BANK OF ENGLAND

Working Paper No. 378

Do supermarket prices change from week to week?

Colin Ellis

November 2009

BANK OF ENGLAND

Working Paper No. 378
Do supermarket prices change from week to week?

Colin Ellis[(1)]

Abstract

This paper examines the behaviour of supermarket prices in the United Kingdom, using weekly scanner data supplied by Nielsen. A number of stylised facts about pricing behaviour are uncovered. First, prices change very frequently in supermarkets, with 40% of prices changing each week, and even controlling for 'temporary' changes, a quarter of prices change each week. Importantly, there is evidence that focusing on monthly observations, rather than weekly ones, overstates the implied stickiness of prices. Second, the probability of price changes is not constant over time — all product categories have declining hazard functions. Third, the range of price changes is very wide, with some very large price cuts and price rises; but despite this, a significant number of price changes are very small. Fourth, there appears to be little link between the frequency and magnitude of price changes — prices that change less frequently do not tend to change by more. Fifth, the strongest correlation between price and volume changes is contemporaneous, suggesting that prices and volumes move together from week to week. And sixth, rough analysis based on simplifying assumptions suggests that consumers are fairly price sensitive: volumes change by more than prices.

Key words: Supermarket prices, behaviour of prices, demand elasticities.

JEL classification: D40, E31.

(1) Formerly Bank of England, now at Daiwa Securities SMBC Europe Ltd. Email: colin.ellis@daiwasmbc.co.uk

The views expressed in this paper are those of the author, and not necessarily those of the Bank of England, the Monetary Policy Committee, the Nielsen Company or Daiwa Securities SMBC Europe Ltd. I am very grateful to Simon Boswell, Alex Bowen, Neal Hatch, Miles Parker, Mike Watkins and especially Philip Bunn and Stephen Hilton, for assistance and useful comments. I hereby claim all mistakes as my own. This paper was finalised on 25 September 2009.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/publications/workingpapers/index.htm

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email mapublications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Contents

Summary

The object of UK monetary policy is to target inflation, as measured by the consumer prices index, the CPI, at 2% a year. In order for policymakers to keep inflation on target, they need to understand how the actual prices in the economy that underlie official inflation measures behave. One central issue is the degree of nominal rigidity in the economy, the extent to which prices and wages are ‘sticky’. That follows if companies are either unable or unwilling to adjust either quickly, perhaps because of costs of adjustment. This stickiness has profound implications for inflation dynamics and therefore for the conduct of monetary policy.

As a result, a key question for policymakers is how often prices change, and by how much. Early work to investigate this phenomenon often focused on examining the behaviour of aggregate inflation rates at the macroeconomic level. But that can potentially be misleading. So recently economists have spent time examining so-called ‘micro-pricing’ data – the prices of individual products, which may be weighted and aggregated to construct the official price indices.

This paper adds to that exploratory effort, and examines how prices behave for around 280 products in 240 different supermarkets across Great Britain. The data cover a recent three-year period, and were kindly made available to the Bank of England by Nielsen, a market research company. In all, the data set accounts for a little under 5% of annual household expenditure. One big advantage of these data is that they are available at a relatively high frequency – Nielsen collect information on a weekly basis, as oppose to the monthly collection of price quotes often used by national statistical offices. By examining prices and volumes over shorter periods, in particular a week rather than a month, we can shed some light on whether evidence from monthly data may overstate the true degree of price stickiness in the economy – as, by construction, a monthly price series can only change a maximum of twelve times a year.

Several interesting features emerge from analysing the data. Prices change quite frequently in supermarkets – as much as 40% a week, even after trying to strip out temporary promotions and sales – and there is also evidence that monthly price observations can overstate the implied stickiness of prices. The range of different prices changes is very wide, with some very large moves but also many small ones, and there appears to be little link between how much a price changes by and how long it has been since the last time it changed. Prices and volumes – the number of goods sold – tend to move together in the data, and there is tentative evidence that consumers may be quite price sensitive, with volumes changing more than one-for-one when prices change. But, importantly, it must be borne in mind that all of these results relate to supermarket prices, rather than other prices, which may exhibit less flexibility.

1 Introduction

UK monetary policy aims to keep inflation on target at 2% a year. So it is important for policymakers to consider how prices behave. In particular, the degree of nominal rigidity in the economy will influence the short-term impact of nominal interest rates on real activity and the response to inflation to monetary policy.

The notion of nominal rigidity is a feature of many economic models. Essentially, these models assume that companies are unable to freely adjust their prices. A variety of mechanisms have been put forward to explain this assumption. These include costs of adjusting prices (Rotemberg (1982) and Mankiw (1985)), staggered price contracts (Taylor (1980)), threshold or so-called 's,S' pricing (Sheshinski and Weiss (1977)), and fixed probabilities of being able to change prices (Calvo (1983)).

One popular pricing model that results from the last approach is the so-called New Keynesian Philips Curve (NKPC). This relates current inflation to future expected inflation and the deviation of marginal cost from its steady-state value. One feature of these models is that, when estimated, they imply price durations – how long, on average, it takes for companies to change their prices.[1] Early estimates of the NKPC implied that firms changed their prices every 15 to 18 months (Gali and Gertler (1999)), although some estimates have suggested that prices change once every two years (Smets and Wouters (2003)).[2]

These timings are somewhat longer than evidence from direct surveys of companies' price setting behaviour – Blinder *et al* (1998) and Druant *et al* (2005) both find that the median price changes once a year in the United States and the euro area, respectively. Other evidence suggests that individual prices may be more flexible than this. In particular, Amirault *et al* (2005) and Bils and Klenow (2004) found that prices change on average every three to four months. And evidence from 300 of the Bank of England's Agency contacts suggests that half of companies change prices at least five times a year (Bank of England (2006)). One important point of note when comparing results from these various different studies is that sectoral coverage can vary significantly, which may have an impact on the resulting estimates of price flexibility. And the role and treatment of sales, or temporary promotions, appears to be important in measuring price duration – excluding sales and promotions, Nakamura and Steinsson (2008) find the median duration of retail prices is between eight and eleven months.

Kehoe and Midrigan (2007) also examine the role of temporary promotions. Excluding their definition of sales, they find that the implied duration of prices is around four to five months, compared to just three weeks if those sales are included. Somewhat uniquely, Kehoe and Midrigan use weekly store-level scanner data to investigate price frequencies, rather than the more common approach of examining monthly micro-price data, for example in Bunn and Ellis (2009) or Baudry *et al* (2007). As Kehoe and Midrigan point out, while scanner-level data may

[1] McAdam and Willman (2007) find that the Calvo probability parameter defines the upper limit of price duration rather than the average, when NKPC is adapted to include a state-dependent price-resetting signal.
[2] These papers essentially estimate macro models from which micro-pricing behaviour is inferred. Boivin *et al* (2007) and Mumtaz *et al* (2009) adopt a different approach, modelling aggregate and disaggregated pricing data simultaneously.

be less comprehensive than official micro-price estimates compiled by national statistical offices, its big advantage is the higher frequency of observations. Monthly data give no indication whatsoever about what happens to prices within each month, regardless of whether they are recorded as a monthly average or a point estimate.

This paper follows a similar approach to Kehoe and Midrigan (hereafter KM), and examines weekly store-level data for a sample of UK supermarket products. One difference between this analysis and that of KM and Chevalier *et al* (2000) is that, whereas their data samples are concentrated around one urban sample (Chicago) and are for one retail chain, our data covers the whole of Great Britain and several different retail chains. One key point of note is that, by construction, this paper (and indeed KM) is focusing on prices in outlets where prices may tend to change more frequently than across the wider economy as a whole. Despite recent advances into less traditional product markets, supermarkets still sell more food than anything else – so food will account for a larger share of sales in these data than in, for example, the UK CPI. That must be borne in mind when comparing these results (and KM's) to other work that uses broader but less frequent price data. Indeed, Greenslade and Parker (2008) suggest this may well be the case in the United Kingdom, finding that retailers change their prices much more frequently than firms in other sectors of the economy.

