UNIVERSITEIT GENT

FACULTEIT ECONOMIE
EN BEDRIJFSKUNDE

TWEEKERKENSTRAAT 2
B-9000 GENT
Tel. : 32 - (0)9 – 264.34.61
Fax. : 32 - (0)9 – 264.35.92

WORKING PAPER

Price rigidity in Europe and the US: A comparative analysis using scanner data

Benjamin Verhelst
Dirk Van den Poel

November 2010

2010/684

D/2010/7012/55

Price Rigidity in Europe and the US: A Comparative Analysis Using Scanner Data *

Benjamin Verhelst † Dirk Van den Poel ‡

November 17, 2010

Abstract

This paper uses scanner data from two large retailers to offer new insights into the extent of price rigidity in Europe and the US. Recent empirical research in this field has made extensive use of monthly data to study price stickiness and to control for the impact of temporary sales. We show that the use of monthly data is potentially highly misleading. We employ scanner data in (bi)weekly frequency and highlight the importance of high frequency data in studying price rigidity. Regular prices show roughly the same degree of flexibility in Europe and the US, in line with recent empirical research, when we study monthly price series derived from our high frequency scanner data. This finding collapses, however, when the original scanner datasets in higher base frequency are examined. Regular prices are then far more flexible in the US than in Europe. This result is robust to the type of sales filter that we apply and the statistic used to capture price rigidity.

JEL: C33, D4, E3, L66
Keywords: price setting, scanner data, frequency of price change, sales filtering

*The authors would like to thank Freddy Heylen, Maarten Dossche and Punnoose Jacob for valuable comments. We acknowledge support from the Belgian Program on Interuniversity Poles of Attraction, initiated by the Belgian State, Federal Office for scientific, technical and cultural affairs, contract UAP No. P 6/07. Benjamin Verhelst acknowledges financial support from the Flemish Fund for Scientific Research (FWO). Any remaining errors are ours.

†Corresponding author, SHERPPA and Department of Social Economics, Ghent University, Benjamin.Verhelst@UGent.be

‡Department of Marketing, Ghent University, Dirk.VandenPoel@UGent.be

1 Introduction

A large literature documents the persistent effects of nominal shocks on real output and inflation (Peersman, 2004; Christiano *et al.*, 2005). Given the key role of price rigidity to explain this persistence, many micro-based models of price setting have been developed for macroeconomic models. The empirical assessment of the price-setting process at the micro level, however, has until recently been very limited.

This paper makes an empirical contribution to the literature in this field by estimating price stickiness in Europe and the US using high frequency supermarket scanner data. Our results qualify earlier findings of Bils and Klenow (2004), Dhyne *et al.* (2005), Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008), and underscore the importance of the frequency at which price quotes are available in estimating a price stickiness parameter. We show that neglecting this frequency issue can lead to spurious results and potentially blur cross-sectional comparisons of price rigidity.

Until the start of the 2000s, most micro-based price studies were based on relatively small samples. Although these studies show that price changes are infrequent, they offer limited understanding of micro price setting at large because they concentrate on a small number of goods sold in one geographical area (Cecchetti, 1986; Lach and Tsiddon, 1992; Kashyap, 1995). A major step forward in the empirical study of price setting at the micro level was made in 2003, when the Eurosystem of Central Banks set up the Inflation Persistence Network (IPN). This led to the collection of large datasets of consumer and producer prices from various euro area countries, used to construct the Consumer and Producer Price Indices (CPI and PPI). Dhyne *et al.* (2005) summarize the results obtained in a number of national studies and build upon these to deduce the specifics of the price setting process in the euro area.

At roughly the same time of the IPN formation, Bils and Klenow (2004) compiled similar price records underlying the computation of the CPI for the United States, collected by the Bureau of Labor Statistics (BLS). Consequently, it became an obvious research objective to compare the price flexibility in Europe and the US based on large datasets of consumer prices, as differences in price stickiness call for different policy measures [1]. As sticky prices raise the output cost of stabilising inflation, a central bank needs to take the stickiness of prices into account when deciding which weights to assign to output and inflation variability in the monetary policy reaction function.

[1] Comparable studies with similar datasets have been conducted by Baharad and Eden (2004) for Israel, Masahiro and Yumi (2007) for Japan, Karadi and Reiff (2007) for Hungary, Coricelli and Horvath (2008) for Slovakia and Gagnon (2009) for Mexico.

Dhyne *et al.* (2005) find an average euro area monthly frequency of price change of around 15 percent, and a weighted median price duration of 10.6 months. The results of Bils and Klenow (2004) for the US, based on a comparable set of product categories, show an average monthly frequency of price change in the neighbourhood of 25 percent and a weighted median price duration of 4.6 months. Based on these estimates, which are calculated from millions of monthly price observations over several years for a broad array of products, euro area prices appear far stickier than US prices. The observation of higher micro price rigidity in Europe than in the US was fully consistent with earlier macroeconomic estimates showing a smaller reaction of inflation to changes in real marginal cost in Europe (Gali *et al.*, 2001).

Nakamura and Steinsson (2008) analyze the frequency of price change more thoroughly and emphasize the importance of product turnover and especially temporary sales promotions in US CPI data. If the associated temporary price changes are believed to be orthogonal to macroeconomic conditions, then regular sales-filtered prices are more interesting than posted prices when studying macroeconomic issues. Consequently, an adequate filtering of the product turnover and sales episodes from the posted price series is required. Whereas Bils and Klenow (2004) eliminate sales by using the proportion of sales in the overall data, Nakamura and Steinsson (2008) correct for sales directly using the detailed price level information of each separate product [2]. After filtering out price changes associated with product turnover and temporary sales promotions, they obtain a frequency of price change in their US data that is much closer to the ones found in comparable euro area data. Hence, the gap in price stickiness between Europe and the US that is found in the posted price series could simply reflect that US retailers use sales promotions more often than their European counterparts.

Klenow and Kryvtsov (2008) look at this issue in a slightly different way, as they consider price changes associated to product turnover to be price changes fair and square, and therefore only filter out temporary sales promotions. On top of that, they also take into account non-adjacent price quotes, whereas Nakamura and Steinsson (2008) only consider adjacent price observations. Nonetheless, the results of Klenow and Kryvtsov (2008) point to the same conclusion as obtained by Nakamura and Steinsson (2008). Whereas the probability of a posted price change is higher in the US than in the euro area, the probability of a regular price change is highly similar in both regions.

[2] The BLS Commodities and Services Substitution Rate Table that is used by Bils and Klenow (2004) to study price rigidity does not contain information on the magnitude of price changes during sales periods. It only contains the share of price quotes that involve some change in price. Because the BLS collects prices net of sales and other promotions, Bils and Klenow (2004) only have the proportion of sales in the overall data at their disposal to construct the regular price series.

Besides CPI data, there is another important source of micro data that offers a vast amount of pricing points over time and for a broad range of products, that being scanner data. Unlike CPI data, which are gathered by government officials, scanner data are obtained through the scanning process of product-specific barcodes that is in common use in supermarkets and drugstores. Scanner data therefore contain all price quotes of items that are actually being bought by consumers. The retailing sector is huge and consumer-retailer interactions mimic conditions present in many other types of purchases, making supermarket data highly valuable to study consumer and producer behaviour (Chevalier *et al.*, 2003; Dossche *et al.*, 2010) [3].

Although CPI data are in general much broader in terms of product coverage than scanner data from any given supermarket, the latter have a very important asset in that they are generally measured on a weekly basis, whereas CPI data are measured monthly, at best. Our purpose is to use the higher frequency of scanner data to revisit the results on price stickiness obtained in CPI-studies. We analyze price records in their base frequency, i.e. the retailer does not change the price of its products in between data points. We compare the frequency of price change that we obtain from our datasets at base frequency with the ones from datasets at monthly frequency that we artificially derive from the original high frequency data. As such, we are able to gauge the effects of short-term pricing tactics on the stickiness of prices and on the validity of the comparison between Europe and the US.

Our results show that posted and regular prices in the US are far more flexible than in the euro area when we use the full amount of information contained in our scanner datasets. However, switching from base frequency data to derived monthly data, the difference largely disappears, leaving the frequency of price change at a similar level. This dichotomy qualifies the results of Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008) and highlights the importance of the frequency at which price quotes are available. Using monthly data to study the impact of temporary sales on the flexibility of prices is potentially highly misleading, and leaves a comparison between regions severely biased in case of differences in short-term pricing tactics.

The remainder of this paper is organized as follows. In section 2, we describe the scanner data and give some descriptive statistics. Section 3 describes the estimation methodology and presents the main results with respect to the frequency of price change. The robustness of these results with respect to the sales filtering method is ascertained. In section 4, we decompose price changes into increases and decreases, and study their frequency. Section 5 looks at the size of price changes. Section 6 concludes.

