BANK OF ENGLAND

Working Paper No. 438

How do individual UK consumer prices behave?

Philip Bunn and Colin Ellis

October 2011

BANK OF ENGLAND

Working Paper No. 438

How do individual UK consumer prices behave?

Philip Bunn[1] and Colin Ellis[2]

Abstract

This paper examines the behaviour of individual consumer prices in the United Kingdom, and uncovers a number of stylised facts about pricing behaviour. First, on average 19% of prices change each month, although this falls to 15% if sales are excluded. Second, the probability of price changes is not constant over time. Third, goods prices change more frequently than services prices. Fourth, the distribution of price changes is wide, although a significant number of changes are relatively small and close to zero. Fifth, prices that change more frequently tend to do so by less. We find that conventional pricing theories struggle to match these results, particularly the marked heterogeneity, which argues against the use of 'representative agent' models.

Key words: Consumer prices, price-setting behaviour.

JEL classification: E31, D40.

(1) Corresponding author. Bank of England. Email: philip.bunn@bankofengland.co.uk
(2) BVCA and University of Birmingham. Email: cellis@bvca.co.uk, c.ellis@bham.ac.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. This work was completed while Colin Ellis was employed at the Bank of England.

We would like to thank the Office for National Statistics (ONS) for making the data available to us. We are also grateful to Neal Hatch, Steve Millard, Nick Oulton, and seminar and workshop participants at the Bank of England for useful comments. All errors remain our own. This paper was finalised on 5 July 2011.

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

It is important for monetary policy makers concerned with meeting an inflation target to consider how prices behave. Nominal rigidities imply that prices cannot freely adjust, and the degree of nominal rigidity in the economy will influence the short-term impact of monetary policy on real activity and hence the response of inflation. This paper uses a database of over 11 million price quotes to investigate how individual consumer prices behaved in the United Kingdom between 1996 and 2006. These are the microdata that underpin the monthly Consumer Prices Index produced by the Office for National Statistics. This work enables us to establish the facts about how frequently consumer prices change and how much they change by when they do change, and it should help us to improve our understanding of the nature of the nominal rigidities that exist in the economy. The results also help to establish which theories of pricing behaviour most closely represent the way in which prices are set in the real world, or at least in the UK economy.

This paper is the first to examine how UK consumer prices behave using the individual price quotes underlying the published aggregate inflation measure that is targeted by the Bank of England. This paper complements similar work on producer prices, which examines how prices behave further up the supply chain, and a recent survey of how firms set prices that was carried out by the Bank.

We find that 19% of consumer prices change each month on average, although this falls to 15% if sales are excluded. There is little evidence to support the presence of downward nominal rigidities in product markets, since 40% of all consumer price changes are decreases. UK consumer prices appear to be slightly more flexible than in the euro area, but they are less flexible than in the United States.

Consumer goods prices change more frequently than those of services, as on average 24% of goods prices change each month, compared with only 9% of services prices. At the component level, the prices of energy goods change the most frequently. The main service sector components all display a similar degree of price stickiness.

The share of prices changing each month varies across different years of our sample. There is some correlation between the share of prices increasing and the aggregate consumer price inflation rate. There are also some seasonal effects: prices are most likely to change in January and April and least likely to change in November and December. For consumer goods prices, the probability of a price change is highest in the month immediately following the previous change. As more time passes since the last price change, the probability of a price changing in any given month declines. For services, prices are most likely to change a year after the previous change, suggestive of annual price reviews. The probability of services prices changing in other months is broadly constant.

The distribution of the size of price changes is wide, although a significant number of changes are relatively small and close to zero. Around 60% of all price changes are between -10% and 10%, and the modal price change is an increase between 1% and 2%. The distribution of the size of consumer price changes narrows a little if sale prices are excluded. There are more small increases in prices and fewer price cuts for services than there are for goods, but there are considerable differences in the shape of the distributions of price changes at the component level. Prices that change more

frequently tend to do so by less. This relationship appears to be particularly strong for services prices, but it also hold for goods prices as well once the effects of sales are taken out.

