2005). Explanations for sales in the industrial organization literature may be grouped into two categories. First, some sales arise due to intertemporal price discrimination of retail outlets (Varian, 1980; Sobel 1984). Price changes due to this type of sales may be largely orthogonal to macroeconomic aggregates. Second, sales are also used as a method for inventory management (Lazear, 1986; Aguirregabiria, 1999). Sales of this type are related to the macroeconomy to the extent that inventory is cyclical. Not all clearance sales are due to variation in aggregate demand. In some products—such as apparel—clearance sales may occur due to unpredictable shifts in tastes rather than shifts in aggregate demand (Pashigian, 1988). Pashigian and Bowen (1991) argue that price discrimination and clearance sales due to uncertainty about tastes account for most sales. Hosken and Reiffen (2001) document that sales are uncorrelated across retail outlets. They interpret this as evidence that retail sales are not primarily driven by changes in wholesale prices.

Sales are different from regular prices along important dimensions. Some features of sales seem inconsistent with the standard menu cost model. Table 7 documents a feature of sales that is particularly difficult to reconcile with the standard menu cost model, namely, the fact that the price of a product usually returns to its original regular price following a sale. The table presents statistics on sales for the 4 Major Groups for which sales are most important. For these 4 Major Groups, prices return to their original regular price between 60.0% and 86.3% of the time after a

²²See the “Supplementary Material” to this paper for a detailed analysis of the frequency of product substitutions in different sectors of the U.S. economy.

one period sale.²³ This fraction is highly negatively correlated with the frequency of regular price change. The fact that prices of Unprocessed Foods return to their original level after sales only 60.0% of the time may seem low. However, given that the frequency of regular price change in Unprocessed Food is 25.0%, for one period sales the probability that the regular price does not change between the month before the sale and the month after the sale is 56.2%.²⁴

We use these statistics to compare the frequency of regular price change during sales and the frequency of regular price change during non-sale periods. We calculate the fraction of one-period sales that have a different regular price immediately following the sale than immediately preceding the sale. From this number we calculate the monthly frequency of price change under the assumption that the hazard of price change was constant during this two month period. The resulting statistic is reported in column 2 of table 7. We find that the frequency of regular price change is similar during sale periods as during other periods in processed and unprocessed food but quite a bit higher in household furnishings and apparel. The simple average of the difference across these 4 Major Groups is 1.8 percentage points.

3.4 Frequency of Price Change: Producer Prices

Table 8 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 seem to have almost completely flexible prices. The median frequency of price change of crude materials in 1998-2005 was 98.9% and corresponding median implied duration is 0.2 months. Sales do not appear to be common in our producer price data set.²⁵ We therefore make no adjustment for sales when analyzing producer prices.

²³These statistics are based on sales for which a regular price for the product is observed after the sale ends. Clearance sales are therefore not included when these statistics are calculated.

²⁴These statistics probably underestimate the fraction of sales that return to the regular price since they are based on monthly statistics. Sales shorter than one month may revert to the original price, and then experience a regular price change.

²⁵The PPI database does not include a sales flag. We used the sales filters described in section 3.8 to assess the importance of sales in the producer price data. These sales filters identified very few sales.

Table 9 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 across sectors. Table 9 also reports the frequency of product substitution for these two digit Major Groups. The frequency of product substitution is close to 2% for all the Major Groups.

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%.

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. Blinder et al. (1998) surveyed firm managers about their pricing practices and found that prices changed on average once a year. Carlton (1986) estimated the rigidity of prices in the Stigler-Kindhal data set. He also found a substantial degree of price rigidity. Most of the prices analyzed in these studies were producer prices.

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 into long-term “implicit contracts” in which observed transaction prices are essentially payments 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 then choose to subsequently uphold them. In this type of situation retail prices would react to changes in manufacturer price 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).

3.5 Frequency of Price Change: CPI vs. PPI

In order to compare price flexibility at the consumer and producer levels, we matched 153 ELI's from the CPI with product codes from the PPI.²⁶ Table 10 presents comparisons between the frequency of price change at the consumer and producer level for the Major Groups in which a substantial number of matches 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.

3.6 The 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.²⁷ With even a modest amount of inflation, these models imply that almost all price changes are price increases. Table 3 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%.²⁸ Table 8 shows that the same pattern emerges for producer prices. The fraction of price changes in producer prices 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 (2006)—it provides strong evidence for the hypothesis

²⁶42 ELI's were matched to PPI categories at the 8 digit product-code level, 71 ELI's were matched to PPI categories at the 6 digit product-code level and 40 ELI's were matched to PPI categories at the 4 digit product-code level. When an ELI was matched to a PPI category at, for example, the 6 digit product code level, the unweighted median of the mean frequency of price change of item codes within that 6 digit product code was used.

²⁷Examples include, Taylor (1980), Calvo (1983), Caplin and Spulber (1987), Dotsey et al. (1999) and Mankiw and Reis (2002). A notable exception is Golosov and Lucas (2006).

²⁸These 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.

that idiosyncratic shocks are an important driving force of price changes.

3.7 The Size of Price Changes

Price adjustment seems 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 11 reports the median absolute size of log changes in consumer prices. For consumer prices excluding sales, the median absolute size of price changes is 8.5%.²⁹ This table also reports the absolute size of price change by Major Group. Price changes that are due to sales are on average much larger than regular price changes. Table 11 reports that the median absolute size of price changes due to sales is 29.5%, more than three times the size of regular prices. We have also calculated the size of price changes 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 11 is that the median size of price decreases is 3.2 percentage points larger than the median size price increases. For finished goods producer prices this difference is 1 percentage points. The median size of price decreases is larger than that of price increases for 10 of 11 Major Groups for consumer price and 11 of 15 Major Groups for producer prices.