[3] Scanner data have been used extensively in macroeconomic research, see e.g. Levy *et al.* (1997), Dutta *et al.* (1999), Eichenbaum *et al.* (2008), Midrigan (2010) and Kehoe and Midrigan (2010).

2 Data

We use two scanner datasets, both originating from large and geographically dispersed retailers. For the US, we use the publicly available scanner data from 86 stores of Dominick's Finer Foods, an important supermarket chain in the Chicago metropolitan area. The time span of these data runs from September 1989 to May 1997. The data have already been used extensively for similar research purposes (Peltzman, 2000; Chevalier *et al.*, 2003; Rotemberg, 2005; Ray *et al.*, 2006; Midrigan, 2010; Kehoe and Midrigan, 2010). For the euro area, we use unique scanner data from 6 stores of an anonymous European retailer. These data run from January 2002 to April 2005. They have only been used before by Dossche *et al.* (2010) to test the existence of the kinked demand curve and to estimate its curvature.

The different timing of the data should not pose any problem arising from varying growth or inflation rates. During a recession, price setters might adapt their prices more swiftly to the economic environment, because they become more attentive and revise their prices more often during tough times (Klenow and Malin, 2010). In this respect, it is important to note that our datasets include both prosperous and recessionary episodes, so that the comparison between them is not affected by a possible increase in the frequency of price change during an economic downturn [4].

Inflation can also spur additional price changes, as price setters want to keep up with the changing macro price level. However, both our data episodes are situated in a low inflation environment, the average aggregate inflation rate during the sample period being 3.43 percent in the US and 2.60 percent in Europe. This small difference can hardly explain any potential gap in price stickiness. Although the composition of price increases and decreases changes with inflation, these movements almost entirely cancel out leaving the frequency of price change unaffected (Dhyne *et al.*, 2005). Hence, low to moderate inflation only shows up at the intensive margin, i.e. the size of price changes, not at the extensive margin (Klenow and Kryvtsov, 2008). Only during high inflation episodes, the frequency of price change might be affected as well (Gagnon, 2009; Wulfsberg, 2009).

An important aspect of these datasets is that they are available in their base frequency. We know that the retailers do not change their prices in between data points. For the US data, the base frequency is weekly, the European data are available at biweekly frequency. Consequently, we do not lose any price information in between price quotes, so no measurement error occurs when estimating price change frequencies.

[4] The US economy suffered from a downturn in 1990-1991, whereas the European data include the downturn of 2001-2003.

We restrict the scanner datasets by selecting only those product categories that are represented in both, as to make them comparable. This is very important because the frequency of price change is known to be extremely heterogeneous across products, implying that a proper assessment and comparison of price stickiness must be based on similar product baskets (Bils and Klenow, 2004; Dhyne *et al.*, 2005; Nakamura and Steinsson, 2008). By doing so, we end up with 16 product categories containing 9261 products for the US dataset, and 18 categories containing 1270 products in the European data [5]. In appendix I, we list all product categories that we include in our European and US datasets, with their respective number of products.

All categories contain a large number of fairly homogeneous products, defined at a very detailed level by their Universal Product Code. One UPC will for example be connected to a package of 20 scrolls of white toilet paper of a certain brand. The raw data consist of individual price trajectories, i.e. sequences of elementary price quotes for a specific product in a specific outlet.

3 Frequency of Price Change

In this section, we present statistics on the frequency of price change for both the US and European retailer. An important aspect of the analysis is the type of price series that we study, as previous research on price setting behaviour using micro data has shown that excluding promotional prices from the data has a major impact on inference about the stickiness of prices. Not surprisingly, there is an ongoing debate in the literature about how to define a sales promotion and whether one should treat regular and sales prices asymmetrically (Eichenbaum *et al.*, 2008).

3.1 Temporary sales

Before going into the identification and filtering of sales prices and the comparison of results on the frequency of price change for posted and regular price series in Europe and the US, we digress on the phenomenon of temporary sales that is apparently very important for price setters, especially in retailing. Bils and Klenow (2004) define a sales price in general as a price that is temporarily lower than the regular price and available to all customers. That begs the question which prices should be taken into account when thinking about the macroeconomic implications of price rigidity and calibrating macroeconomic models. Assuming that temporary sales have macro content, they have to remain part

[5] The difference in the number of product categories stems from the fact that some of them are defined more broadly in US data compared to the classification scheme of the euro area retailer. The product coverage of both our constructed datasets is however highly similar.

of the price series in order to find the appropriate measure of price stickiness for macro policy purposes. If on the other hand they are considered orthogonal to macroeconomic aggregates, they do not represent a actual form of rigidity and they can be left out. Standard staggered price models focus on permanent price changes and either drop temporary price changes from the data, the temporary-changes-out approach, or leave them in and treat them as permanent, the temporary-changes-in approach. Kehoe and Midrigan (2010) propose a menu cost model that includes motives for both permanent and temporary price changes by including a parameter on the technology of price adjustment.

To decide about the macro content of sales prices, the obvious thing to do is to look at the motives of price setters to offer temporary sales, a topic that has been studied extensively in the industrial organization literature. Adapting to idiosyncratic shocks that hit the firm is a first straightforward motive to set a sales price. In other words, it is a simple way for the firm to adapt prices continuously to cyclical fluctuations in its costs or demand. A second motive for temporary sales that is often cited is intertemporal price discrimination, i.e. retailers lower their markups and prices of a certain product in periods of high price elasticity of demand for that product, as to maximize the opportunity to gain market share (Kehoe and Midrigan, 2010). Loss-leader behaviour on behalf of competing retailers is a closely related motive for temporary sales, where they discount items in high relative demand, even if this does not coincide with a period of high aggregate demand, as this signals to consumers that all unadvertised products are sold at their reservation price (Chevalier *et al.*, 2003).

All three of these sales motives can broadly be described as orthogonal to macroeconomic aggregates. Consequently, if these motives are believed to be important, using regular prices to calibrate macroeconomic models is advocated. Additionally, due to the transience of price adjustment associated with sales, a given number of price changes due to sales yield much less aggregate price adjustment than the same number of regular price changes, further strengthening the case for studying regular price series. However, there might also be motives for temporary sales that do contain macro content, for example when inflation has been low or when excess inventory builds up (Klenow and Kryvtsov, 2008). In the latter case, clearance sales serve as a perfect method for inventory management in case of shifts in aggregate demand or unpredictable shifts in tastes (Nakamura and Steinsson, 2008). When inventory is large, reducing inventory reduces costs, so that the expected marginal cost of selling an item is lower in periods of high inventory (Hosken and Reiffen, 2004). In this case, the magnitude and duration of temporary sales do respond to macroeconomic shocks, hence they should not be excluded from the analysis (Bils and Klenow, 2004).

As the macroeconomic relevance of temporary sales is not the main issue of the paper, we will provide estimates of the frequency of price change for both posted and regular prices. This has the advantage that the results can be used in a flexible way in a variety of economic settings, and are easy to compare with previous studies. An additional advantage is that statistics for both posted and regular prices are important if you want to calibrate a hybrid model in which firms face a different cost for a temporary versus permanent price change, as in Kehoe and Midrigan (2010).

To filter temporary sales from the posted price series, we need a specific definition of what constitutes a sale before we can design a filter algorithm. We therefore have to decide on three different dimensions of a sales definition. The first aspect to be considered is the symmetry of a sale, i.e. should the price before and after the sales episode be the same or is it allowed to be higher after the sale? The second facet of the definition is the maximum number of periods that a sale can be in effect. The third aspect to consider is if a sale price should remain fixed or is allowed to vary during the sales episode. The sales filter can then be parameterized based on the choice with respect to these three dimensions of the definition.

Bils and Klenow (2004) define a temporary sale quite simply as a price change in one month that is changed back to the original price the next month. To be viewed as a sale, a price decrease thus has to be followed by an identical price increase in the next period, i.e. only symmetrical V-shaped price changes lasting one period are classified as sales. Hosken and Reiffen (2004) define a sale in a similar, straightforward way as occurring if the price falls by at least some fixed percentage (e.g. 10 or 20 percent) between periods t-1 and t and then rises by at least that percentage between periods t and t+1. The difference with the definition of Bils and Klenow (2004) is that both symmetrical and asymmetrical V-shaped price changes are covered by this definition. A product is thus recorded as being on sale in month t if the prices in month t-1 and month t+1 are both significantly higher than the price charged in month t, but not necessarily the same. It counts brief discounts followed by changes in the regular price as sales. Price decreases that last more than one month are again not treated as sales. Campbell and Eden (2007) identify a sale in a similar but slightly more restrictive fashion as a price decline of 10 percent or more in a given week that the store completely reverses within two weeks, instead of one month. All prices between the initial decline and the reversal are then branded as sale prices. The more restrictive nature of this sale classification is only possible because they have weekly instead of monthly data.