Like KM and Chevalier *et al*, this analysis finds that 'raw' supermarket prices change very frequently, but that some of those changes can be accounted for by temporary discounts. However, even after adjusting for sales prices change very frequently, at least once a month on average. The overall picture is of a considerable degree of price flexibility in the supermarket sector. The next section describes the data used in this analysis, and the following two sections describe results. Finally, the paper concludes.

2 Data

Nielsen is a market research company that provides clients with analysis of sales trends and promotional impacts. To provide this service they collect data from a nationwide network of Electronic Point of Sale (EPoS) checkout scanners which represent sales at 65,000 supermarket and convenience stores in Great Britain. Nielsen maintains a detailed database of different products, covering selling prices, volumes sold, and promotional activities.

This paper uses a bespoke data set created from the Nielsen database. It covers around 240 different supermarkets located throughout Great Britain, covering the largest retailers.[3] In total, just over 280 distinct products are included in the data set; however, not all stores stock all products, and some products appear intermittently. The individual products were chosen, with advice from Nielsen staff, both with consideration to brand importance (see Nielsen (2007)), data availability, and to try to get a broad range of different types of goods.

The data set covers selling price and the quantities sold over a three-year period: the data start in the week of 19 February 2005 and end in the week of 9 February 2008. It is worth bearing in

[3] Specifically, Tesco, Asda, Sainsbury's, Morrison's, Somerfield and Waitrose.

mind that any results from this analysis are conditional on this sample – in particular in terms of the shocks that hit the UK economy and how they played out over this period. In all, there are just under 5½ million individual price observations, or roughly 35,000 different price observations each week. The price observations are 'average' prices for each week: this means that temporary changes in prices, such as selling damaged goods more cheaply, will appear in the data. To the extent that these represent genuine price changes, these observations are useful – they are direct evidence on how easy it is for firms to change prices. But to the extent that these changes represent changes in quality, it may be more preferable to exclude these if we want to focus on underlying prices, for example for identical products.[4] I took a deliberate decision not to censor or restrict data any further, partly on the basis of these considerations – the aim was to get the 'cleanest' version of the raw data without truncating any of the price distribution. Even so, the duration of price trajectories – how long an item is in the data set – varies from product to product and store to store, reflecting both the availability and seasonal demand for various items. But across the data set as a whole, on average over 90% of the sample (by sales values) is observed each week.

Nielsen break the products down into ten different categories: Alcohol; Bakery; Confectionary; Dairy; Fresh (eg fruit and vegetables); Frozen; Grocery; Household; Personal (eg health care); and Soft Drinks. In all, the data set accounts for a little under 5% of annual household expenditure.

Sales values for each category, as a proportion of total sales in the data set, are shown in Table 1: the Fresh category clearly dominates. In the results that follow, this high weight must be borne in mind.

Table 1: Share of sales by product category

Category	Frequency		Sales	
	Number	*Percentage of total*	*£ million*	*Percentage of total*
Alcohol	319,195	5.6	6,449	5.9
Bakery	161,087	2.8	1,624	1.5
Confectionary	544,268	9.6	2,318	2.0
Dairy	614,746	10.8	12,778	11.7
Fresh	1,030,831	18.1	61,890	56.7
Frozen	255,294	4.5	1,986	1.8
Grocery	1,234,536	21.7	10,323	9.5
Household	408,352	7.2	3,430	3.1
Personal	492,449	8.7	2,460	2.3
Soft Drinks	621,778	10.9	6,033	5.5
Total	5,682,536	100	109,110	100

[4] Averaging could have other implications: for example, multi-buys will reduce average selling prices (which are excluded from ONS data), and a price cut may appear in two separate observations if it happens mid-week. Averaging could also exaggerate the role of temporary promotions, depending on how exactly they are implemented.

2.1 *Accounting for temporary changes in prices*

As previous work has noted, temporary sales are likely to play a role in any analysis of price changes. While these can be a genuine indication of price flexibility – if firms can change their prices easily and swiftly, that suggests costs of changing prices are low – it is also possible that 'noise' in the data may reflect measurement issues (eg multi-buys or quality changes). In addition, if we are interested in focusing on underlying price changes – those longer-frequency changes that may be more likely to reflect macroeconomic conditions – there is a case for smoothing through some of the volatility that temporary price changes will cause. In this paper, headline results from the data – those including all price changes – will be presented alongside results that adjust for temporary changes in prices.

Those temporary changes will be accounted for using three different methods. One method is to examine the so-called 'reference price' put forward by Eichenbaum *et al* (2008). This notional 'reference price' is simply the modal price within a given quarter. If temporary discounts are important, using reference prices will clearly wash a significant degree of variability out of the raw price data – in particular, reference prices will not change at all within the three-month window.[5] If anything, this could serve to understate the degree of price flexibility in the economy.

However, the notion of a reference price is a very powerful tool for examining an important question in this paper. One of the key reasons for examining scanner data from supermarkets is that it is available on a significantly higher frequency than other micro data sources. The official micro-price data underpinning the producer prices index (PPI) or CPI will only have one set of observations for each reporting period – each price is only recorded once a month. By construction, this will miss any intra-month variation in prices that may, at the same time, be picked up in our weekly scanner data. Comparing these scanner data with both CPI and PPI micro data is not straightforward, as the former is based on single price observations on a given day of the month, while the latter is based on the notion of an average monthly price. Given that the underlying scanner data are essentially average weekly prices, that suggests they sit somewhere in between the two, although distinguishing precisely where would require a number of (untestable) assumptions.

But what the scanner data can provide is an indication of how much is lost by moving from weekly observations to a single monthly estimate for prices. By constructing monthly reference prices based on the supermarket returns, we can examine what happens to the implied frequency of price duration, compared to the weekly results. The null hypothesis is clearly that it should increase – by construction, monthly data can change less frequently than weekly data. But how much? Scanner data reference prices will shed light on how much we may miss by focusing on the official monthly price data.

In addition to constructing reference prices, this paper will use two other metrics to wash out some of the short-term variation in the data, which could reflect temporary sales. The first is KM's notion of a 'regular price'. This classifies price reductions as short-term sales or discount

[5] The choice of a three-month window seems arbitrary.

offers if they are reversed sufficiently quickly, within some defined period.[6] For this paper, regular prices are calculated using a window of five weeks, following KM's analysis. So if a price cut is reversed within five weeks, it is excluded from the data. The resulting 'regular price' series is thereby generated from the observed price series, smoothing through these short-term price changes.

One point of note is that this definition of 'regular price' has a similar defining characteristic as a 'reference price' – namely the window of observation and calculation. Both concepts are 'time dependent' in the sense that they are determined by a (subjective) length of time which is used to calculate the adjusted price series. In real life, the pattern of discounts could be heterogeneous between and within product categories, which could make this blanket 'common-window' approach inappropriate.

In light of this, the final method for treating temporary price deviations is free from this consideration of 'window length'. In this paper, I define a 'price reversal' as occurring when prices move either up or down, before exactly reversing at some later (unbounded) point in time.[7] These price reversals can then be excluded from the data, so that only the remaining observations are treated as price changes.

The unbounded nature of this 'price reversal' concept could potentially lead to long periods of price reversal; but if KM's finding of an average sales duration of two weeks holds in the Nielsen data, then the frequency of price changes should be similar either using 'regular prices' or excluding 'price reversals'. However, the notion of a 'price reversal' also overlaps with the concept of a 'reference price': if most deviations from some 'normal' prices are temporary discounts that are reversed, adjusting for these reversals should yield a clear picture of what that 'normal' price is. If excluding price reversals does drive the frequency of price change sufficiently higher, that could provide some justification for using quarterly reference prices.