Our results on the behaviour of UK consumer goods prices are similar to those from previous work on UK producer prices (which covers only goods and not services). This suggests that there are few pricing frictions between the production and retail sectors in the United Kingdom.

Our findings from the microdata are not consistent with any one theory of price-setting. The marked heterogeneity that we observe in the behaviour of prices in different parts of the economy suggests that different theoretical models may better explain how prices are determined in different sectors. This would argue against the use of 'representative agent' models. The challenge is to develop a new theory of price-setting that better fits the stylised facts observed in these micro-studies while also fitting the properties of the aggregate macrodata.

1 Introduction

Nominal rigidities imply that prices cannot freely adjust to changes in supply and demand. These nominal rigidities are important because they allow changes in monetary policy to affect real output. With fully flexible prices, changes in monetary policy will immediately be offset by changes in nominal variables such as inflation with no effect on real output. The degree of nominal rigidity is a key part of the monetary transmission mechanism and understanding more about these rigidities is therefore highly relevant for monetary policy makers who are setting monetary policy to meet an inflation target. Understanding more about price dynamics may be important in judging what the appropriate policy response is developments in the economy.

A range of different mechanisms for modelling price stickiness have been put forward. These can be categorised under two main headings: state-dependent and time-dependent pricing models. In state-dependent models the decision to change price depends on the state of the economy and the market faced by firms. Firms are typically assumed to incur some type of cost when changing their price. Examples of these types of costs include relatively small fixed costs of changing price – or menu costs (Mankiw (1985)), or disutility associated with making large price changes (Rotemberg (1982)). Whatever the nature of these costs, prices change at irregular intervals because a firm will only change its price when the cost of a price adjustment is outweighed by the extra profit the firm can make by changing price. In a time-dependent model, the probability of a price change depends only on the time since the previous price change. This class of models includes the popular Calvo (1983) model in which firms are assumed to have a fixed probability of changing price in any given period.

A range of different approaches have been taken to estimate the degree of price stickiness. At a macro level the New Keynesian Phillips Curve (NKPC), which relates current inflation to expected future inflation and the deviation of marginal cost from its steady state, incorporates an assumption about the nature of price-setting behaviour. By estimating the NKPC it has been possible to estimate how often prices change. Using the Calvo assumption that prices have a constant probability of changing in each period, Gali and Gertler (1999) estimated that firms change price every 15 to 18 months in the United States. Smets and Wouters (2003) found that the implied duration between price changes in the euro area was longer at two and half years.

At a micro level, the two main approaches to measuring the extent of nominal rigidities have been to survey firms to ask them how often they change their prices, or to study how often prices actually do change using micro data sets. By looking at micro price data we can establish what the facts are about how often prices change using very large data sets. In practice, pricing surveys tend to use much smaller samples, but they can help to tell us how firms set prices and the reasons why they change prices. Pricing surveys can inform us about how frequently firms review prices as well as how often they actually change them, but using micro price data allows the analysis of pricing behaviour at different points in time rather than just in one particular period.

Implied durations from either type of micro-based source tend to be shorter than the macro-based estimates. In a recent survey of UK firms, Greenslade and Parker (2008) find that the median UK firm reviews price twice a year, but only changes price once a year. But the distribution of price durations is bimodal, with a large number of companies changing their prices even more frequently

than this. The UK survey results of Greenslade and Parker (2008) are consistent with similar evidence for other countries. Fabiani *et al* (2006) and Blinder *et al* (1998) also find that the median firm changes price once a year in the euro area and the United States respectively. Estimates based on price microdata used in the construction of aggregate inflation indices tend to imply that prices change more frequently than these survey results suggest. Using the microdata that underlies the euro-area CPI, Dhyne *et al* (2006) find that 15% of consumer prices change each month, while the median price duration is around eleven months. Evidence from analysis of the US CPI microdata suggests that consumer prices change more frequently in the United States: Bils and Klenow (2004) find that 26% of prices change each month.[1] Price quotes for use in constructing aggregate inflation indices are typically only collected once a month, which may mean that the true frequency with which prices change is underestimated because intramonth prices changes are excluded. Studies based on higher frequency data, such as scanner data (Kehoe and Midrigan (2007) and Chevalier *et al* (2000)) suggest that prices change more frequently than estimates based on microdata used to construct inflation data, albeit using a less representative sample. Recent evidence for the United Kingdom based on weekly scanner data found that 40% of supermarket prices change each week (Ellis (2009)).