3.8 Alternative Measures of Sales

Up until now we have used the BLS sale flag to identify sales. In doing so, we follow a number of previous papers (Bils and Klenow, 2004; Klenow and Kryvtsov, 2005). An alternative (and complementary) approach to identifying sales is to look for “V-shaped” patterns in the data and identify these patterns as sales. An important conceptual difference between this “sales filter” approach and our previous approach is that clearance sales are not defined as “sales” according to this approach.

There are two main empirical drawbacks of the sale filter approach as a mechanism for identifying V-shaped sales. First, since 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

²⁹This statistic is calculated by finding the average log change in price by ELI and then taking the weighted median across ELI’s.

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. Another type of event that would not be captured by a V-shaped filter is if the good is on sale twice in a row when the BLS price collector samples it but at a different sale price—say a 30% discount and then a 50% discount.

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—e.g., from \$2.49 to \$2.59. For this reason, sales filters will indicate that gasoline is on sale a significant fraction of the time, while the BLS sale flag indicates that there are virtually no sales in the gasoline category.

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 12 reports results for two types of sales filters, which we refer to as sales filter 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. Sales filter A is designed to also remove price patterns in which a sale is followed by a change in the regular price, i.e., 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 1 and 5 months. For example, for the 2 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 B with a window of 5 months is 12.3% for the 1998-2005 period. The median frequency of price change based on the more complex sale filter A is 11.0% over this period. This statistic is similar to the weighted median frequency of price change that uses the sales flag to exclude all sales except for clearance sales. 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, if one assumes a window of one month 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 between 11.4% and 16.4%.

4 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 (2006). We analyze whether the facts established in the preceding section are consistent with this model and what they imply about the values of its key parameters.

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

$$y_t(z) = A_t(z)L_t(z), \quad (1)$$

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

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

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 which determines the "size of the market" for the firm's good. In order 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

$$\frac{W_t}{P_t} = \frac{\theta - 1}{\theta}, \quad (3)$$

where W_t denotes nominal wage rate in the economy at time t .³⁰

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

$$\Pi_t(z) = C \left(\frac{p_t(z)}{P_t} \right)^{-\theta} \left(\frac{p_t(z)}{P_t} - \frac{\theta - 1}{\theta} \frac{1}{A_t(z)} \right) - \frac{\theta - 1}{\theta} K I_t(z), \quad (4)$$

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

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

³⁰In 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 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 non-neutrality is small, variation in C_t will be small and the real wage will be approximately constant.

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 β . The value function of the firm is given by the solution to

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

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: β , θ , K/C , μ , ρ , σ_ϵ and σ_η . We set the monthly discount factor equal to $\beta = 0.96^{1/12}$. We choose $\theta = 4$ to roughly match estimates from the industrial organizations literature on markups of price over marginal costs.³¹ We estimate $\mu = 0.0021$ and $\sigma_\eta = 0.0032$ from data on the CPI from 1998-2005. This sample period was chosen to correspond to the more recent sample period for which we report result from the CPI Research Database.

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 price changes in 1998-2005. The parameter values that imply that the model matches the data along these three dimensions are $K/C = 0.0245$, $\rho = 0.660$, $\sigma_\epsilon = 0.0428$.³² 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 non-neutrality. We minimize the squared deviations of the model implied values of these moments from their estimated values in the data. Using this loss

³¹Berry et al. (1995) and Nevo (2001) find that markups vary a great deal across firms. The value of θ we choose implies a markup similar to the mean markup estimated by Berry et al. (1995) but slightly below the median markup found by Nevo (2001). Broda and Weinstein (2006) estimate elasticities of demand for a large array of disaggregated products using trade data. They report a median elasticity of demand below 3. Midrigan (2005) uses $\theta = 3$ while Golosov and Lucas (2006) use $\theta = 7$. The value of θ is not important for the points we make in this paper.

³²Were we to assume $\theta = 10$, our estimate of K/C would rise to 0.07. All other results would be essentially unaffected.

function, the parameters are well identified. Since we are able to exactly match all three parameters, the relative weights on different parameters in the loss function do not matter.

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.

5 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. This empirical result, which we document in this section, provides a natural test for our menu cost model. 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.³³ Figures 2 and 3 plot the annual evolution of these four series for consumer prices along with the evolution of CPI inflation.³⁴ 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.

The correlation between the frequency of price increases and inflation is 0.81. Furthermore, a regression of the median frequency of price increases on aggregate inflation over the period 1988-2005 indicates that a 1 percentage point increase in the inflation rate is associated with approximately a 1 percentage point increase in the median frequency of price increases. A back-of-the-envelope calculation indicates that the variation in the frequency of price increases is large enough to account for most of the variation in aggregate inflation. A 1 percentage point increase in the monthly frequency of price increases, is associated with an increase of 0.1 percentage points in monthly inflation (since the average size of price increases is approximately 10%). This corresponds

³³Gagnon (2005) has emphasized the importance of distinguishing between price increases and price decreases in this context. Analyzing disaggregated Mexican consumer price data, he found that at low levels of inflation (<10% per year) the overall frequency of price change covaries little with inflation because movements in the frequency of price decreases partly offset movements in the frequency of price increases.

³⁴As in section 3, these statistics are calculated by first calculating the mean frequency (size) within each ELI and then finding the weighted median across ELIs.

to approximately a 1 percentage point increase in annual inflation.