3 Analytical results

This section describes the analytical results from examining prices in the data set. Unless otherwise stated, the results presented are weighted by sales values for individual items. The results are grouped into five broad categories, covering: frequency; hazard functions; magnitude; frequency and magnitude; and links between prices and volumes.

3.1 Headline price change frequencies

On an unweighted basis, roughly 40% of prices change each week in the data set. On a weighted basis, this is pushed up to 60% (Table 2). One implication of this result is that those items that consumers spend more on tend to change price more frequently. One category, 'Fresh', exhibits more price changes than other categories by some margin, and accounts for a significant part of the higher weighted frequency of price changes. Excluding 'Fresh' products, the weighted frequency of price changes is 40% a week.

[6] The algorithm is described in detail in the annex.
[7] This is calculated by considering price changes in pence, as percentages will vary for rises/falls of the same absolute magnitude.

Table 2: Frequency of price changes

Prices changing each week	*Per cent*
Unweighted	40.9
Weighted	60.0
- excl. Fresh	40.4

There is some variation in the frequency of price changes across categories, with a maximum of 75% a week for Fresh and a minimum frequency of 29% a week for Dairy (Table 3). The frequency of price changes in the data set is split roughly half and half between rises and falls, with no evidence of downward nominal rigidity.

Table 3: Frequency of price changes by product category

Fraction of prices changing each week

Category	*Per cent changing*	*Per cent rising*	*Per cent falling*
Alcohol	58.0	29.0	29.0
Bakery	48.5	24.8	23.7
Confectionary	32.2	16.5	15.7
Dairy	28.6	15.7	12.9
Fresh	75.0	37.4	37.6
Frozen	32.4	16.2	16.1
Grocery	38.8	20.0	18.8
Household	35.7	17.8	17.9
Personal	40.9	20.5	20.4
Soft Drinks	55.1	27.9	27.2

This degree of price flexibility is significantly higher than found in many studies of monthly data, suggesting that examining prices at even a relatively high (monthly) frequency can yield a misleading picture: there is significant variation in prices within the month that monthly data simply do not capture.

However, a significant proportion of these price changes may be attributed to temporary discounts, under either the 'regular price' or 'price reversal' methods. Focusing on the regular prices method reduces the frequency of price changes to 37% a week, whereas the price reversals method reduces the frequency to 45%. If we then exclude 'Fresh' products as well, we find that around a quarter of prices change each week in the data set using both methods (Table 4). The consistency of results from these two different approaches is encouraging and suggests that they may well manage to account for temporary discounts: if the results had been very different, that would have raised concerns that we were not adjusting the data appropriately.

Table 4: Frequency of price changes, adjusting for temporary changes

Prices changing each week	*Per cent*
Total	60.0
- excl. price reversals	*45.3*
- based on 'regular' prices	*37.3*
Excluding Fresh products	40.4
- excl. price reversals	*27.0*
- based on 'regular' prices	*24.3*

The majority of price reversals are decreases (followed by increases), which is consistent with most temporary price changes being short-term promotions or sales. One important result from the data is the finding that, on average, the duration of a temporary discount or sale – as defined either using regular prices or by excluding price reversals – is between two and three weeks. That is very close to KM's findings, and, together with the fact that temporary price changes appear to account for between a quarter and two fifths of all price changes, provides a good guide to the impact of temporary discounts in the data set.

These results suggest that, as in other studies, sales can account for a significant proportion of relatively high-frequency price changes. But unlike other studies, such as KM and Nakamura and Steinsson (2008), excluding sales does not have as large an impact on the implied duration of price changes. Excluding price reversals, 45% of prices still change each week in the data sample (27% excluding Fresh), an implied price duration of a little over two weeks (just under four weeks, excluding Fresh). Looking at regular prices, 24% of these change each week when Fresh is excluded, implying a similar duration of around a month. This is markedly shorter than other estimates in the literature. As mentioned earlier, that is partly likely to reflect the fact our data come from supermarkets, rather than other types of retail output.

3.2 How misleading are longer-frequency estimates?

One key benefit of the Nielsen data, as discussed earlier, is its weekly frequency. Many other pricing studies rely on monthly frequency data at best – which, by construction, will limit the implied frequency of price changes. The Nielsen data can offer some insights here – in particular, how big any distortion may be from focusing on monthly data. This can be examined by using Eichenbaum *et al*'s (2008) notion of reference prices – the modal price within a defined period. Two different reference windows were used: a quarterly one, and a monthly one.

In constructing reference prices, one interesting observation was that several items in the data set exhibited multiple modes within three-month periods. In dealing with these, the reference price picked was the highest (maximum) mode within each quarter or month, on the basis that most temporary promotions and discounts were likely to result in lower prices. Table 5 presents results from the reference price series: even excluding the 'Fresh' category, 50% of reference prices change each quarter, implying an average duration of six months.

Table 5: Reference prices

Reference prices changing each period	*Per cent*
Quarterly reference prices	
Total	68.7
Excluding Fresh	50.3
Monthly reference prices	
Total	64.0
Excluding Fresh	44.0

For comparison, 64% of monthly reference prices changed each month, or 44% excluding Fresh products. That implies an average price duration of just over two months, compared with the duration of around half a month implied by the weekly data (again excluding Fresh products). So by moving from a weekly to a monthly frequency – but using the same underlying data – the implied frequency of price adjustment has quadrupled. Indeed, the implied duration from monthly reference prices is twice that found using the regular price of price reversal adjustments, even excluding Fresh products. This strongly suggests that existing duration estimates that are based on monthly frequency data run the risk of overstating the degree of nominal rigidity in the economy. However, our monthly reference prices are still more flexible than new evidence from UK CPI data, which suggest that goods prices change around once every four months (see Bunn and Ellis (2009)). That discrepancy is likely to reflect the source and nature of our data (supermarkets with frequent price changes).

These results clearly demonstrate how the frequency of the underlying data matters – by using long-frequency data, the implied price duration can be considerably higher, and by focusing on such data, be it monthly or quarterly, we may miss much of the higher-frequency variation that is actually present in prices.

3.3 Seasonal factors and price durations

Given these concerns about the impact of looking at monthly averages, the weekly data are likely to offer the best guide to seasonal patterns and price duration. Indeed, the frequency of weekly price changes is broadly constant by calendar month (Table 6), suggesting that seasonal factors do not play a significant role. This is in contrast to Nakamura and Steinsson (2008), who found the frequency of price changes to be highly seasonal in the United States.[8]

[8] The seasonality of the size (as opposed to the frequency) of price changes is discussed later.

Table 6: Frequency of price changes by calendar month

Fraction of prices changing each week

Category	*Per cent changing*	*Per cent rising*	*Per cent falling*
January	61.5	29.9	31.6
February	58.5	29.4	29.1
March	59.1	31.6	27.5
April	58.1	28.5	29.5
May	60.0	32.4	27.7
June	61.5	29.1	32.4
July	60.6	28.5	32.1
August	59.4	29.4	30.0
September	58.5	29.2	29.2
October	60.8	30.9	29.9
November	61.6	32.2	29.4
December	60.8	31.4	29.4

How long do prices tend to persist at a given level? It turns out that the distribution of price durations is highly skewed (Chart 1). In the data set, roughly 80% of prices have changed in the previous week, but the smaller tail of the distribution is very long. Although the average duration of prices is 1.6 weeks (including Fresh), the median duration is just 1 week.