Midrigan (2010) also uses weekly price quotes but defines a temporary sale in a less restrictive way as a price decrease of any size that is reversed in one of the four weeks following the original price cut. This definition not only covers V-shaped price changes but also gradual price decreases, provided these are eventually followed by a price increase after at most four weeks following the first price cut. The price is therefore allowed to vary during the sales episode. A price decrease that lasts more than four weeks is always treated as a regular price change, even if it is reversed afterwards. This filter artificially introduces a new sale in case a price cut is gradually reversed. To correctly identify and filter out all sales, the algorithm has to be repeated three times in order to eliminate sales that have been gradually implemented.

To visualize this peculiarity of the filtering process, we show in figure 1 a fictitious example of a price pattern in which a sale is gradually reversed. After the first loop of the filtering algorithm, the sale price at time 4 is replaced with the last observed regular price. We now have a sale price at time 5, which will be filtered out when the algorithm is repeated a second time. This in turn generates a new sale at time 6 and the procedure continues until each successively introduced sale is filtered out. The number of times the algorithm has to be repeated to correctly filter out all sales depends on the maximum duration of a temporary sale, in this case four periods.

We choose to apply four different sales filters to the posted price series, all similar to the one used by Midrigan (2010). All our filters define a sale as a price decrease of any size that lasts a maximum of four weeks following the initial price cut. As a robustness check, we also tested several filters that allow sales to last six or eight weeks, but the results were not significantly different from the ones that we will present below. Apparently, the bulk of temporary sales are very short-lived. Capping off the length of a sales episode at four weeks, we then design a symmetric and an asymmetric filter for the case where we impose a fixed sales price for the duration of the sale on the one hand, and for the case where the sale price is allowed to vary during the sale on the other hand.

We thus end up with four different filtered series. When presenting the results, we will refer to these four filters as *symmv4fix*, *asymmv4fix*, *symmv4flex* and *asymmv4flex*, respectively. If a price quote is defined as a temporary sale by the filter of our choosing, we replace this sale price with the last observed regular price before the sale. Intuitively, we would suggest the asymmetric filter that allows for a varying sales price, i.e. *asymmv4flex*, as the one that most effectively distills a regular price from the posted price series. The latter also most closely resembles the one of Midrigan (2010). To visualize the effect that this sales filter has on the price trajectory of an arbitrary product, we show in figure 2 the posted and regular price series of a product in the laundry detergent category of our European scanner dataset over a time period of 76 biweeks. The dotted line gives the posted price, whereas the full line signifies the regular price. The difference between both

Figure 1: Filtering out a gradually reversed temporary sale

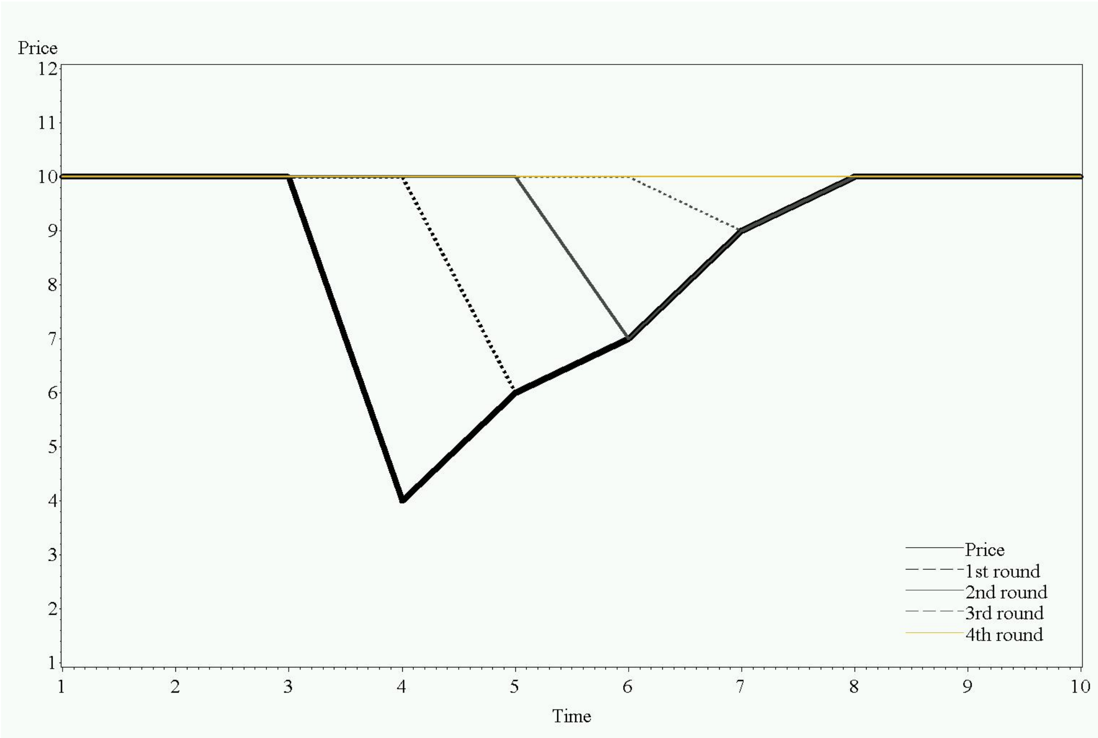

clearly visualizes the temporary sales. In this case, the majority of sales are one-period symmetric V-shaped sales. The pattern shows that the typical retail product is at its regular price most of the time, and deviations from that regular price are downward and short-lived (Hosken and Reiffen, 2004).

Whereas Klenow and Kryvtsov (2008) fill up gaps in their price series, we choose to work with adjacent prices only, following Nakamura and Steinsson (2008). Another difference in their analysis is the problem of voluntary and forced product substitution that is typical for CPI data [6]. Klenow and Kryvtsov (2008) include price changes that are associated with forced product turnover in their analysis, whereas Nakamura and Steinsson (2008) only compare regular prices in between item substitutions, hence excluding all substitution-related price changes. In our scanner data, we can steer clear of this

[6]Voluntary product substitution originates from item rotation in the data set, e.g. each product could be included in the dataset for a period of five years before being replaced, or when the statistical office decides to revise its sample of products because the consumption habits have changed (Dhyne *et al.*, 2005). Forced product substitution occurs when a product is no longer sold and the government official who collects the price quotes chooses a replacement item, in many cases a newer version of the discontinued product.

Figure 2: Posted and regular price series

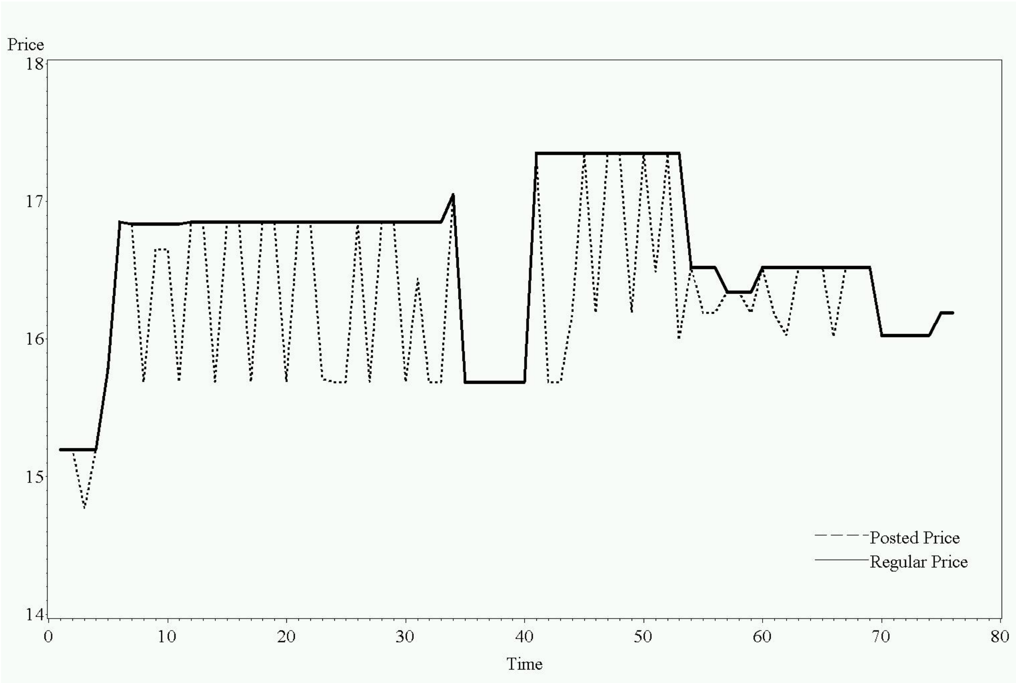

discussion, as there is no voluntary product substitution due to item rotation, and forced product substitution leads to a new price trajectory attached to a different Universal Product Code.