Table 13 conveys through regressions what figures 2 and 3 convey graphically. We regress the four components at the ELI-level on 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 the results for the 1988-1997 period, which are most pronounced, are heavily influenced by the behavior of inflation and the frequency of price change in 1990; and more generally that these results should be interpreted with caution given the small amount of inflation variability over the period we consider.³⁵ Yet, as we discuss below, the empirical behavior of the frequency of price change appears to be broadly consistent with the simple menu cost model.

Figures 4 and 5 compare the evolution of these variables in the model to their 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 and size of price increases and decreases by year. Just as in the data, the frequency of price 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 (2006) 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 towards 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.

Figure 5 shows that the model also matches the fact that the median size of price decreases is

³⁵The 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. We also considered specifications where we include a dummy variable for the year 1990. In these specifications, the coefficient fell to 0.68 (0.36) for the frequency of price change and was virtually unchanged at 0.97 (0.18) for the frequency of price increases, in the regressions for the 1988-1997 period.

larger than that for price increases. Ellingsen et al. (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. The relative size of price increases and price decreases also depends on the steady state rate of inflation. As the steady state rate of inflation rises the size of price increases eventually becomes larger than the size of price decreases.³⁶

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 4 shows that for the economy as a whole we do not find evidence of this phenomenon. 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.³⁷

The finding that the frequency of price change covaries more with variation in the inflation rate than the size of price changes is consistent with a number of previous empirical studies. Vilmunen and Laakkonen (2004) and Gagnon (2005) provide direct evidence for this phenomenon. Lach and Tsiddon (1992), Cecchetti (1986), Kashyap (1995), and Goette et al. (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.

Klenow and Kryvtsov (2005) find that most of the variation of aggregate inflation stems from variation in the average size of price changes. At first glance, our results may seem to contradict their results. Notice, however, that the average size of price change may be decomposed as $s_{all} = f_u s_u - f_d s_d$, 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 price changes.

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.

³⁶A alternative explanation for the fact that price decreases are larger than price increases in the data is that we may have failed to filter out all sales.

³⁷There 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; while the frequency of price change for travel services rose again monotonically from approximately 20% in 1988 to 50% in 2005.

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 6 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.³⁸ The average size of sales has also increased substantially over the sample period in all of the categories except for household furnishings. The increase is most dramatic in processed food, where the size of sales has nearly doubled from about 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 1960's.

Table 14 presents the results of regressions of the frequency and size of sales on CPI inflation, ELI fixed effects and a time trend. We do not find robust evidence of a relationship between either the size or frequency of sales and aggregate variables. For both 1988-1997 and 1998-2005, we find a negative coefficient on the inflation rate, but neither coefficient is statistically significant at the 5% level. This suggests that a small effect may exist, 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.

6 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. Analyses of price change behavior often discuss the existence of a “pricing season”. Yet the magnitude of this phenomenon, and the extent to which pricing seasons are coordinated across firms, have not previously been documented for the U.S. economy. We find

³⁸The 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 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.

a substantial seasonal component of price changes for the U.S. economy, for both consumer and producer goods.

Figure 7 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. The decline in the frequency of price increases between the first and fourth quarter is 1.9 percentage points, or 25%. In contrast, price decreases decline by 0.6 percentage point, or 18%, between the first and last quarter.³⁹

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 of price change comes from the frequency of price increases.

Olivei and Tenreyro (2005) 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 of an allocative effect of seasonality in wage setting or other components of costs.

The seasonal pattern in sales is very different from the seasonal pattern in regular price changes. Figure 8 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,

³⁹Álvarez et al. (2005b) find that prices are significantly more likely to change in January in the Euro Area.

January and June than in the months with the least sales. However, even in these other months, more than 25% of price quotes are sales in Apparel. The yearly winter and summer sales are clearly not the only sales in Apparel. This pattern has remained roughly unchanged between 1988-1997 and 1998-2005 while the overall level of sales in Apparel has increased dramatically. We find much less seasonality in sales in other Major Groups.

7 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 $\lambda(t) = P(T = t | T \geq t)$. 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. More generally, the shape of the hazard function is influenced by the relative size of transient and permanent shocks to marginal costs. Non-stationarity in marginal costs—e.g. 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 9 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.⁴⁰ In contrast, the Calvo model assumes a flat hazard function of price change.

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

⁴⁰The 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.

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 leads to a downward bias in the estimated slope of the hazard function. Even if the hazard functions of all the goods are flat or upward sloping, heterogeneity in the level of the hazard function of different products can cause the estimated hazard function to be downward sloping.⁴¹

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).⁴² 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 non-parametrically. Specifically, we assume that the hazard function is

$$\lambda_i(t|x_{i,j}) = \nu_i \lambda_0(t) \exp(x_{i,j} \beta) \quad (7)$$

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 non-parametric baseline hazard function with dummies for each month, $x_{i,j}$ is a vector of covariates for the j th observation of products i and β is a vector of parameters.⁴³ We assume that $\nu_i \sim \text{Gamma}(1, \sigma_\nu^2)$.⁴⁴ 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 .⁴⁵ We estimate the model by maximum likelihood. We truncate the price spells at 18 months and drop left censored spells.⁴⁶

For consumer prices, we have approximately 2.75 million price spell observations after dropping left-censored spells. When we exclude sales as well, we are left with 1.65 million observations. We divide the data set into groups at the level of Major Groups. Figure 10 plots the baseline hazard

⁴¹See Kiefer (1988) for a survey of hazard function estimation.