Chart 1: Distribution of price duration

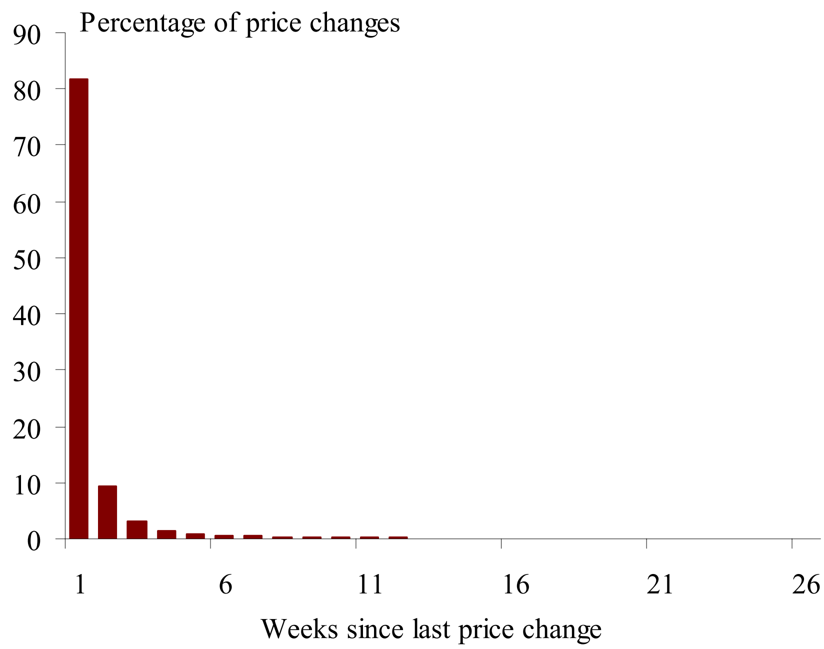

The frequency of price changes also varies somewhat over time within the data sample: this is not surprising, as any change in headline inflation must be accounted for either by more frequent or larger changes in prices (or a shift in the weights towards items with higher inflation rates). Of course, the scanner data do not cover the full range of prices in the CPI or RPI – and while there is evidence of some (lagging) relationship between the two, the strongest correlation between the difference in price change frequency from quarter to quarter is with the general public's perceptions of inflation, as gauged by the median response in the Bank/GfK NOP survey (Chart 2).[9] This could suggest that the public's inflation perceptions are influenced by those prices that they observe most frequently (see Driver and Windram (2007)).

Chart 2: Average frequencies and inflation perceptions

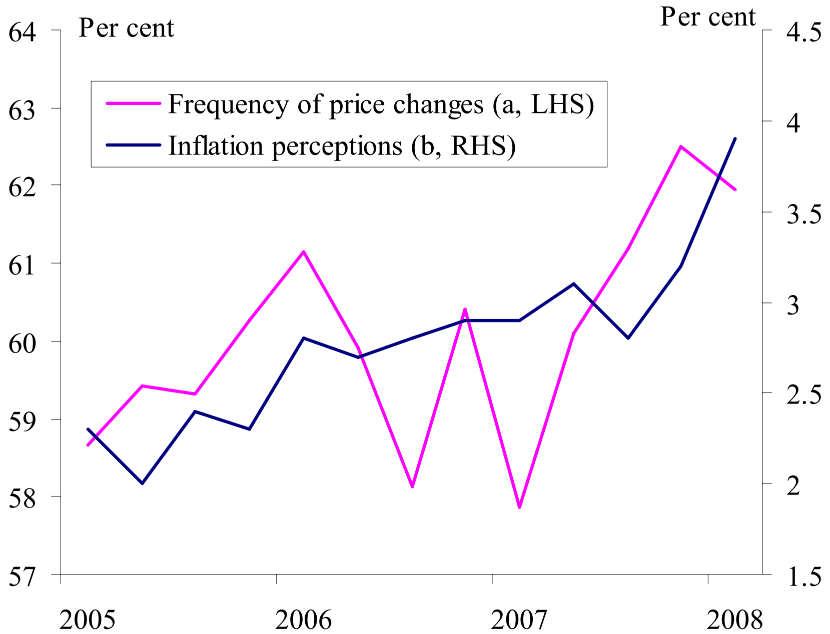

(a) Average percentage of prices changing each week.
(b) Survey median.

However, it is important to remember that looking at all of the price changes in this manner will double-count items that have multiple changes in the data set. In order to address this, more formal hazard analysis is considered in the next section.

3.4 Hazard functions

Based on the high-frequency Nielsen data, supermarket prices appear to be very flexible indeed. But in part this could reflect products with frequently changing prices appearing many times. In order to investigate this, more formal hazard functions were calculated using the price data. Hazard functions estimate the probability of a price changing at some point in time, given when the previous change in price occurred.

[9] For more information on the Bank/GfK NOP survey, see Benford and Driver (2008).

Chart 3: Hazard function for all data

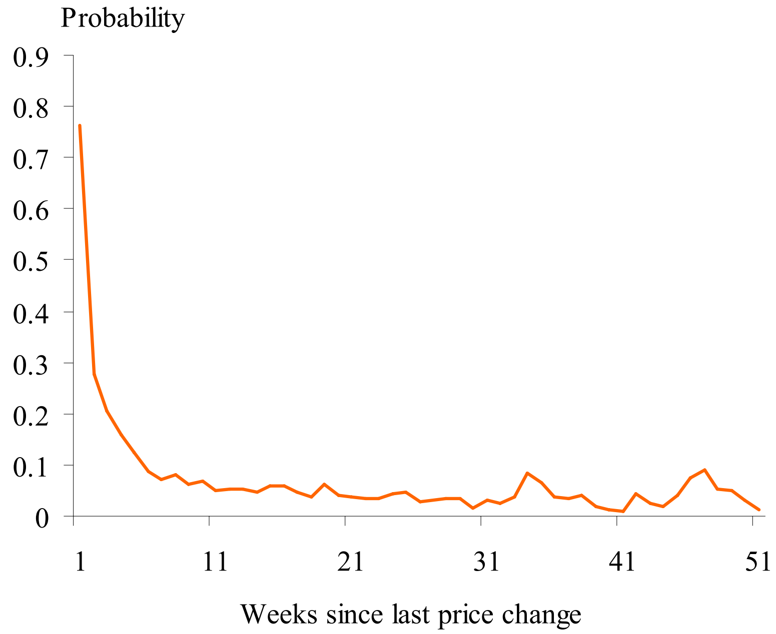

The resulting hazard function for the entire data set is shown in Chart 3. The function is sharply downward sloping, as KM found in their analysis. This argues very strongly against any uniform time-dependency framework for price-setting: under these frameworks, the frequency of price adjustment is invariant with regard to price duration, and this is not evident either in these data or in several of the other studies previously mentioned.[10]

One concern here may be that raised by Fougère *et al* (2005), who find that aggregate hazard functions can be misleading, and that estimating functions for disaggregated groups of products can lead to very different inference. In the case of the Nielsen data, in actual fact hazard functions for the different product categories are broadly similar (Charts 4a and 4b), with no category exhibiting a constant probability of price adjustment. That could reflect the fact our data come from supermarkets rather than other outlets – Fougère *et al* also find that hazard functions differ across outlet types, noting in particular that supermarkets tend to exhibit decreasing hazard functions. Yet perhaps that is not surprising given that they also estimate that 'flexible' prices account for 80% of all supermarket prices. Indeed, these downward-sloping disaggregated hazard functions are consistent with Fougère *et al*'s finding of marked flexibility in supermarket prices.

[10] Alvarez *et al* (2005) get around this problem by positing groups of firms with different frequencies of price adjustment in order to match the data.