3.2 Frequency vs. Duration Approach

As to make our results comparable with previous research, we will capture price rigidity by the duration of a price spell in months [7]. There are however two approaches to obtain this price stickiness parameter. The duration approach measures the duration of price spells directly from the data, whereas the frequency approach aims at calculating the frequency of price change which can then be inverted to obtain the duration of a price spell [8]. We will follow the latter approach for two reasons. First of all, gaps in price trajectories are a serious concern in the duration approach, as it is impossible to derive the spell length directly if price quotes are missing. One way to tackle this problem is to carry forward the last available price to fill in gaps, but this generates an upward bias for the duration of price spells (Baudry *et al.*, 2004; Klenow and Kryvtsov, 2008). In

[7] A month is defined here as a period of exactly four weeks. Our results will therefore slightly overestimate duration. This can however easily be taken into account if one wants to compare our results to the ones obtained in earlier studies that work with monthly data.

[8] See Baudry *et al.* (2004) and Veronese *et al.* (2005), among others, for an analysis of both approaches.

the frequency approach, a long and uninterrupted span of time series is not a necessary condition to calculate the frequency of price change. This is a very important asset of the frequency approach in our setting, as scanner data usually contain a reasonable amount of gaps in their price series, for example due to stockouts.

Secondly, the duration approach suffers from serious censoring issues that complicate the measurement of price spell durations. The problem of censoring reflects the fact that in the process of price collection the true time of beginning/ending of the first/last price spell might not correspond to the one observed in the dataset, as it comes before/after the first/last price observation, leading to a downward bias in the average duration of a price spell when you ignore censoring (Veronese *et al.*, 2005). Dropping the first and last price spell is not a satisfactory solution for the problem, because these spells are more likely to be long, again implying a downward bias in the average duration of a price spell. In the frequency approach on the other hand, no explicit treatment of censoring is called for, in turn allowing us to use all available information in the dataset as only observations that involve transitions from or to unobserved prices have to be dropped. Appendix II details the computation process of the different statistical measures that we apply to calculate the frequency of price change.

In order to obtain estimates for the duration of a price spell using the frequency approach, we need to invert the estimates of the frequency of price change. In this respect, it is important to note that we use data in their base frequency. As a consequence, we can calculate the duration of a price spell simply by inverting the frequency parameter, call it λ. The implied duration of a price spell therefore equals $1/\lambda$, expressed in time units corresponding to the intervals at which price quotes are available. Working with data that are not in base frequency complicates the picture, because if prices can change at any point in between data points, then the instantaneous probability of a price change is $-ln(1-\lambda)$ and the implied mean time between price changes is $-1/ln(1-\lambda)$ (Bils and Klenow, 2004). This is the way frequencies and durations are linked when using e.g. CPI data, and it is the formula that we will use when we work with our derived monthly scanner data.

3.3 Parameter Choice

There are a number of statistics that can be used to capture the implied duration of a price spell. We have to take two important aspects into account when choosing a statistic. The first one is the choice between the mean or median duration. Both represent a valuable measure of price stickiness, although their value can be significantly different depending on the structure of the data. As we can see in figures 3 and 4, the distribution of price spell durations in both the European and US data is very right skewed, contain-

Figure 3: Distribution of biweekly price spell duration (EU)

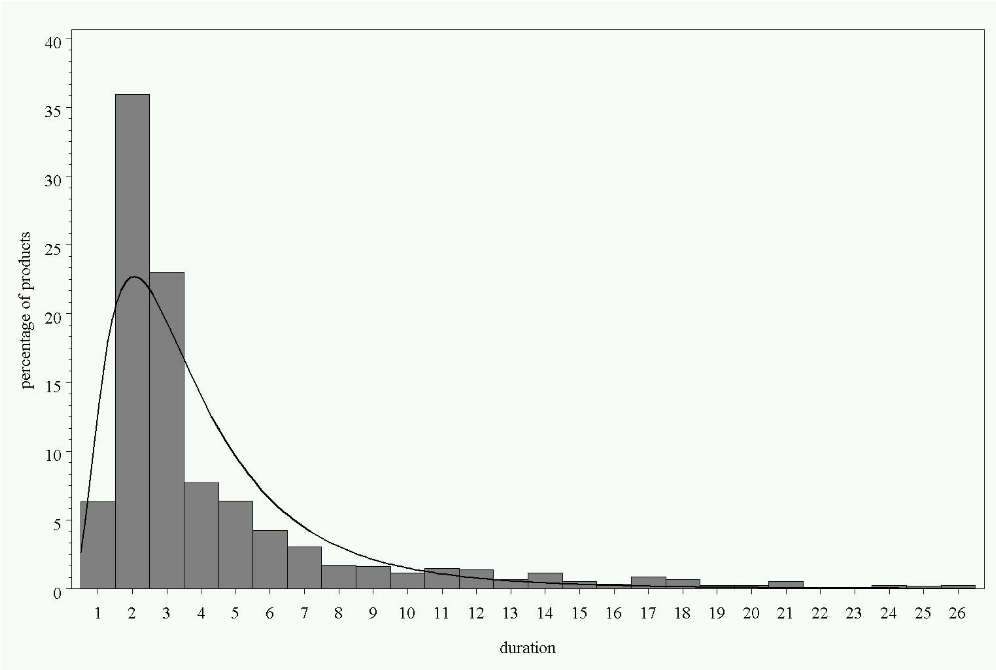

ing a lot of very short price spells and a long tail at the right with outliers. Consequently, the median duration will be much lower than the mean. We expect the skewness of the distribution to diminish when we gradually filter out sales, bringing the mean and median durations closer to each other. Nonetheless, in an environment with this type of uneven distribution, the median captures the true amount of rigidity better than the mean.

On top of that, simply inverting the mean frequency will not give us the true mean duration, because of a combination of this severe product heterogeneity and Jensen's inequality (Dhyne *et al.*, 2005). The latter states that, given a random variable X and a convex function f, $E[f(x)] \geq f(E[X])$. Duration being a convex function, we can rewrite the general rule as $E[1/\lambda] \geq 1/E[\lambda]$. In words, Jensen's inequality states that the inverse of the mean frequency is always smaller than the mean over all inverse frequencies for each product, unless all products have the same duration. The latter captures the true average duration of a price spell, whereas the former is a pseudo-average duration. The difference can be sizeable, e.g. Dhyne *et al.* (2005) show that in a log-normal case with a variance of the distribution of duration equal to unity, the average duration is close to three times higher than the pseudo-average duration.

Figure 4: Distribution of weekly price spell duration in weeks (US)

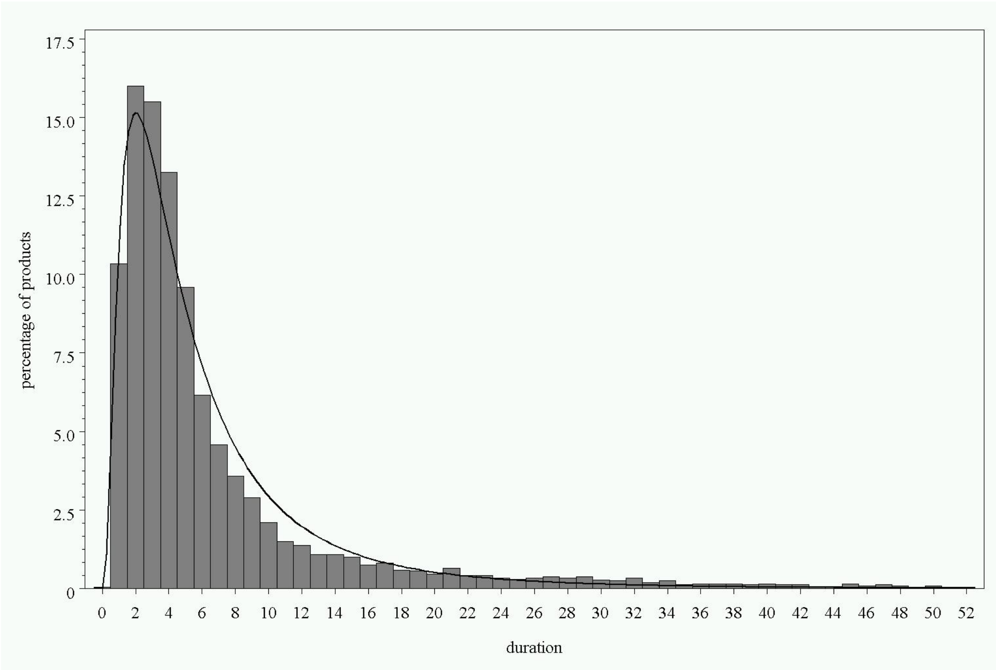

In figures 3 and 4, we overlay the histogram of duration with the log-normal distribution where the threshold parameter θ is assumed to be zero, and the scale and shape parameters are estimated from the data using maximum likelihood. The histogram of duration can clearly be captured by a log-normal distribution. Taking all of this into account, we therefore use the median as our preferred statistic, although we checked the means as well as a robustness check of our results and the conclusions of our analysis remain unaffected (results available upon request).