⁴²An example of a “product” is 16oz Kraft Singles sold at a particular supermarket in New York.

⁴³The only covariates we consider are seasonal month dummies.

⁴⁴We have estimated the model with $\nu_i \sim N(1, \sigma_\nu^2)$. The results are virtually identical.

⁴⁵See Honore (1993) for a discussion of identification results for multiple spell duration models.

⁴⁶In the presence of heterogeneity, discarding left-censored spells leads us to disproportionately drop price spells arising from subjects with low values of ν_i , since 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.

function from the model described by equation (7) for Processed Food and Services.⁴⁷ Each panel 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 for the Major Groups. This pattern holds 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 8 subgroups. For the major groups in which sales occur frequently (i.e., Processed and Unprocessed Food, Household Furnishings and Apparel), the hazard function including sales is much more steeply downward sloping than the hazard function of regular prices. Models of price change designed to match the behavior of prices including sales must generate this steeply downward sloping hazard function. For Services, we estimate a large spike in the hazard function at 12 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 12 month spike is completely absent in most other Major Groups. We do not include information about the standard errors of our estimates in figure 10 because the standard errors are very small in most cases due to the large number of observations.

For producer prices, we have 1.95 million price spells after dropping left-censored spells. 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 are similar for producer prices as they are for consumer prices. The hazard functions are downward sloping for the first few months, then mostly flat except for a large 12 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 12 months. Interestingly, the 12 month spike in the hazard function is a much more pervasive phenomena in producer prices than in consumer prices.

The existing evidence on the shape of the hazard function is mixed. Empirical support for upward sloping hazard functions appears to arise mostly in studies in which almost all price changes are increases, indicating few idiosyncratic shocks (Goette et al. 2005; and Cecchetti, 1986), or in

⁴⁷The extended version of this paper—available on our websites—reports plots of the hazard function of eight Major Groups for consumer prices and another eight Major Groups for producer prices.

periods of very high inflation (Gagnon, 2005).⁴⁸ These facts line up well with the basic intuitions about the shape of the hazard function that arise from the menu cost model.

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 while in the model it is small and rising sharply. We have considered an extension of our benchmark model with heteroskedastic shocks to marginal costs. This model generates a downward sloping hazard function in the first few months that lines up with the data. In a general equilibrium setting, the model with heteroskedastic shocks to marginal costs yields similar behavior for macroeconomic aggregates as the model with homoskedastic shocks to marginal costs. Another way to generate a downward sloping hazard function is time variation in the cost of changing prices. This type of model generates substantially more monetary non-neutrality than a general equilibrium version of our benchmark model (Nakamura and Steinsson, 2006). Finally, sufficiently large and transient shocks to marginal costs can generate a downward sloping hazard function.

8 Conclusion

In this paper, we present new evidence on price adjustment in the U.S. economy. Using BLS micro-data we document that the median frequency of non-sale 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 price change for finished goods producer prices is roughly 11% per month. We argue that it is important to analyze regular price changes, price changes due to product introduction and temporary sales separately since these three types of price changes have quite distinct empirical features and are likely to have different macroeconomic implications. We show that one third of price changes are price decreases. Combined with the large average absolute size of price changes, the large number of price decreases observed in the data provides clear evidence that idiosyncratic shocks are an important source of price changes.

We find that the frequency of price increases covaries strongly with inflation while the frequency

⁴⁸The results from some of these studies are hard to interpret since they use the conditional logit formulation. This formulation is biased in this application (see Willis, 2006).

of price decreases and the size of price increases and price decreases do not. We show that this pattern is consistent with the implications of a benchmark menu cost model. We document a dramatic secular rise in the frequency and size of sales in several sectors of the U.S. economy. We do not find robust evidence that sales respond to aggregate variables. We find that the frequency of price change is highly seasonal. It is highest in the 1st quarter and lowest in the 4th quarter. Furthermore, in consumer prices the frequency of price change is highest in the first month of each quarter and falls monotonically within quarter.

Finally, we estimate the hazard function of price change for consumer and producer prices accounting for heterogeneity at the product level. We find that this hazard function is slightly downward sloping for the first few months and then flat (except for a large spike at 12 months in consumer services and all producer prices). This pattern is not consistent with our benchmark menu cost model. The calibrated benchmark model yields a hazard function that is sharply upward sloping in the first few months and does not imply a spike at 12 months. The spike at 12 months may be evidence of a time-dependent element of price setting. The fact that the empirical hazard is large and falling in the first few months may be evidence of heteroskedasticity in marginal costs or time variation in the cost of adjusting prices.