Previous estimates of price stickiness in the United Kingdom based on the microdata used to construct the Producer Price Index (Bunn and Ellis (2010)) found that an average of 26% of manufacturing output prices change each month. There is no previous work using UK consumer prices. Our paper is the first to do this using data that has been made available for the first time by the Office for National Statistics (ONS). This adds to our previous work on producer prices because it allows us to establish how prices behave in the measure of inflation targeted by the Bank's Monetary Policy Committee. Analysing consumer prices will tell us how prices behave further down the supply chain than our work on producer prices, since this will cover the prices charged by retailers selling goods directly to consumers rather than the prices charged by firms manufacturing goods. Looking at consumer prices also enables us to study the behaviour of UK services prices for the first time. The CPI data set is much larger than the producer price data and our sample has a longer time series dimension than the data used in our producer price work.

Several interesting stylised facts emerge from our analysis of UK consumer prices. First, on average 19% of prices change each month, although this falls to 15% if sales are excluded. Second, the probability of price changes is not constant over time. Third, the probability of prices changing varies significantly between goods and services prices, and within this there is considerable heterogeneity between components. Fourth, the distribution of price changes is wide, although a significant number of changes are relatively small and close to zero. Fifth, prices that change more frequently tend to do so by less.

The next section of the paper discusses the details of our data. We then present the results of our analysis and draw out the stylised facts that emerge. This starts with the frequency of price changes and hazard functions, and is followed on by analysis of the size and distribution of price changes and the relationship between the frequency and size of price changes. We then look at the implications of our results for pricing theory and monetary policy. Finally, we conclude.

[1] A similar figure is found by Nakamura and Steinsson (2008) using a larger and more recent sample.

2 Data

The data we use in this study are individual consumer prices, collected by the Office for National Statistics. These price quotes are weighted and aggregated to form consumer price indices (CPIs). The aggregate CPI is calculated from locally collected data, where ONS price collectors go into shops and record selling prices, and centrally collected data, where the ONS collect 'national' prices from particular companies. One example of a centrally collected price would be the price of a national newspaper. Unfortunately, only locally collected prices were available for this study. Not having access to the full CPI sample is not a problem unique to our study. Similar analysis of the microdata underlying consumer price indices for Belgium (Aucremanne and Dhyne (2004)) and Spain (Alvarez and Hernando (2004)) were also restricted to locally collected data. And studies in Germany (Hoffmann and Kurz-Kim (2006)), Italy (Veronese *et al* (2005)) and the Netherlands (Jonker *et al* (2004)) were only able to use a subsample of 50 representative products.

The locally collected data we did use covers around two thirds of the aggregate CPI by weight. Our sample covers the period between 1996 and 2006 and includes a total of just over 11 million individual price quotes.[2] Prices are recorded monthly on the CPI collection day, which is usually the second or third Tuesday of the month. Since price quotes are only collected once a month, our analysis will exclude any price changes that take place between the two collection days. By excluding these intramonth price changes we may be underestimating the true frequency of price changes and therefore overestimating the duration between price changes.

Our price quotes are those collected by the ONS: they are the price of a particular product in a particular shop in a given month. Each of these product-locations (which are subsequently referred to as items) is tracked individually, and in total there are just under 600,000 different items in the data set. As with other microdata, the panel is not balanced: new products enter and existing products exit and the locations at which prices are collected also change. The sample is updated annually, in February, to ensure that it remains representative, although there is still some rotation in the intervening period because the prices of specific items may no longer be available. Unfortunately, the precise reason for an item leaving the sample is not available in the underlying data set. There is also significant attrition in the data as individual products are modified, dropped or change location: only 96 items are in the data every month across the whole eleven year sample. In all, around 700 items are in the data for more than ten years and 17,000 are present for at least five years. The mean number of months in the sample for an item is 19, the median is 13. Unless otherwise stated, the results presented in the paper are on a weighted basis, where these weights were supplied by the ONS. The weights represent the share of each item in the locally collected CPI in each month.[3]

ONS themselves examine the microdata on prices carefully ahead of aggregating to produce a headline price index. As such, there were a small percentage of prices in the microdata that were not used in the constriction of the headline CPI data – we exclude these observations too. These tend to

[2] The CPI series starts in 1996.