A second aspect that we have to take into account is the level of aggregation at which we analyse our data. We have two workable levels of aggregation at our disposal, namely the product and product category level. At the product category level, the number of frequency parameters from which to calculate medians is rather limited, putting the soundness of the estimates at risk. We therefore propose to calculate medians at the product level. To this end, we compute the mean frequency of price change for each separate product across time and stores, take the median of all those frequency parameters and invert it to obtain the median implied duration of a price spell at the product level. Appendix II shows how we compute the mean frequency of price change for a specific product i across n stores and τ time periods.

The results at the product category level are calculated in the same way, computing the mean frequency of price change for each product category across time and stores, then taking the median over all categories and inverting it to obtain the median implied duration of a price spell at the product category level. We will present the latter results as a robustness check, because the statistics at the product level could potentially be biased as they overweigh categories which have a large number of very similar products (Eichenbaum *et al.*, 2008).

3.4 Results

The results in table 1 show that the posted price of the median product lasts slightly longer in our European data compared to the US, 1.5 versus 1.3 months respectively. In other words, half of prices last less than 1.5 and 1.3 months in the European and US scanner data, respectively. In the European data, after every period of two weeks, 34 percent of products witness a price change, whereas in the US, 18.5 percent of products see their price change every other week. When we look at the median product category, the difference in price spell duration between Europe and the US is more pronounced, 1.6 months versus 1.0 months respectively.

If we take out sales using our four preferred filters, we see that also regular prices are clearly more flexible in the US scanner data, irrespective of the strictness of the filtering method. With the most extensive method, filtering out all symmetric and asymmetric sales with a fixed or flexible sales price of up to 4 weeks, the US regular price for the median product lasts 3.3 months, whereas the median duration in Europe amounts to 6 months. At every biweekly price revision, only 8.4 percent of European products witness a regular price change, whereas 7.6 percent of US regular prices change at the weekly price revision. The difference in the extent of price stickiness is also present at the product category level, albeit less pronounced, 3.3 months versus 4.5 months. Using the full informational content of our pricing data, we can thus conclude that both posted and regular prices are more rigid in our US data compared to the data of the European retailer.

As expected, we find that the mean duration of posted prices is three to four times higher than the pseudo-mean duration for both our datasets, a direct consequence of the extensive product heterogeneity in combination with Jensen's inequality. At the product category level, the difference is of course much smaller because the more disaggregated are the data, the greater is the level of heterogeneity, and the greater is the difference between mean and pseudo-mean duration. If we look at the filtered data at the product level, the average duration is still twice as high as the pseudo-average duration, pointing to the fact that the hetereogeneity in the frequency of regular price change is still sizeable, albeit smaller than for posted prices.

Table 1: Median implied duration of a price spell in months, base frequency

	Product level		Category level	
	EU	US	EU	US
No filter	1.5	1.3	1.6	1.0
Symmv4fix	5.0	2.0	3.1	1.6
Asymmv4fix	5.6	2.3	4.0	2.0
Symmv4flex	5.2	2.7	3.3	2.4
Asymmv4flex	6.0	3.3	4.5	3.3

Note: Symmv4fix filters all symmetric V-shaped sales of up to 4 periods with a fixed sales price. Asymmv4fix filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed sales price. Symmv4flex filters all symmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price. Asymmv4flex filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price.

Compared to previous research on the extent of price stickiness, our estimates of duration are very low. The most important reason to explain the high frequency of price change is that we study retail prices. This implies that some very sticky classes of consumer goods are not included in our sample, mainly services but also apartment rents or restaurant meals for example [9]. Another factor is the type of retailer, as large supermarkets like the ones we study change their prices more swiftly compared to corner shops (Dhyne *et al.*, 2005). As the duration estimates of regular prices are closer to what has been found in the literature compared to the results for posted prices, our retailers also seem to use temporary sales more extensively compared to the average in the economy.

In table 2, we present the results of a similar analysis on monthly price series that we directly derived from our high frequency scanner data by withholding only the first price observation of each month. As such, we mimic the structure of typical monthly data sets like the one used to form the CPI. When we now look at the median duration of a price spell for posted prices, there is no significant difference between Europe and the US. This result holds both at the product and the category level. The conclusion that US retail prices are more flexible than their European counterparts therefore collapses in this setting. On top of that, prices in general appear to be more sticky when monthly data are studied. This is a direct consequence of the fact that temporary sales are very

[9] Services display much more price rigidity due to lower volatility of consumer demand and input costs. The latter is due to a higher share of labour as an input factor, the cost of which is much more stable than that of intermediate goods (Bils and Klenow, 2004; Dhyne *et al.*, 2005).

short-lived and hence not picked up in monthly data. It is therefore an inherent problem of price data that are not available in their base frequency, as they do not provide direct evidence about the critical issue of how many temporary price changes happen in between data points (Kehoe and Midrigan, 2010).

Switching our attention to regular prices, it is important to note that in this stage of the analysis, we filter sales at a monthly frequency, i.e. we first transform the high frequency data into monthly price series and only then filter out sales. Consequently, there is no longer a difference between fixed and flexible sales episodes, as temporary sales of up to 4 weeks last just one period when filtering sales at a monthly frequency. For the median product, we see that also regular prices display more or less the same flexibility in our European and US data, duration being only slightly higher in Europe than in the US. For the median product category, prices appear to be even more rigid in the US than in Europe, albeit only marginally so, contrary to our conclusions for the data in base frequency. This clearly signifies a non-negligible amount of measurement error when using monthly data. Comparing regions with respect to the stickiness of prices based on this kind of data should therefore be performed with due caution.

Table 2: Median implied duration of a price spell in months, monthly frequency (time aggregation before filtering)

	Product level		Category level	
	EU	US	EU	US
No filter	2.2	2.3	2.1	2.1
Symmv4	4.8	4.2	3.1	3.5
Asymmv4	5.2	4.7	4.0	4.2

Note: Symmv4fix filters all symmetric V-shaped sales of up to 4 periods with a fixed sales price. Asymmv4fix filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed sales price. Symmv4flex filters all symmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price. Asymmv4flex filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price.

Our data also offer the possibility to filter out sales at weekly/biweekly frequency and then time-aggregate the filtered weekly data into monthly observations. Midrigan (2010) uses this method to find an estimate for price stickiness that can be used to calibrate a macroeconomic model, because time-aggregating first and then eliminating sales, as we have done to obtain the results of table 2, can produce spurious price changes if stores periodically put their prices on sale, at regular intervals. Consequently, the latter method is to be preferred to the one above, but take into account that this is not possible when one starts out with monthly data, say the CPI data that have been used widely to study price stickiness.

Looking at the results in table 3, we see that the duration of a regular price for the median product is lower in the US compared to Europe, as we have found in our high frequency data, but the difference here is smaller in size. The estimate of duration for the European data using our most strict sales filter is 6.1 months, very similar to the one we found in table 1. The estimated duration of the median product in the US data on the other hand, is now 4.7 months, considerably higher than the 3.3 months from table 1. At the product category level, prices appear slightly more rigid in the US than in Europe, contrary to what we obtained using data in base frequency.

We could therefore conclude that the method of filtering out sales at base frequency and then time-aggregating, does a better job in calculating the frequency of price change and the duration of a price spell than the method that does the job the other way around, but that it is still a shaky basis to obtain reliable estimates upon which regional comparisons can be drawn. Apparently, the problem remains that there is a lot of information contained in the short-term price series that gets lost when the high frequency data are aggregated into monthly price series. Time-aggregation therefore biases the estimates of duration.

Our results highlight an important caveat in the recent empirical literature that studies price stickiness using micro data. If the latter are collected at a monthly frequency, a lot of price changes get lost in between data points leading to a potentially severe bias in the estimation of the frequency of price change and the duration of a price spell. A comparison between two regions based on monthly data can lead to highly misleading results in case there is a difference in short-term pricing tactics, as time-aggregation takes away a lot of relevant information to take the latter into account. This caveat not only holds true for posted prices, but also for regular prices because a typical sales filter takes away a lot of variation on the downside of the regular price, but not on the upside. Consequently, retailers that change their prices more actively both upwards and downwards will appear more rigid relative to retailers with a stable price setting policy, when the data being used are not in their base frequency.

Looking at our data, the flexibility of US retail prices lies in their short-term, weekly movements. Because these are absent in the European data as a policy choice by the retailer, studying the data in base frequency leads to significantly different results compared to the same data in monthly frequency. Taking this result to the macroeconomic level, it implies that flexible short-term pricing strategies do not show up in the statistics, and countries where these are extensively used to react vividly to changing market conditions, could therefore mistakenly be classified as rigid.