References

- AGUIRREGABIRIA, V. (1999): “The Dynamics of Markups and Inventories in Retail Firms,” *Review of Economic Studies*, 66, 275–308.
- ÁLVAREZ, L. J., P. BURRIEL, AND I. HERNANDO (2005a): “Do Decreasing Hazard Functions for Price Changes Make Sense?,” Working Paper No. 461, European Central Bank.
- ÁLVAREZ, L. J., E. DHYNE, M. M. HOEBERICHTS, C. KWAPIL, H. L. BIHAN, P. LUNNEMANN, F. MARTINS, R. SABBATINI, H. STAHL, P. VERMEULEN, AND J. VILMUNEN (2005b): “Sticky Prices in the Euro Area: A summary of New Micro Evidence,” Working Paper No. 563, European Central Bank.
- ARELLANO, M. (1987): “Computing Robust Standard Errors for Within-Groups Estimators,” *Oxford Bulletin of Economics and Statistics*, 49(4), 431–434.
- BAHARAD, E., AND B. EDEN (2004): “Price Rigidity and Price Dispersion: Evidence from Micro Data,” *Review of Economic Dynamics*, 7(3), 613–641.
- BARRO, R. J. (1972): “A Theory of Monopolistic Price Adjustment,” *Review of Economic Studies*, 39(1), 17–26.
- (1977): “Long Term Contracting, Sticky Prices and Monetary Policy,” *Journal of Monetary Economics*, 3, 305–316.
- BAUMGARTNER, J., E. GLATZER, F. RUMLER, AND A. STIGLBAUER (2005): “How Frequently do Consumer Prices Change in Austria?,” Working Paper No. 523, European Central Bank.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BILS, M., AND P. J. KLENOW (2002): “Some Evidence on the Importance of Sticky Prices,” NBER Working Paper No. 9069.
- (2004): “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112(5), 947–985.
- BLINDER, A. S., E. R. D. CANETTI, D. E. LEBOW, AND J. B. RUDD (1998): *Asking About Prices*. Russell Sage Foundation, New York, New York.
- BRODA, C., AND D. E. WEINSTEIN (2006): “Globalization and the Gains from Variety,” *Quarterly Journal of Economics*, 121(4), 541–585.
- CALVO, G. A. (1983): “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–398.
- CAMPBELL, J. R., AND B. EDEN (2004): “Rigid Prices: Evidence from U.S. Scanner Data,” Working Paper, Vanderbilt University.
- CAPLIN, A., AND D. SPULBER (1987): “Menu Costs and the Neutrality of Money,” *Quarterly Journal of Economics*, 102(4), 703–725.

- CARLTON, D. W. (1979): “Contracts, Price Rigidity and Market Equilibrium,” *Journal of Political Economy*, 87(5), 1034–1062.
- (1986): “The Rigidity of Prices,” *American Economic Review*, 76(4), 637–658.
- CARVALHO, C. (2006): “Heterogeneity in Price Stickiness and the New Keynesian Phillips Curve,” Working Paper, Princeton University.
- CECCHETTI, S. G. (1986): “The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines,” *Journal of Econometrics*, 31, 255–274.
- DHYNE, E., L. J. ÁLVAREZ, H. L. BIHAN, G. VERONESE, D. DIAS, J. HOFFMANN, N. JONKER, P. LUNNEMANN, F. RUMLER, AND J. VILMUNEN (2006): “Price Setting in the Euro Area and the United States: Some Facts From Individual Consumer Price Data,” *Journal of Economic Perspectives*, 20(2), 171–192.
- DIAS, D. A., C. ROBALO MARQUES, AND J. M. SANTO SILVA (2005): “Time or State Dependent Price Setting Rules? Evidence from Portuguese Micro Data,” Working Paper No. 511, European Central Bank.
- DOTSEY, M., R. KING, AND A. WOLMAN (1999): “State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output,” *Quarterly Journal of Economics*, 114(2), 655–690.
- ELLINGSEN, T., R. FRIBERG, AND J. HASSLER (2006): “Menu Costs and Asymmetric Price Adjustment,” Working Paper, Stockholm School of Economics.
- FABIANI, S., M. DRUANT, I. HERNANDO, C. KWAPIL, B. LANDAU, C. LOUPIAS, F. MARTINS, T. MATHA, R. SABBATINI, H. STAHL, AND A. STOKMAN (2004): “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.
- FOUGÉRE, D., H. L. BIHAN, AND P. SEVESTRE (2005): “Heterogeneity in Consumer Price Stickiness: A Microeconomic Investigation,” Working Paper No. 536, European Central Bank.
- GAGNON, E. (2005): “Price Setting Under Low and High Inflation: Evidence from Mexico,” Working Paper, Northwestern University.
- GOETTE, L., R. MINSCH, AND J.-R. TYRAN (2005): “Micro Evidence on the Adjustment of Sticky-Price Goods: It’s How Often, Not How Much,” Discussion Paper, University of Copenhagen.
- GOLOSOV, M., AND R. E. LUCAS (2006): “Menu Costs and Phillips Curves,” Working Paper, MIT.
- GOPINATH, G., AND R. RIGOBON (2006): “Sticky Borders,” Working Paper, Harvard University.
- HECKMAN, J. J., AND B. SINGER (1986): “Econometric Analysis of Longitudinal Data,” in *Handbook of Econometrics, Volume III*, ed. by Z. Griliches, and M. D. Intrilligator, pp. 1689–1763. Elsevier Science Publishers.
- HOBIIJN, B., F. RAVENNA, AND A. TAMBALOTTI (2006): “Menu Costs at Work: Restaurant Prices and the Introduction of the Euro,” *Quarterly Journal of Economics*, 121(3), 1103–1131.