Chart 4a: Hazard functions for product categories

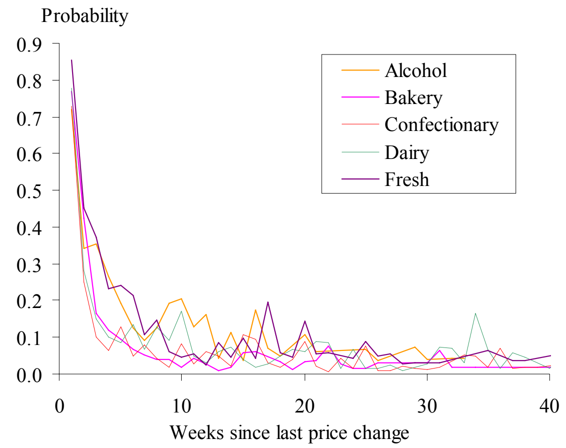

Chart 4b: Hazard functions for product categories

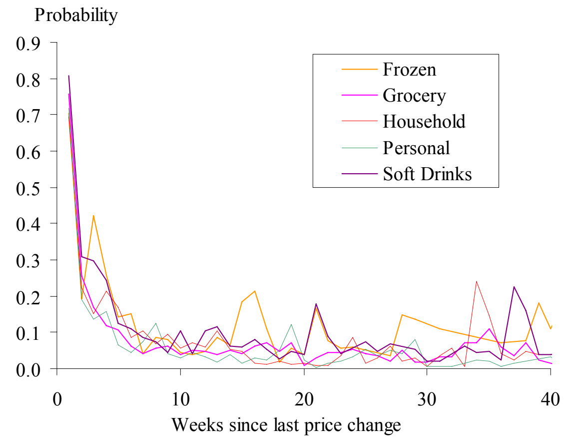

3.5 *Magnitude of price changes*

The results so far suggest that prices change quite frequently. But all of the analysis has so far been restricted to analysis of frequency – the results have just observed prices changing without considering how much they have changed *by*. This section examines the magnitude of price changes in the data set.

Across the data set as a whole, the size of changes varied markedly, from around -33% to +45% (Table 7). However, these data represent the tails of the distribution – most price changes were much smaller in size, with the interquartile range being just 6.7 percentage points (pp). Chart 5 plots the distribution of price changes for all observations. The high proportion of small price changes suggests that fixed menu costs are not widespread, while the observed large price changes are contrary to what might be expected if firms faced quadratic costs of adjusting prices.

Table 7: Magnitude of price changes

	Per cent
Mean	0.9
Median	0.2
1st percentile	-32.7
5th percentile	-15.2
10th percentile	-9.1
25th percentile	-3.2
75th percentile	3.5
90th percentile	10.7
95th percentile	18.2
99th percentile	45.2

Chart 5: Distribution of all price changes

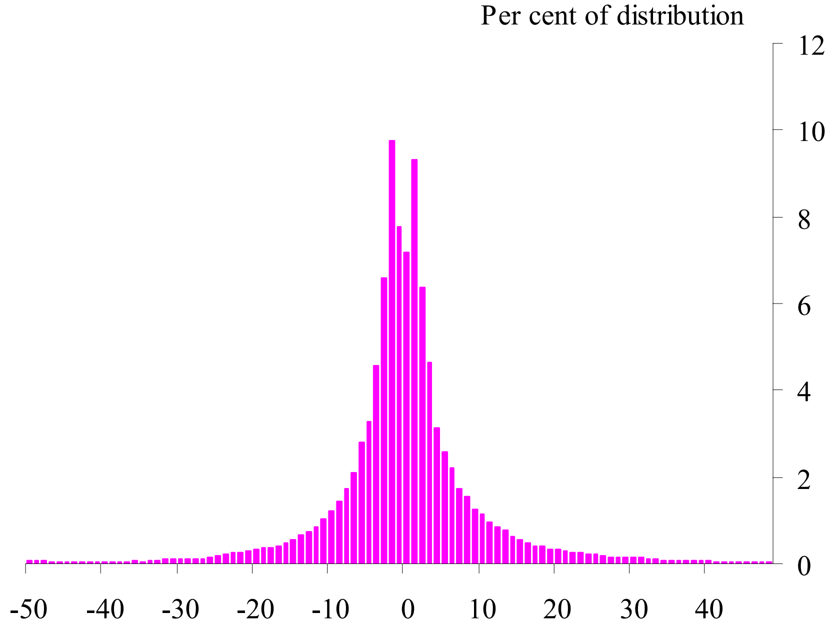

Within different product categories there is some degree of variation in the distribution of price changes. Table 8 shows percentiles of the price change distribution for different categories – Soft Drinks appear to have the widest distribution of price changes, followed by Frozen. Alcohol and Bakery have the slimmest distributions.

Table 8: Percentiles of price change distribution by category

Percentage change in prices

Category	*Mean*	*Median*	*Percentiles* *5th*	*25th*	*75th*	*95th*
Alcohol	0.5	0.0	-12.5	-1.6	2.0	14.1
Bakery	0.8	0.9	-12.2	-2.2	3.2	14.5
Confectionary	1.2	0.5	-19.0	-3.3	4.3	21.1
Dairy	1.3	0.7	-13.0	-1.6	2.8	16.3
Fresh	0.7	-0.2	-14.1	-3.4	3.6	16.7
Frozen	1.9	0.4	-29.0	-2.8	5.5	33.3
Grocery	1.2	0.4	-20.0	-2.3	3.4	22.2
Household	1.6	-0.2	-18.9	-1.3	2.2	19.0
Personal	1.3	0.1	-18.7	-3.4	3.9	21.4
Soft Drinks	2.0	0.4	-29.7	-3.6	4.7	36.8

Interestingly, the distribution of price changes varies relatively little by calendar month – as with the frequency results, this suggests relatively little role for seasonal effects (Table 9). These results suggest that supermarkets are not the typical venues for large seasonal sales, which may be more apparent in other CPI categories such as furniture. One observation is that the average price change increases over the sample, but the median is more stable (Chart 6). This suggests that the distribution of price changes became more skewed over time in the Nielsen data.

Table 9: Percentiles of price change distribution by calendar month

Percentage change in prices

Month	*Mean*	*Median*	*Percentiles* *5th*	*25th*	*75th*	*95th*
January	0.3	-0.4	-18.6	-3.9	3.5	18.4
February	0.8	0.2	-12.7	-2.9	3.1	15.8
March	1.3	0.6	-13.6	-2.7	3.8	17.2
April	0.6	-0.3	-16.0	-3.5	3.3	18.9
May	1.9	0.7	-13.9	-2.8	4.9	22.3
June	0.0	-0.5	-21.6	-4.1	3.3	17.7
July	0.7	-0.5	-16.0	-3.4	3.3	18.9
August	0.8	-0.2	-13.3	-2.9	3.3	16.7
September	0.7	0.1	-13.2	-3.0	3.2	15.5
October	1.1	0.3	-13.8	-2.9	3.8	18.4
November	1.3	0.5	-13.0	-2.8	3.9	19.0
December	1.0	0.4	-16.3	-2.9	3.6	17.6

Chart 6: Summary measures of price changes

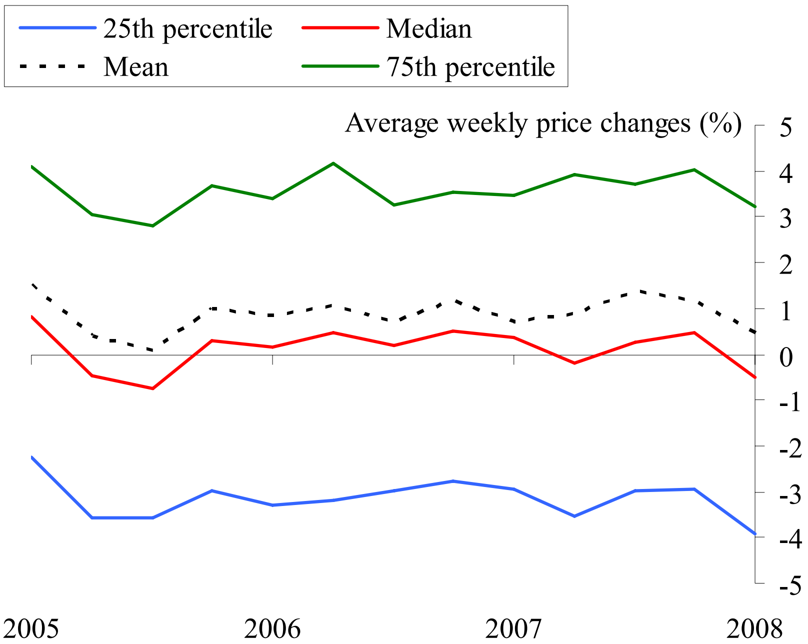

3.6 Frequency and magnitude of price changes

Having examined the frequency and magnitude of price changes separately, this section considers linkages between the two. If prices are set intermittently, larger price changes may occur when the duration of the previous price is large.