Table 3: Median implied duration of a price spell in months, monthly frequency (filtering before time aggregation)

	Product level		Category level	
	EU	US	EU	US
No filter	2.2	2.3	2.1	2.1
Symmv4fix	5.4	3.1	3.3	2.6
Asymmv4fix	5.9	3.4	3.8	2.8
Symmv4flex	5.8	4.1	3.5	3.5
Asymmv4flex	6.1	4.7	4.0	4.2

Note: Symmv4fix filters all symmetric V-shaped sales of up to 4 periods with a fixed sales price. Asymmv4fix filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed sales price. Symmv4flex filters all symmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price. Asymmv4flex filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price.

A more insistent short-term pricing strategy in one region compared to another might be due to a number of reasons. First, it can be an optimizing choice on behalf of the retailer in response to more flexible factor markets. Additionally, higher perceived price elasticity of demand, for example due to more competitive product markets with lower wholesale markups, can force the price setter to adapt prices more swiftly in response to price changes at competing retailers. A third potential rationale for adapting prices more promptly could lie in a lower menu cost, for example due to a more technologically efficient, electronic price tagging system.

4 Frequency of Price Increases vs. Decreases

Looking at the median duration of a price spell does not give any insight into the composition of the price changes. Therefore, we break up all price changes into increases and decreases, and table 4 presents the median frequency of both. We only present these results for an analysis of the scanner data in base frequency, as we have shown this to be the only accurate basis to perform frequency estimation. Therefore, the presented

European frequency statistics are biweekly, whereas the US results give the percentage of prices that increase and decrease per week.

Price increases are slightly more frequent than decreases when we look at the posted, unfiltered data. This is true for both our European and US data. In Europe, 17.4 percent of products increase in price at the biweekly revision, whereas 16.5 percent witness a price decrease. For the US, the equivalent frequencies at the weekly revision time are 9.8 percent and 8.7 percent. Interestingly, as the frequency of price change decreases by filtering out temporary sales along our four preferred dimensions, the percentage of price decreases vis-á-vis increases in the retail data falls in both Europe and the US, but they nonetheless remain important for the flexibility of prices.

Table 4: Median frequency of price increases and decreases

	EU		US	
	Up	Down	Up	Down
No filter	0.174	0.165	0.098	0.087
Symmv4fix	0.062	0.042	0.067	0.055
Asymmv4fix	0.056	0.035	0.057	0.049
Symmv4flex	0.059	0.039	0.049	0.038
Asymmv4flex	0.053	0.032	0.041	0.031

Note: Symmv4fix filters all symmetric V-shaped sales of up to 4 periods with a fixed sales price. Asymmv4fix filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed sales price. Symmv4flex filters all symmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price. Asymmv4flex filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price.

These results establish the prevalence of price decreases in retailing, even when temporary sales are filtered out. Hence, we find no support for downward rigidity in the price level. Price decreases are thus not just a reflection of temporary sales episodes, but a robust stylized fact in retail pricing. This confirms earlier results from the literature (Baudry *et al.*, 2004; Dhyne *et al.*, 2005).

5 Size of Price Changes

Besides the frequency of price change, the size of price changes is another dimension of the pricing behaviour of firms. It is not possible to give some sensible insight into the micro foundation of the inflation rate without looking at the frequency and size of price changes in conjunction. Different combinations along those two dimensions can lead to the same inflation rate, as it can be decomposed in the following way (Klenow and Kryvtsov, 2008):

$$\pi_t = \lambda_t^+ .dp_t^+ - \lambda_t^- .dp_t^-$$

with π_t the inflation rate in period t, λ_t^+ and dp_t^+ the frequency and size of price increases, and λ_t^- and dp_t^- the frequency and size of price decreases. In this section, we will look into the size of price changes. Combining this analysis with the frequency of price increases and decreases from section 4 will give some insight into the micro level roots of the inflation rate.

Table 5 displays the mean and median size of price increases and decreases in our European and US scanner data. In appendix II, we show how these size parameters are computed. It is important to note that we compute the size of price increases and decreases as the difference of logarithm, so that the two successive price changes recorded during a temporary sale are equal in absolute terms (Dhyne *et al.*, 2005).

A first point to note in table 5 is that posted price changes are sizable if you compare them with the low inflation rates that are prevalent in western economies during the last decades. In line with the more aggressive short-term pricing strategy with respect to the frequency of price change, we see that the US retailer, which reconsiders its prices more aptly, not only changes its prices more often, but also by more than the European retailer. Focusing on the size of a price increase for the median product, we find it to be 10.6 percent in the US versus 7.2 percent in Europe. The same applies for the median size of a price decrease, with an even higher dispersion of 12.8 percent versus 7.2 percent. These results are in the same ballpark as the ones found in previous research on the size of price changes (Dhyne *et al.*, 2005; Nakamura and Steinsson, 2008; Klenow and Kryvtsov, 2008).

Secondly, we notice that the median size of a price change is lower than the mean, implying that the distribution of price changes is skewed towards small price changes with a limited number of big price changes in the tail. This is especially the case in the US. A third point to note when looking at both the mean and median size of posted price changes is that price increases and decreases are of the same order of magnitude in Europe, whereas price decreases are slightly larger in size than price increases in the US. Turning our attention to regular prices, we clearly see that the size of regular price changes is smaller than the size of posted price changes. In other words, temporary sales are larger in size than regular price changes, a feature that is common to both retailers.

Table 5: Size of price changes

	EU				US			
	Increase		Decrease		Increase		Decrease	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
No filter	9.630	7.208	9.654	7.155	17.963	10.648	19.888	12.783
Symmv4fix	8.480	5.841	8.527	5.585	14.669	7.372	17.251	9.554
Asymmv4fix	7.344	4.927	7.790	4.564	12.460	6.100	13.639	6.720
Symmv4flex	8.536	5.859	8.581	5.580	14.326	6.761	17.336	8.952
Asymmv4flex	7.287	4.853	8.060	4.689	13.055	5.749	16.641	8.058

Note: Symmv4fix filters all symmetric V-shaped sales of up to 4 periods with a fixed sales price. Asymmv4fix filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed sales price. Symmv4flex filters all symmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price. Asymmv4flex filters all symmetric and asymmetric V-shaped sales of up to 4 periods with a fixed or flexible sales price.

The difference between the mean and median size is even more pronounced for regular than for posted price changes, pointing to the fact that the bulk of regular price changes are small. This applies to both price increases and decreases. In the US, the median size of regular price decreases is larger than that of regular price increases. This is most apparent when we also filter out all asymmetric sales, because large price increases that follow immediately after a temporary sales price and overshoot the previous regular price are in these cases partially filtered out, transforming a lot of large price increases into small ones.

Along with the high frequency of price decreases that we showed in section 4, the relatively large size of price changes signifies an important role for large but relatively transient idiosyncratic shocks in the pricing strategy of a typical retailer (Bils and Klenow, 2004). That poses a serious problem for a lot of price rigidity models which assume that price changes only occur in response to aggregate shocks [10]. In an inflationary environment, almost all price changes should then have to be increases (Nakamura and Steinsson, 2008).

Although the mean and median price change are quite sizeable, a lot of price changes are small. To visualize this, figures 5 and 6 show a histogram for the size of posted and

[10]See for example Taylor (1980), Calvo (1983), Caplin and Spulber (1987), Dotsey *et al.* (1999), Chari *et al.* (2000), Mankiw and Reis (2002), Burstein (2006) among others. Golosov and Lucas (2007), Gertler and Leahy (2008) on the other hand propose models that do incorporate a role for idiosyncratic shocks.

regular price changes, respectively, in the US dataset [11]. We see in figure 5 that more than 30 percent of posted price changes is smaller than 5 percent in absolute value, and 19 percent of them is smaller than 3 percent. The histogram for posted price changes is very symmetrical, and is captured quite well by a normal distribution with an estimated mean $\mu = 0,202$ and standard deviation $\sigma = 17,708$. This pattern is mainly due to the fact that symmetrical V-shaped sales are quite common, and the price decreases and increases due to this type of temporary sales are equal in size.