- HONORE, B. (1993): "Identification Results for Duration Models with Multiple Spells," *Review of Economic Studies*, 60(1), 241–246.
- HOSKEN, D., AND D. REIFFEN (2001): "Pricing Behavior of Multiproduct Retailers," Working Paper.
- (2004): "Patterns of Retail Price Variation," *Rand Journal of Economics*, 35(1), 128–146.
- JONKER, N., C. FOLKERTSMA, AND H. BLIJENBERG (2004): "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.
- KACKMEISTER, A. (2005): "Yesterday's Bad Times are Today's Good Old Times: Retail Price Changes in the 1890's were Smaller, Less Frequent, and More Permanent," Finance and Economics Discussion Series, Federal Reserve Board.
- KASHYAP, A. K. (1995): "Sticky Prices: New Evidence from Retail Catalogs," *Quarterly Journal of Economics*, 110, 245–274.
- KEHOE, P., AND V. MIDRIGAN (2007): "Sales, Clustering of Price Changes, and the Real Effects of Monetary Policy," Working Paper, University of Minnesota.
- KIEFER, N. M. (1988): "Economic Duration Data and Hazard Functions," *Journal of Economic Literature*, 26(2), 646–679.
- KLENOW, P. J., AND O. KRYVTSOV (2005): "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation," Working Paper, Stanford University.
- KLENOW, P. J., AND J. L. WILLIS (2006): "Real Rigidities and Nominal Price Changes," Federal Reserve Bank of Kansas City Working Paper.
- KONIECZNY, J. D., AND A. SKRZYPACZ (2005): "Inflation and Price Setting in a Natural Experiment," *Journal of Monetary Economics*, 52(3), 621–632.
- LACH, S., AND D. TSIDDON (1992): "The Behavior of Prices and Inflation: An Empirical Analysis of Disaggregated Price Data," *Journal of Political Economy*, 100(2), 349–389.
- LANCASTER, T. (1979): "Econometric Methods for the Duration of Unemployment," *Econometrica*, 47(4), 939–956.
- (1990): *The Econometric Analysis of Transition Data*. Cambridge University Press, Cambridge.
- LAZEAR, E. P. (1986): "Retail Pricing and Clearance Sales," *American Economic Review*, 76, 14–32.
- MANKIW, N. G., AND R. REIS (2002): "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117(4), 1295–1328.
- MEYER, B. D. (1986): "Semiparametric Estimates of Hazard Models," Mimeo, MIT.
- (1990): "Unemployment Insurance and Unemployment Spells," *Econometrica*, 58(4), 757–782.

- MIDRIGAN, V. (2005): “Menu Costs, Multi-Product Firms, and Aggregate Fluctuations,” Working Paper, Ohio State University.
- MOULTON, B. R., AND K. E. MOSES (1997): “Addressing the Quality Change Issue in the Consumer Price Index,” *Brookings Papers on Economic Activity*, (1), 305–349.
- NAKAMURA, E., AND J. STEINSSON (2006): “Monetary-Non-Neutrality in a Multi-Sector Menu Cost Model,” Working Paper, Harvard University.
- NEVO, A. (2001): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 69(2), 307–342.
- OLIVEI, G., AND S. TENREYRO (2005): “The Timing of Monetary Policy Shocks,” Working Paper, Federal Reserve Bank of Boston.
- PASHIGIAN, B. P. (1988): “Demand Uncertainty and Sales: A Study of Fashion and Markdown Pricing,” *American Economic Review*, 78(5), 936–953.
- PASHIGIAN, P. B., AND B. BOWEN (1991): “Why Are Products Sold on Sales? Explanation of Pricing Regularities,” *Quarterly Journal of Economics*, 106, 1015–1038.
- SHESHINSKI, E., AND Y. WEISS (1977): “Inflation and Costs of Price Adjustment,” *Review of Economic Studies*, 44(2), 287–303.
- SOBEL, J. (1984): “The Timing of Sales,” *Review of Economic Studies*, 51, 353–368.
- STIGLER, G. J., AND J. K. KINDAHL (1970): *The Behavior of Industrial Prices*. Columbia University Press, New York, N.Y.
- TAUCHEN, G. (1986): “Finite State Markov-Chain Approximation to Univariate and Vector Autoregressions,” *Economics Letters*, 20(2), 177–181.
- TAYLOR, J. B. (1980): “Aggregate Dynamics and Staggered Contracts,” *Journal of Political Economy*, 88, 1–23.
- U.S. DEPARTMENT OF LABOR (1997): *BLS Handbook of Methods*. Government Printing Office, Washington, D.C.
- VARIAN, H. R. (1980): “A Model of Sales,” *American Economic Review*, 70, 651–659.
- VERMEULEN, P., D. DIAS, M. DOSSCHE, E. GAUTIER, I. HERNANDO, R. SABBATINI, AND H. STAHL (2006): “Price setting in the euro area: Some stylised facts from Individual Producer Price Data and Producer Surveys,” Working Paper.
- VILMUNEN, J., AND H. LAAKKONEN (2004): “How Often Do Prices Change in Finland? Micro-Level Evidence from the CPI,” Working Paper, Bank of Finland.
- WILLIS, J. L. (2006): “Magazine Prices Revisited,” *Journal of Applied Econometrics*, 21, 337–344.

Table 1: Median Frequency of Price Change in the CPI

	Median Frequency		Median Implied Duration	
	88-97	98-05	88-97	98-05
	%		Months	
Panel A: Including Sales				
1. Excluding Substitutions	20.3	19.4	4.4	4.6
2. Including Substitutions	21.7	20.5	4.1	4.4
Panel B: Excluding Sales and Substitutions				
3. Contiguous observations	11.1	8.7	8.5	11.0
4. Carry regular price forward during sales and stockouts	11.2	9.0	8.4	10.6
5. Estimate freq. of price change during sales	11.5	9.6	8.2	9.9
6. Estimate freq. of price change during sales and stockouts	11.9	9.9	7.9	9.6
Panel C: Excluding Sales, Including Substitutions				
7. Contiguous observations	12.7	10.9	7.4	8.7
8. Carry regular price forward during sales and stockouts	12.3	10.6	7.6	8.9
9. Estimate freq. of price change during sales	12.8	11.3	7.3	8.3
10. Estimate freq. of price change during sales and stockouts	13.0	11.8	7.2	8.0

All frequencies are reported in percent per month. Implied durations are reported in months. "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 ELI's within the Major Group using CPI expenditure weights. The implied duration is equal to $-1/\ln(1-f)$, where f is the median frequency of price change.