Chart 7 is a scatter plot of average and median price changes against the duration of the previous price that was set. One important point to note is that there are relatively few observations beyond three weeks duration, as most prices change more frequently than that. There is little sign that longer-lasting prices are (eventually) changed by greater amounts than prices with shorter durations. Indeed, econometric investigation confirmed that there were no significant or stable relationships between the magnitude and frequency data in Chart 7.

Chart 7: Frequency and magnitude of price changes

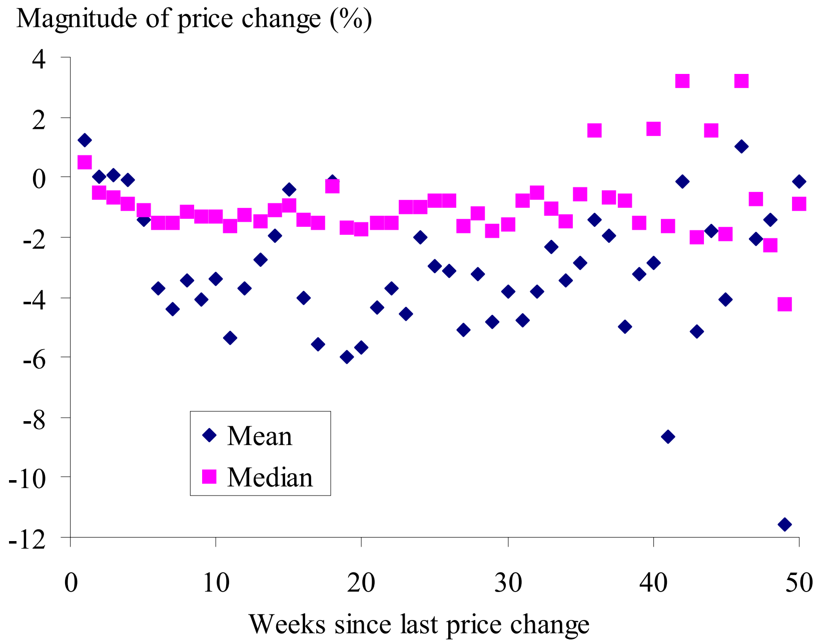

In part, this could reflect large positive and negative changes offsetting each other. So Chart 8 plots similar series, but this time for the average and median absolute price change. Once again, there is no sign of a stable relationship, either from the chart or econometric analysis based on the same data.

Chart 8: Frequency and absolute magnitude of price changes

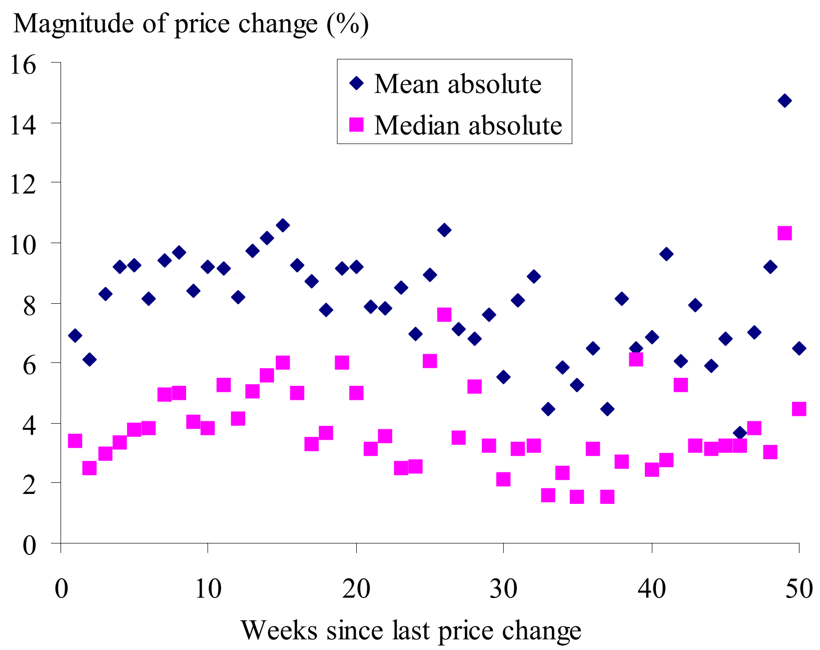

4 Comparing movements in volumes and prices

4.1 Changes in volumes

One advantage of the Nielsen data set is that sales data are included alongside prices. This enables some examination of the relationship between changes in price and changes in volume.

Table 10 reports summary statistics on the distribution of volume changes in the data set. It is readily apparent that volume changes are more dispersed than price changes (Table 7): the interquartile range is 24.7pp, compared to 6.7pp for price changes. The skew in the distribution of volume changes also appears to be larger than for prices, reflected in the wider gap between the mean and the median.

Table 10: Percentage changes in volume

	Per cent
Mean	22.4
Median	-0.3
5th percentile	-44.3
10th percentile	-28.6
25th percentile	-11.5
75th percentile	13.2
90th percentile	40.0
95th percentile	74.5

This larger skew is also evident when the whole distribution of volume changes is plotted (Chart 9), particularly in the longer right-hand tail. Large changes in volume are evident in all product categories (Table 11), albeit to a lesser extent for Bakery, Dairy and Fresh, suggesting these categories exhibit less volatile demand, consistent with less substitution between individual products.

Chart 9: Distribution of volume changes

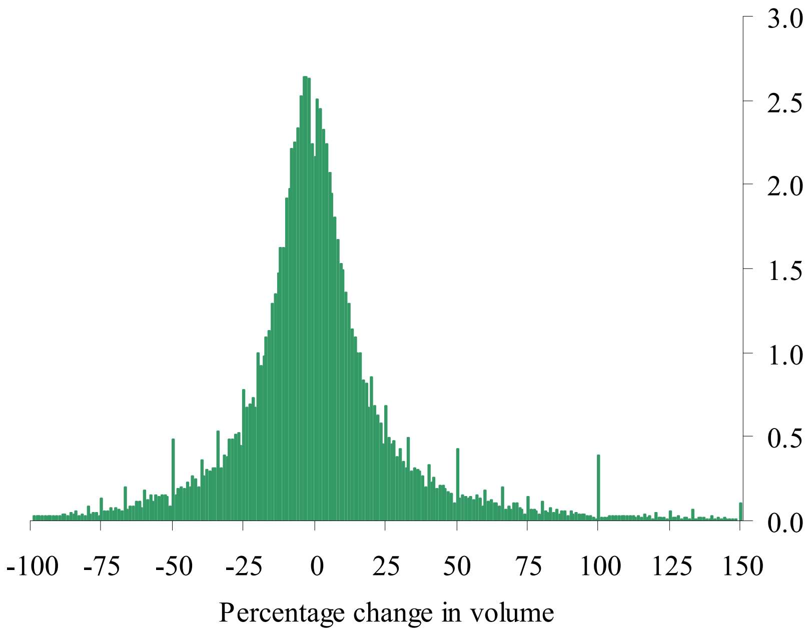

Table 11: Changes in volume by product category

Percentage change in volumes

Category	*Mean*	*Median*	*Percentiles* *5th*	*25th*	*75th*	*95th*
Alcohol	41.20	1.8	-66.7	-27.3	37.8	200.0
Bakery	7.40	1.1	-39.1	-11.6	14.7	60.0
Confectionary	22.70	2.3	-60.0	-21.1	29.0	133.3
Dairy	8.60	0.2	-33.5	-8.7	9.9	46.2
Fresh	9.10	-0.4	-30.9	-9.1	9.9	46.4
Frozen	34.00	-1.2	-58.3	-20.9	25.0	136.8
Grocery	33.90	0.2	-53.4	-17.0	20.0	106.9
Household	225.60	-0.8	-80.0	-27.3	33.3	454.4
Personal	26.30	1.6	-62.5	-27.8	38.5	166.7
Soft Drinks	36.20	-0.7	-60.9	-22.3	27.5	150.7

4.2 Correlations between volumes and prices

The previous section examined the magnitude of volume changes from week to week. One interesting question is whether these changes in volume lead or lag price changes. In practice, prices and volumes are jointly determined by the interaction of supply and demand. But it may be the case that consumers are surprised by changes in prices, and take time to adjust their spending patterns, or that producers take time to change their prices in response to shifts in demand.