Figure 5: Size distribution of posted price changes (US)

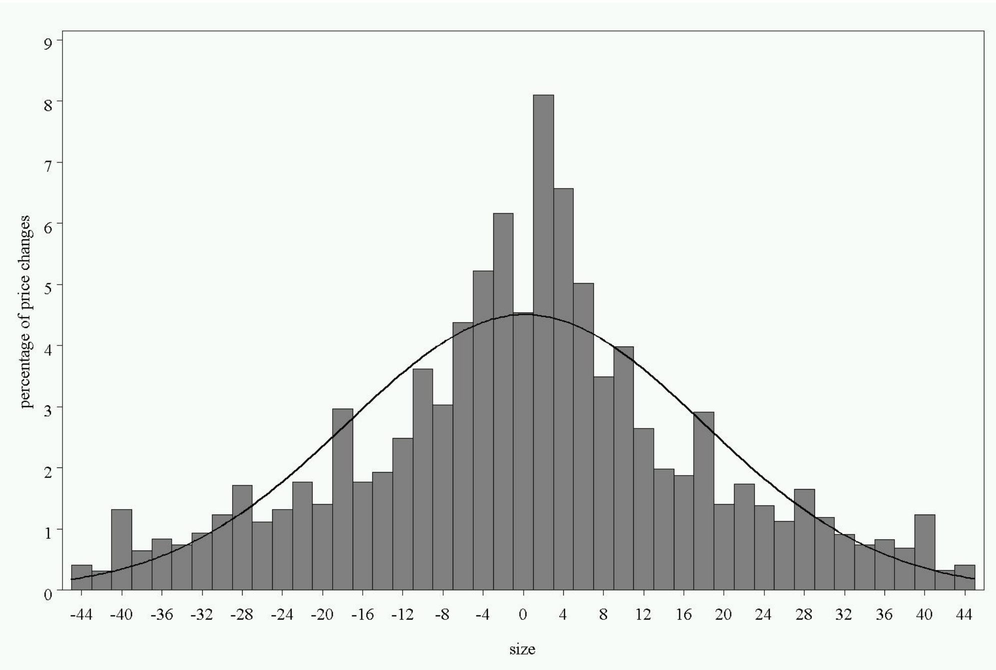

Figure 6 provides a histogram of the size of regular price changes in the US dataset. Small price changes are even more prevalent in the regular price series, thereby confirming that price changes attached to temporary sales are larger in size than regular price changes. Slightly less than half of regular price changes are smaller than 5 percent in absolute value, and 28 percent of them is smaller than 3 percent. The distribution of the size of regular price changes is again quite symmetrical, with an estimated mean of the overlaying normal distribution equal to 0.260 and a standard deviation of 13.834.

[11]The pattern for the size of price changes is highly similar in the European dataset (results available upon request).

Figure 6: Size distribution of regular price changes (US)

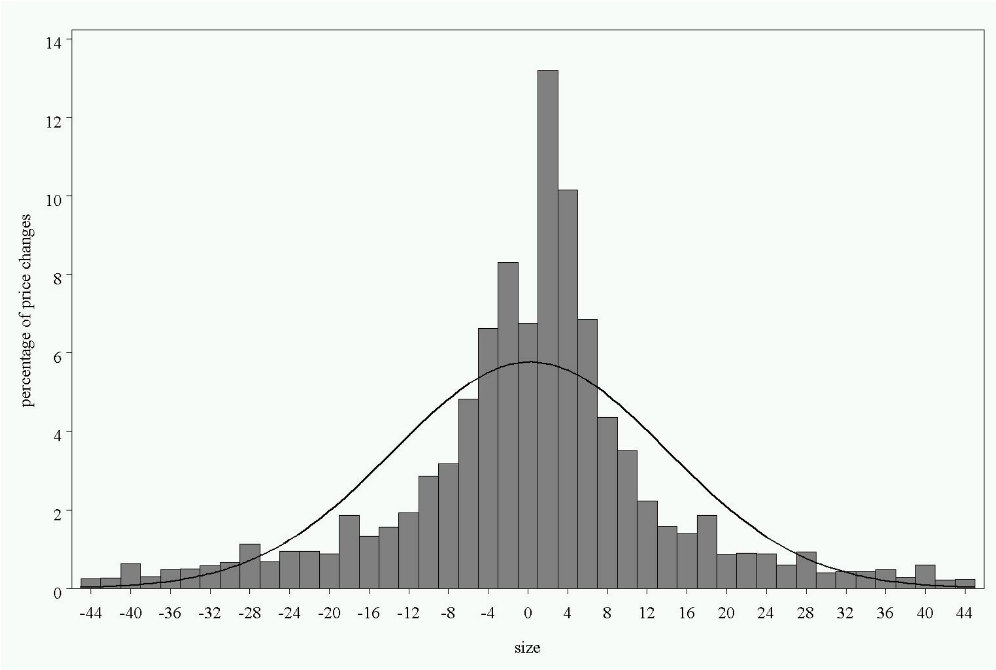

The higher incidence of small price changes in the regular price series compared to the posted price series is also apparent when we look at the kurtosis of both distributions. The size distribution of posted price changes has a kurtosis coefficient of 3.033, which is very close to the coefficient of 3 for the normal distribution. The size distribution of regular price changes, on the other hand, is clearly leptokurtic, with a kurtosis coefficient of 4.572. The latter distribution therefore displays more weight in the vicinity of zero and the tails are less fat compared to the normal distribution and the size distribution of posted price changes. The high prevalence of small price changes in the data is hard to reconcile with a standard menu cost model (Midrigan, 2010).

6 Conclusion

Price rigidity is a catalyst that transforms nominal shocks into real effects, and therefore plays a vital role in macroeconomic theory and policy. This paper adds to the recent progress in empirical research by estimating price rigidity from high-frequency scanner data of two large retailers, one in Europe and one in the US. The results on the extent of price stickiness that we find in micro data can be highly useful to calibrate business cycle models and support the policy making process.

Recent empirical results based on CPI data show that the frequency of price change is higher in the US than in Europe, but that the difference largely disappears when regular prices are analyzed. Our results on high-frequency scanner data show that the empirical literature fails to incorporate short-term pricing dynamics into the analysis. The scanner data at our disposal have an important advantage in that they are available in their base frequency, i.e. prices do not change in between data points. Consequently, there is no source of measurement error in the data gathering process. This enables us to incorporate all short-term movements into the measured price stickiness parameters.

A comparative analysis between our scanner data in base frequency on the one hand, and data at monthly frequency that we generated from our high-frequency price series on the other hand, shows us that the use of monthly data to study the impact of temporary sales on the flexibility of prices is potentially highly misleading, as the majority of sales episodes are very short-lived, a matter of weeks rather than months. Whereas regular prices in the European and US datasets indeed show more or less the same amount of flexibility when we study the derived monthly data, this finding collapses when the original scanner data in base frequency are examined. In the latter case, also regular prices appear to be far more flexible in the US data than in the European data, pointing to a more aggressive short-term pricing strategy in the US, a feature that is not captured by monthly data.

This result is robust to the type of sales filter that we apply to obtain our regular price series, filtering symmetric and/or asymmetric V-shaped sales with a fixed or flexible sales price and a duration of the sales episode of up to four weeks. The result also holds irrespective of the statistic used to capture price rigidity, be it the mean or the median across products or product groups. Monthly data might thus not be capable of capturing the true amount of price rigidity, making a comparison between regions severely biased in case of differences in short-term pricing tactics.

Additionally, the analysis of our scanner data confirms some stylized facts often cited in the literature: sales are an important driving force of micro price flexibility, at least in retailing; the frequency of price change is extremely heterogeneous across products; price changes are on average much larger than needed to keep up with aggregate inflation, although many small price changes do occur; price decreases are not merely a reflection of sales episodes, but a robust fact in micro price setting, providing no proof of downward rigidity in retail prices.

References

Baharad, E. and Eden, B. (2004). Price Rigidity and Price Dispersion: Evidence from Micro Data. *Review of Economic Dynamics*, **7** (3), 613–641.

Baudry, L., Le Bihan, H., Sevestre, P. and Tarrieu, S. (2004). Price Rigidity: Evidence from the French CPI Micro-Data. *ECB Working Paper*, **No. 384**, 74p.

Bils, M. and Klenow, P. (2004). Some Evidence on the Importance of Sticky Prices. *Journal of Political Economy*, **112** (5), 947–985.

Burstein, A. (2006). Inflation and Output Dynamics with State-Dependent Pricing Decisions. *Journal of Monetary Economics*, **53** (7), 1235–1257.

Calvo, G. (1983). Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics*, **12** (3), 383–398.

Campbell, J. and Eden, B. (2007). Rigid Prices: Evidence from US Scanner Data. *Federal Reserve Bank of Chicago - Working Paper Series*, **No. 2005-08**, 41p.

Caplin, A. and Spulber, D. (1987). Menu Costs and the Neutrality of Money. *Quarterly Journal of Economics*, **102** (4), 703–726.

Cecchetti, S. (1986). The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines. *Journal of Econometrics*, **33** (3), 255–274.

Chari, V., Kehoe, P. and McGrattan, E. (2000). Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem? *Econometrica*, **68** (5), 1151–1179.

Chevalier, J., Kashyap, A. and Rossi, P. (2003). Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data. *American Economic Review*, **93** (1), 15–37.

Christiano, L., Eichenbaum, M. and Evans, C. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, **113** (1), 1–45.

Coricelli, F. and Horvath, R. (2008). Price Setting and Market Structure: An Empirical Analysis of Micro Data. *William Davidson Institute Working Paper*, **No. 938**, 35p.