Table 2: Frequency of Sales

	1988-1997	1998-2005
Expenditure weighted:		
Fraction of Price Changes Due to Sales	21.2	21.5
Fraction of Price Quotes with Sales	6.6	7.4
Weighted by Number of Observations:		
Fraction of Price Quotes with Sales	10.3	12.1

All statistics are reported as percentages. For the statistics in the first row, we first calculate the fraction of price changes due to sales for each ELI and then take an expenditure weighted mean across ELIs. The same procedure is used in the second row to calculate the expenditure weighted fraction of price quotes with sales. In the last row, we apply the same procedure as in the second row except that values for each ELI are weighted by the number of price observations for that ELI.

Table 3: Frequency of Price Change by Major Group in 1998-2005

Major Group	Weight	Regular Prices				Prices				Sales	
		Median Freq.	Impl.Dur	Mean Freq.	Frac. Up	Median 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

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. "# Obs." denotes the number of price observations for each 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 ELI's within the Major Group using CPI expenditure weights. The other median statistics in this table are calculated in an analogous manner. "Median Dur." is equal to $-1/\ln(1-f)$, where f is the median frequency of price change. "Median Ch.+Sub." denotes the median of the frequency of price change including price changes associated with substitutions. "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.

Table 4: 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

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.

Table 5: Mean Frequency of Price Change in the CPI

	Mean Frequency		Mean Implied Duration	
	88-97	98-05	88-97	98-05
	%		Months	
Panel A: Including Sales				
1. Excluding Substitutions	23.9	26.5	8.3	9.0
2. Including Substitutions	25.2	27.7	7.5	7.7
Panel B: Excluding Sales and Substitutions				
3. Contiguous observations	18.7	21.1	11.6	13.0
4. Carry regular price forward during sales and stockouts	18.6	20.9	11.0	12.3
5. Estimate freq. of price change during sales	19.0	21.3	11.2	12.5
6. Estimate freq. of price change during sales and stockouts	18.9	21.5	10.8	11.7
Panel C: Excluding Sales, Including Substitutions				
7. Contiguous observations	20.4	22.8	9.3	9.8
8. Carry regular price forward during sales and stockouts	19.7	22.0	9.6	10.4
9. Estimate freq. of price change during sales	20.8	22.8	9.2	9.8
10. Estimate freq. of price change during sales and stockouts	20.7	23.1	9.0	9.3

All frequencies are reported in percent per month. Implied durations are reported in months. "Mean Frequency" denotes the weighted mean frequency of price change. It is calculated by first calculating the mean frequency of price change for each ELI and then taking a weighted mean across ELI's using CPI expenditure weights. "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.

Table 6: Frequency of Substitution and Price Change by Category

Major Group	weight	Subs.	Pr.Ch. + Prod. Intro.		Price Change	
		Freq.	Freq. Reg.	Freq.	Freq. Reg.	Freq.
Processed Food	8.2	1.3	11.1	26.6	10.5	25.9
Unprocessed Food	5.9	1.2	25.6	37.9	25.0	37.3
Household Furnishing	5.0	5.0	9.3	23.1	6.0	19.4
Apparel	6.5	9.9	10.7	36.9	3.6	31.0
Transportation Goods	8.3	10.2	36.6	36.6	31.3	31.3
Recreation Goods	3.6	6.3	9.5	16.7	6.0	11.9
Other Goods	5.4	1.0	15.6	17.1	15.0	15.5
Utilities	5.3	0.6	38.4	38.4	38.1	38.1
Vehicle Fuel	5.1	0.2	87.6	87.6	87.6	87.6
Travel	5.5	1.9	42.6	43.6	41.7	42.8
Services (excl. Travel)	38.5	0.9	7.5	7.5	6.1	6.6

The sample period is 1998-2005. "Subs. Freq." gives the median average 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 5 months or less when the product is unavailable at the time of sampling but subsequently becomes available. "Pr. Ch. + Prod. Intro." indicates the median average monthly frequency of price change adjusted for product turnover according to the formula $(1-f)(1-pc)$ where $f=0.75f_{sub}$ and f_{sub} is the frequency of product substitution discussed above and the formula relating f and f_{sub} is discussed in the text. "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 ELI's and then calculating the expenditure-weighted median across ELI's. "Weight" denotes the expenditure weight of the ELI. "CDF" denotes the cumulative distribution function of the frequency of regular price change.

Table 7: Sales and Prices During Sales

	Freq. Reg. Price Ch.	Freq. Price Ch.	Frac. Return	Frac. of Sales	Freq. Price Ch.
		During One	After One	that Last One	Dur. One Period
		Period Sales	Period Sales	Period	Sales/Missing
Processed Food	10.5	11.4	78.5	64.7	11.1
Unprocessed Food	25.0	22.5	60.0	63.2	22.1
Household Furnishings	6.0	11.6	78.2	43.3	9.4
Apparel	3.6	7.1	86.3	35.8	5.9

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 frequency of regular price changes during one month 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/Missing" 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.