Chart 10: Cross-correlations between price and volume changes

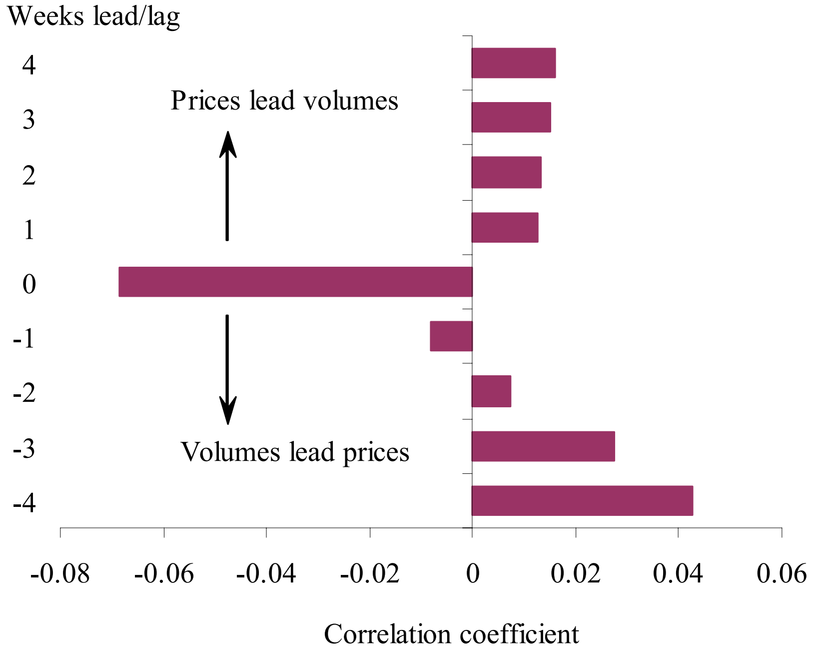

Do prices tend to move at the same time as volumes? Simple pair-wise correlations return relatively low estimates, but in fact given the large sample sizes these were often significant (Chart 10). The positive correlation where volume leads price is consistent with stronger (weaker) demand leading to higher (lower) prices; whereas the negative contemporaneous correlation is consistent with customers responding quickly – ie in the same week – and increasing (decreasing) purchases of items with price cuts (rises). However, while these correlations are useful in showing how the data behave, without more detailed information on these changes, we cannot establish whether price or volume changes (or both!) reflect either demand or supply shocks. As such, these results should not be overinterpreted. In the data, prices and volumes move coincidently; but further inference requires more information or more assumptions about the underlying behaviour of firms and consumers.

4.3 Constructing proxies for product elasticities

The strongest correlation between prices and volumes occurs contemporaneously, indicating that prices and volumes tend to move together. But how much do volumes change when prices change, and *vice versa*? Ideally, the way we should answer this question would be to construct formal price elasticities of demand (PEDs).

Technically, in order to examine how demand (volume) responds to price changes, other variables such as income, preferences, expectations and seasonal purchases should be taken into account. All of these factors could result in the demand curve shifting either out or in. In these instances, *ceteris paribus*, we would observe changes in prices and quantities while the elasticity of demand – the slope of the demand curve – may actually be unchanged. In the same manner, we should also control for movements in the supply curve. By accounting for all of these factors, we can attempt to isolate shifts in the demand curve from movements *along* the demand curve – which will provide genuine measures of the elasticity of demand.

Unfortunately, these demand variables are not readily available on a weekly basis. Instead, the assumption I make here is that the high frequency of the data itself acts as a natural control for shifts in the demand curve. The underlying assumption is that, for most households, income and preferences (etc) do not change from week to week. This implies that weekly changes in prices and volumes are more likely to reflect movements along the demand curve – as, by assumption, the determinants of demand to not change as often.

Of course, in the limit this assumption is almost certainly wrong. In any given week, some households will experience large changes in their circumstances – for example, losing their jobs. And in other weeks seasonal effects may drive the consumption patterns of many consumers. Both of these instances – low-frequency changes in households' circumstances, or temporary seasonal changes in demand – would affect the weekly price and volume data, making it difficult to recover a formal PED.

However, these factors do not affect all households for every week in the data sample – Christmas only happens once a year and only a fraction of workers lose their jobs in any given week. As such, where these large shifts do occur – be they either seasonal fluctuations or large changes in some individuals' circumstances – the resultant impact on prices and volumes should show up in the tails of the respective distributions for the data sample as a whole. As such, by focusing on the median of the distribution, rather than the mean, the impact of infrequent changes in these potentially large demand factors can be excluded. Essentially, this approach assumes that median weekly (high-frequency) price and volume changes are supply driven, rather than reflecting demand – or, put another way, it assumes that changes in the demand determinants drive large rather than small changes in prices and volumes at the individual store and product level. Clearly this is still an oversimplifying assumption – but it does offer a way of calculating approximate PEDs by exploiting the weekly frequency of the Nielsen data.

The approximate PEDs (*apeds*) are therefore calculated as the medians of the distribution of the ratio of volume changes to price changes:

$$aped_i = \frac{\%\Delta volume_i}{\%\Delta price_i}$$

Table 12 presents summary statistics for the distribution of approximate PEDs by product category. The estimated positive elasticities at the higher end of the distributions – which imply that volumes rise when prices rise – is likely to precisely reflect changes in the determinants of demand other than price, such as income or seasonal effects, that are not controlled for. Indeed, the lack of any formal controls for these non-price factors contributes to the considerable variability in the data, although some of this variability could also reflect genuine instability – ie the response of sales to a given price change may not be constant in the data sample, either over time or by magnitude.

Table 12: Approximate price elasticities of demand by product category

Category	*Percentiles* 5th	25th	50th	75th	95th
Alcohol	-190.7	-22.8	-5.1	5.8	135.2
Bakery	-24.9	-6.0	-1.6	2.2	17.4
Confectionary	-54.2	-10.6	-2.7	2.9	36.3
Dairy	-32.9	-6.4	-1.2	2.5	24.1
Fresh	-16.2	-2.9	-0.5	2.0	14.0
Frozen	-70.8	-11.7	-2.6	1.8	37.5
Grocery	-44.8	-8.2	-2.2	2.6	34.5
Household	-268.5	-23.0	-3.4	7.6	119.5
Personal	-101.3	-13.0	-3.0	4.7	72.9
Soft Drinks	-43.0	-9.3	-2.7	1.2	32.4

As discussed, in order to make any inference about elasticities it is most sensible to focus on the median estimates in Table 12. For most products, these indicate relatively elastic demand, as the magnitudes of the PEDs are greater than one. This indicates that volumes tend to change by proportionately more than prices, consistent with the distribution of volume changes being more dispersed than the distribution of price changes (Charts 5 and 9). The exception is for 'Fresh' products – this suggests consumers buy a steadier volume of fruit and vegetables as prices change, compared to other products. However, this may reflect 'Fresh' products being defined in terms of large catch-all categories, whereas other product types are more precisely branded, reflecting greater product differentiation. Interestingly, those products that tend to be more storable over time – such as alcohol and household goods – exhibit higher elasticities, consistent with consumers 'stocking up' when prices fall. So while these results are clearly approximations at best, and must be treated with caution, there is a sensible economic interpretation of the pattern of estimates that is uncovered.