Dhyne, E., Alvarez, L., Le Bihan, H., Veronese, G., Dias, D., Hoffman, J., Jonker, N., Lünneman, P., Rumler, F. and Vilmunen, J. (2005). Price Setting in the Euro Area: Some Stylized Facts from Individual Consumer Price Data. *ECB Working Paper*, **No. 524**, 50p.

Dossche, M., Heylen, F. and Van den Poel, D. (2010). The Kinked Demand Curve and Price Rigidity: Evidence from Scanner Data. *Scandinavian Journal of Economics*, **112** (4), 723–752.

Dotsey, M., King, R. and Wolman, A. (1999). State-Dependent Pricing and The General Equilibrium Dynamics of Money and Output. *Quarterly Journal of Economics*, **114** (2), 655–690.

Dutta, S., Bergen, M., Levy, D. and Venable, R. (1999). Menu Costs, Posted Prices, and Multiproduct Retailers. *Journal of Money, Credit and Banking*, **31** (4), 683–703.

Eichenbaum, M., Jaimovich, N. and Rebelo, S. (2008). Reference Prices and Nominal Rigidities. *NBER Working Paper*, **No. 13829**, 53p.

Gagnon, E. (2009). Price Setting During High and Low Inflation: Evidence from Mexico. *Quarterly Journal of Economics*, **124** (3), 1221–1263.

Gali, J., Gertler, M. and Lopez-Salido, D. (2001). European Inflation Dynamics. *European Economic Review*, **45** (7), 1237–1270.

Gertler, M. and Leahy, J. (2008). A Phillips Curve with an Ss Foundation. *Journal of Political Economy*, **116** (3), 533–572.

Golosov, M. and Lucas, R. (2007). Menu Costs and Phillips Curves. *Journal of Political Economy*, **115** (2), 171–199.

Hosken, D. and Reiffen, D. (2004). Patterns of Retail Price Variation. *RAND Journal of Economics*, **35** (1), 128–146.

Karadi, P. and Reiff, A. (2007). Menu Costs and Inflation Asymmetries: Some Micro Data Evidence. *MPRA Paper*, **No. 7102**, 28p.

Kashyap, A. (1995). Sticky Prices: New Evidence from Retail Catalogs. *Quarterly Journal of Economics*, **110** (1), 245–274.

Kehoe, P. and Midrigan, V. (2010). Prices Are Sticky After All. *Federal Reserve Bank of Minneapolis - Research Department Staff Report*, **No. 413**, 62p.

Klenow, P. and Kryvtsov, O. (2008). State-Dependent or Time-Dependent Pricing: Does It Matter for Recent US Inflation? *Quarterly Journal of Economics*, **123** (3), 863–904.

— and Malin, B. (2010). Microeconomic Evidence on Price-Setting. *Prepared for the Handbook of Monetary Economics (Benjamin Friedman and Michael Woodford, editors).*

LACH, S. and TSIDDON, D. (1992). The Behavior of Prices and Inflation: An Empirical Analysis of Disaggregated Price Data. *Journal of Political Economy*, **100** (2), 349–389.

LEVY, D., BERGEN, M., DUTTA, S. and VENABLE, R. (1997). The Magnitude of Menu Costs: Direct Evidence from a Large US Supermarket Chain. *Quarterly Journal of Economics*, **112** (3), 791–825.

MANKIW, G. and REIS, R. (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, **117** (4), 1295–1328.

MASAHIRO, H. and YUMI, S. (2007). Price Setting in Japan: Evidence from CPI Micro Data. *Bank of Japan Working Paper Series*, **No. 07-E-20**, 61p.

MIDRIGAN, V. (2010). Menu Costs, Multi-Product Firms and Aggregate Fluctuations. *Econometrica, forthcoming.*

NAKAMURA, E. and STEINSSON, J. (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models. *Quarterly Journal of Economics*, **123** (4), 1415–1464.

PEERSMAN, G. (2004). The Transmission of Monetary Policy in the Euro Area: Are the Effects Different Across Countries? *Oxford Bulletin of Economics and Statistics*, **66** (3), 285–308.

PELTZMAN, S. (2000). Prices Rise Faster than They Fall. *Journal of Political Economy*, **108** (3), 466–502.

RAY, S., CHEN, H., BERGEN, M. and LEVY, D. (2006). Asymmetric Wholesale Pricing: Theory and Evidence. *Marketing Science*, **25** (2), 131–154.

ROTEMBERG, J. (2005). Customer Anger at Price Increases, Changes in the Frequency of Price Adjustment and Monetary Policy. *Journal of Monetary Economics*, **52** (4), 829–852.

TAYLOR, J. (1980). Aggregate Dynamics and Staggered Contracts. *Journal of Political Economy*, **88** (1), 1–23.

VERONESE, G., FABIANI, S., GATTULLI, A. and SABBATINI, R. (2005). Consumer Price Behaviour in Italy: Evidence from Micro CPI Data. *ECB Working Paper*, **No. 449**, 46p.

WULFSBERG, F. (2009). Price Adjustments and Inflation - Evidence from Norwegian Consumer Price Data 1975-2004. *Norges Bank Working Paper*, **No. 2009/11**, 34p.

Appendices

I. Description of the datasets

EU		US	
Product Group	# Products	Product Group	# Products
Bath soap	182	Bath soap	600
Bathroom tissues	30	Bathroom tissues	128
Canned soup	5	Bottled juices	511
Canned tuna	46	Canned soup	445
Cereals	49	Canned tuna	278
Coke	39	Cereals	490
Dish detergent	43	Cheeses	667
Dishwasher tablets	43	Crackers	330
Emmental	56	Dish detergent	287
Fruit juice	54	Grooming products	1599
Grooming products	243	Laundry detergent	582
Gruyere	19	Refridgerated juices	228
Laundry detergent	94	Snack crackers	425
Lemonade	33	Soaps	337
Smoked salmon	18	Soft drinks	1746
Snack crackers	136	Toothpastes	608
Soaps	34		
Toothpastes	146		
Total	1270	Total	9261

II. Definition of the different statistical measures

In order to calculate our statistics with respect to the stickiness of prices in Europe and the US, we create four binary variables based on the price trajectories in our scanner datasets, in a similar way as in Dhyne *et al.* (2005). Assume that p_{ijt} is the price of a product i, sold at store j, at time t. Because we only include adjacent price observations in our analysis, we first create an indicator variable I_{ijt}^{obs}, attached to p_{ijt}, that takes a value of 1 if a price quote for product i at store j from the previous period is observed and 0 otherwise:

$$I_{ijt}^{obs} = \begin{cases} 1 & \text{if } p_{ij,t-1} \text{ is observed} \\ 0 & \text{if } p_{ij,t-1} \text{ is not observed} \end{cases}$$

Studying only adjacent price observations, we will only withhold price quotes for which $I_{ijt}^{obs} = 1$ to calculate the price stickiness parameters. The second binary variable that we create is a price change indicator for product i at store j at time t:

$$I_{ijt} = \begin{cases} 1 & \text{if } p_{ijt} \neq p_{ij,t-1} \\ 0 & \text{if } p_{ijt} = p_{ij,t-1} \end{cases}$$

In the same way, we also create a price increase and price decrease indicator for product i at store j at time t:

$$I_{ijt}^{+} = \begin{cases} 1 & \text{if } p_{ijt} \succ p_{ij,t-1} \\ 0 & \text{if } p_{ijt} \leq p_{ij,t-1} \end{cases}$$

$$I_{ijt}^{-} = \begin{cases} 1 & \text{if } p_{ijt} \prec p_{ij,t-1} \\ 0 & \text{if } p_{ijt} \geq p_{ij,t-1} \end{cases}$$

Adding these four indicator variables to our datasets, calculation of the frequency of price change for product i across all n stores and τ time periods is straightforward:

$$\lambda_i = \frac{\sum_{j=1}^{n} \sum_{t=2}^{\tau} I_{ijt}}{\sum_{j=1}^{n} \sum_{t=2}^{\tau} I_{ijt}^{obs}}$$

The frequency of price increase and decrease for product i across all n stores and τ time periods are again calculated in a similar way:

$$\lambda_i^+ = \frac{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}^+}{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}^{obs}}$$

$$\lambda_i^- = \frac{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}^-}{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}^{obs}}$$

The size of price changes, price increases and price decreases are then easily calculated as follows:

$$dp_i = \frac{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}\ (ln\ p_{ijt} - ln\ p_{ij,t-1})}{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}}$$

$$dp_i^+ = \frac{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}^+\ (ln\ p_{ijt} - ln\ p_{ij,t-1})}{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}^+}$$

$$dp_i^- = \frac{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}^-\ (ln\ p_{ij,t-1} - ln\ p_{ijt})}{\sum_{j=1}^{n}\sum_{t=2}^{\tau} I_{ijt}^-}$$