Table 8: Frequency of Price Change for Producer Prices

	Finished Goods		Intermediate Goods		Crude Materials	
	88-97	98-05	88-97	98-05	88-97	98-05
Median Freq. of Change	10.6	10.8	11.4	13.3	73.5	98.9
Median Implied Duration	8.9	8.7	8.3	7.0	0.8	0.2
Median Frac. of Increases	65.3	60.6	61.1	58.4	48.4	56.1
Mean Freq. of Change	25.2	24.7	21.7	26.7	78.0	86.0

Frequencies are reported in percent per month. Implied durations are reported in months. Fractions are reported in percentages. The median frequency of price change is calculated by first calculating the mean frequency of price change for each cell code, then taking an unweighted median within 4-digit commodity code and then taking a value weighted median across 4-digit commodity codes. The median implied duration is $-1/\ln(1-f)$, where f is the median frequency of price change. The mean frequency of price change is a value weighted mean across the 4-digit commodity code statistics discussed above. "88-97" and "98-05" denote the time periods 1988-1997 and 1998-2005, respectively.

Table 9: Frequency of Price Change by Major Group for the Finished Goods PPI

Category Name	Weight	Med. Freq.	Med. Freq.	Frac. Up
		Price Ch.	Substitutions	
Farm Products	1.6	87.5	2.2	48.6
Processed Foods and Feeds	22.4	26.3	2.6	57.8
Textile Products and Apparel	3.6	2.3	2.9	49.7
Hides, Skins, Leather, and Related Products	0.3	3.8	2.9	80.0
Fuels and Related Products and Power	20.8	48.7	2.3	54.1
Chemicals and Allied Products	2.8	6.1	2.9	61.6
Rubber and Plastic Products	1.8	3.2	2.9	83.8
Lumber and Wood Products	0.1	1.3	3.0	86.6
Pulp, Paper and Allied Products	3.0	4.4	2.5	74.9
Metals and Metal Products	1.1	3.8	2.5	72.2
Machinery and Equipment	13.0	3.7	2.7	71.0
Furniture and Household Durables	5.6	5.1	2.4	78.6
Nonmetallic Mineral Products	0.1	4.1	2.3	67.0
Transportation Equipment	16.8	27.3	3.0	53.7
Miscellaneous Products	6.9	16.5	1.8	81.3

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 4-digit commodity code and then taking a value weighted median across 4-digit commodity codes within the Major Group. "Frac Up" denotes the median fraction of price increases. It is calculated in an analogous manner to the median frequency of price change.

Table 10: Frequency of Price Change: Comparison of CPI and PPI Categories

Category	Num. of Matches	Frequency			Implied Duration		
		CPI w/ sales	CPI non-sale	PPI	CPI w/ sales	CPI non-sale	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

CPI regular prices denote consumer prices excluding sales. "Num. of Matches" denotes the number of ELIs matched to 4, 6 or 8-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 4-digit commodity code and then taking a value weighted median across 4-digit commodity codes. All statistics are for the period 1998-2005.

Table 11: Absolute Size of Price Changes

Major Group	Weight	Regular Prices			Sales			All Prices
		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

The sample period is 1998-2005. "Regular prices" denote 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.

Table 12: Frequency of Price Change for Sales Filters 1998-2005

	No Substitutions		w/ Substitutions	
	A	B	A	B
Window:				
1 month	13.3	15.3	14.7	16.4
2 months	12.5	14.2	14.7	15.5
3 months	11.9	13.4	14.1	15.1
4 months	11.4	12.5	13.6	14.7
5 months	11.4	12.5	13.3	14.6
Price Changes	19.4		20.5	
Reg. Price Changes	8.7		10.9	
Reg. Price Ch. + Clear	10.7		13.0	

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 ELI's 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 1 and 5 months. See appendix A for a detailed description of the sale filter algorithm.

Table 13: Regressions of Frequency and Size of Consumer Price Changes on Inflation


Dependent Variable	Regular Prices		Prices	
	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)

The table reports the results of regressions of the median frequency and absolute size of log price increases and decreases at the ELI level on the aggregate CPI inflation rate (log change over 12 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 in Arellano (1987), where the standard errors are clustered by year. A star denotes significance at the 5% level.


Table 14: Regressions of Frequency and Size of Sales on Inflation

Dependent Variable	1988-1997	1998-2005
Consumer Price ELI Level:		
Frequency of Sales	-0.34 (0.17)	-0.24 (0.20)
Size of Sales	-0.19 (0.44)	0.45 (0.43)

The table reports the results of regressions of the frequency and absolute size of sales at the ELI level on the aggregate CPI inflation rate (log change over 12 months). 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 in Arellano (1987), where the standard errors are clustered by year. A star denotes significance at the 5% level.



Panel A




Panel B

Figure 1 : Construction of Price Change Variables With and Without Sales


Each panel reports the first 10 observations for a hypothetical price series. The top row of each panel records the values of the sales flag for the 10 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.

Figure 2: Inflation and the Frequency of Regular Price Change for Consumer Prices



The figure plots the annual evolution of the weighted median frequency of regular price increases and decreases along with the CPI inflation rate.

Figure 3: Inflation and the Size of Regular Price Changes for Consumer Prices



The figure plots the annual evolution of the weighted median absolute size of log regular price increases and decreases along with the CPI inflation rate.




Figure 4: Frequency of regular price increases and decreases in the data and in the model.






Figure 5: Size of regular price increases and decreases in the data and in the model.

Figure 6: Evolution of the Frequency of Sales




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.

Figure 7: Frequency of Regular Price Increases and Decreases by Month for Consumer Prices



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

Figure 8: Seasonality of the Frequency of Sales



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.




Figure 9: 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%.





Figure 10: Hazard Function for Consumer Prices