5 Conclusions

This paper has examined how prices in UK supermarkets behave, using scanner data from Nielsen that are available on a weekly frequency – in all, the data set accounts for a little under 5% of annual household expenditure. This paper therefore adds to the growing literature of micro-pricing studies.

Using these data, several interesting features emerge about how prices behave. First, prices change very frequently in supermarkets – 40% of prices change each week (excluding 'Fresh' items). Some of these price changes are likely to be temporary – but even when we control for these by excluding price reversals or smooth through temporary price falls and look at 'regular' prices, we still find that roughly a quarter of prices change each week. Importantly, there is also evidence that focusing on monthly observations, rather than weekly ones, overstates the implied stickiness of prices. Overall, the results suggest that prices in supermarkets exhibit a greater degree of flexibility than may be evident in other sectors. Second, the probability of price changes is not constant over time – all product categories have declining hazard functions. Third, the range of price changes is very wide, with some very large price cuts and price rises; but despite this, a significant number of price changes are very small. Fourth, there appears to be little link between the frequency and magnitude of price changes – prices that change less frequently do not tend to change by more. Fifth, volume changes exhibit at least as much variation as price changes, and the strongest correlation between the two is contemporaneous, suggesting that prices and volumes move together from week to week. And sixth, rough analysis based on simplifying assumptions suggests that consumers are fairly price sensitive.

Overall, these results suggest that price stickiness is not a key factor in the UK supermarket sector: indeed, nominal rigidities appear to be quite limited. Prices and volumes change frequently, potentially by quite a lot, and there is significant heterogeneity in the patterns of both prices and volumes.

6 Annex: The 'regular price' algorithm

This annex describes the 'regular price' algorithm used to define temporary sales in the main paper. It is based on Kehoe and Midrigan (2007).

In their paper, KM construct the regular price series, P_t^R, from the original price series, P_t, as follows. Whenever the actual price series falls, ie $P_t < P_{t-1}$, check if the actual price rises above its current (new) level over the next five weeks: ie, check if $P_{t+j} \geq P_t$ for $j \leq 5$. If it does, then define J as the first time the actual price rises above its current (new) level, P_t.

To construct the regular price series, replace $P_t, P_{t+1},\ldots,P_{t+J-1}$ with P_{t-1}. If the price never rises above P_t within the next five weeks, leave P_t unchanged. This is repeated at different horizons, to allow for a variety of sales patterns. For example, if the actual price series was:

100, 95, 96, 97, 98, 99

then using the algorithm at $t+1$ would yield

100, 100, 96, 97, 98, 99

but using it up to $t+5$ would yield

100, 100, 100, 100, 100, 99

In this way the algorithm was 'repeated' by sequentially examining prices over decreasing window. So in the first instance, the algorithm examined if $P_{t+5} > P_t$, and replaced $P_t, P_{t+1},\ldots, P_{t+4}$ if they were (each) below P_{t+5} and P_{t-1}. In the second instance, the procedure was repeated by examining if $P_{t+4} > P_t$, and replacements ensued on a similar basis. As five-week sales periods were already replaced in the first instance, there is no risk of double-counting as the inequality conditions will not be met in the second instance. The resulting series for the sample data is the same as is shown above, but is also robust to other patterns: eg 100, 90, 85, 102, 98 becomes 100, 100, 100, 102, 98.

References

Alvarez, L, Burriel, P and Hernando, I (2005), ‘Do decreasing hazard functions for price changes make any sense?’, *ECB Working Paper No. 461*.

Amirault, D, Kwan, C and Wilkinson, G (2005), ‘A survey of the price-setting behaviour of Canadian companies’, *Bank of Canada Review*, Winter.

Bank of England (2006), *Agents’ summary of business conditions*, October.

Baudry, L, Le Bihan, H, Sevestre, P and Tarrieu, S (2007), ‘What do thirteen million price records have to say about consumer price rigidity?’, *Oxford Bulletin of Economics and Statistics*, Vol. 69, pages 139-83.

Benford, J and Driver, R (2008), ‘Public attitudes to inflation and interest rates’, *Bank of England Quarterly Bulletin*, Vol. 48, No. 2, pages 148-56.

Bils, M and Klenow, P (2004), ‘Some evidence on the importance of sticky prices’, *Journal of Political Economy*, Vol. 112, pages 947-85.

Blinder, A, Canetti, E, Lebow, D and Rudd, J (1998), *Asking about prices: a new approach to understanding price stickiness*, New York: Russell Sage Foundation.

Boivin, J, Giannoni, M and Mihov, I (2007), ‘Sticky prices and monetary policy: evidence from disaggregated US data’, *National Bureau of Economic Research Working Paper No. 12,824*.

Bunn, P and Ellis, C (2009), ‘Price-setting behaviour in the United Kingdom: a microdata approach’, *Bank of England Quarterly Bulletin*, Vol. 49, No. 1, pages 28-36.

Calvo, G A (1983), ‘Staggered prices in a utility maximising framework’, *Journal of Monetary Economics*, Vol. 12, pages 383-98.

Chevalier, J, Kashyap, A and Rossi, P (2000), ‘Why don’t prices rise during periods of peak demand? Evidence from scanner data’, *NBER Working Paper No. 7981*.

Driver, R and Windram, R (2007), ‘Public attitudes to inflation and interest rates’, *Bank of England Quarterly Bulletin*, Vol. 47, No. 2, pages 208–23.

Druant, M, Fabiani, S, Hernando, I, Kwapil, C, Landau, B, Loupias, C, Martins, F, Mathä, T Y, Sabbatini, R, Stahl, H and Stokman, A C J (2005), ‘The pricing behaviour of firms in the euro area: new survey evidence’, *ECB Working Paper No. 535*.

Eichenbaum, M, Jaimovich, N and Rebelo, S (2008), ‘Reference prices and nominal rigidities’, *NBER Working Paper No. 13829*.

Fougère, D, Le Bihan, H and Sevestre, P (2005), 'Heterogeneity in consumer price stickiness; a microeconometric investigation', *ECB Working Paper No. 536*.

Gali, J and Gertler, M (1999), 'Inflation dynamics: a structural econometric analysis', *Journal of Monetary Economics*, Vol. 44, pages 195-222.

Greenslade, J and Parker, M (2008), 'Price-setting behaviour in the United Kindgom', *Bank of England Quarterly Bulletin*, Vol. 48, No. 4, pages 404-15.

Kehoe, P and Midrigan, V (2007), 'Sales and the real effects of monetary policy', *Federal Reserve Bank of Minneapolis Working Paper No. 652*, April.

Mankiw, N G (1985), 'Small menu costs and large business cycles: a macroeconomic model of monopoly', *Quarterly Journal of Economics*, Vol. 100, pages 529-37.

McAdam, P and Willman, A (2007), 'State-dependency and firm-level optimization: a contribution to Calvo price staggering', *ECB Working Paper No. 806*.

Mumtaz, H, Zabczyk, P and Ellis, C (2009), 'What lies beneath: what can disaggregated data tell us about the behaviour of prices?', *Bank of England Working Paper No. 364*.

Nakamura, E and Steinsson, J (2008), 'Five facts about prices: a re-evaluation of menu-cost models', Columbia University, *mimeo*.

Nielsen (2007), 'Top 100 grocery brands', March, checkoutmagazine.co.uk.

Rotemberg, J (1982), 'Sticky prices in the United States', *Journal of Political Economy*, Vol. 90, pages 1,187-211.

Sheshinski, E and Weiss, Y (1977), 'Inflation and costs of price adjustment', *Review of Economic Studies*, Vol. 44, pages 287-303.

Smets, F and Wouters, R (2003), 'An estimated stochastic dynamic general equilibrium model of the euro area', *Journal of the European Economic Association*, Vol. 1, pages 1,123-75.

Taylor, J (1980), 'Aggregate dynamics and staggered contracts', *Journal of Political Economy*, Vol. 88, pages 1-23.