[3] ONS collect larger samples of prices for some types of products where they believe it is necessary to produce a reliable estimate of the average price. Weighting the results avoids biasing them towards these types of products. The weights are based on expenditure, and they represent the individual weight of each particular item in the aggregate CPI in each month.

be outliers or 'zero' price quotes. We also dropped observations where there was no price quote for the corresponding item in the previous month, since for the month that followed we were not able to identify whether the price had changed. This is arguably akin to left-censoring the data, although only by one observation for each item. Our cleaned sample represents approximately 85% of the full set of locally collected CPI data.

Given that our sample covers only locally collected CPI data it may not be fully representative of all consumer prices. One potential source of bias in our results is that some prices are more likely to be collected locally than centrally, which will mean that our microdata sample will have a higher weight in those items than in the published CPI. Chart 1 shows the average weight within our microdata sample (our cleaned locally collected data) of each component of CPI compared to the weights within the published CPI data.[4] Some components such as food and non-alcoholic beverages have a higher weight in our microdata than in published CPI, because these prices are more likely to be collected locally than centrally. The only component for which we have very little microdata is communication, but as this component only accounts for around 2% of published CPI on average it should not be large enough to significantly affect our results. We also have no locally collected microdata on education prices (included under miscellaneous services in Chart 1), but again this only accounts for a small part of published CPI, around 1% over our sample period.

Chart 1: Weights by CPI component

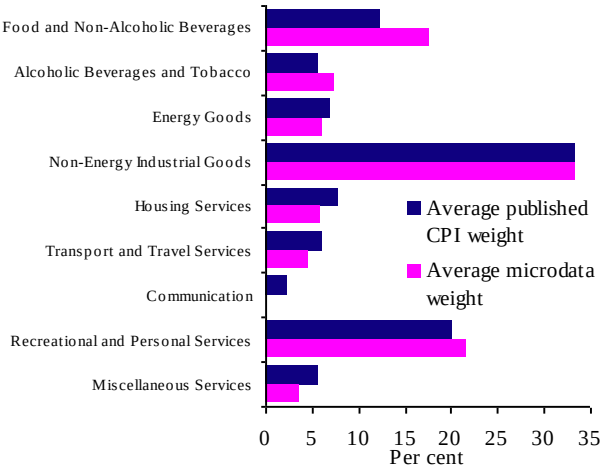

Chart 2: Percentage of price quotes that are sales by calendar month

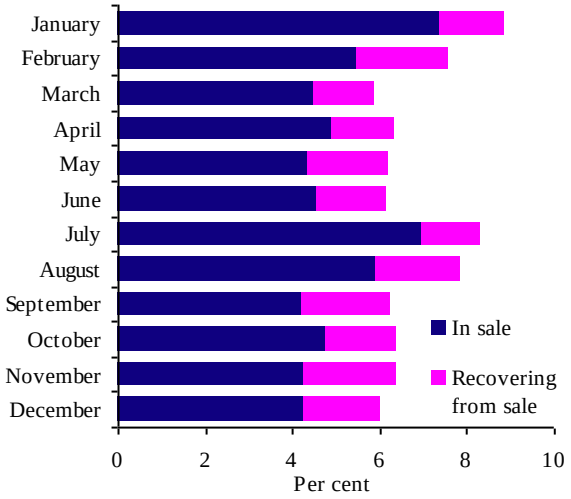

In general, the differences between the microdata and the published weights by component are relatively small. In the analysis of our results we carry out some robustness checks to assess if this has much effect on our results. A second type of bias could occur in our results if the locally collected prices within each component are not representative of the centrally collected items. In many cases, it may be reasonable to assume that they are representative. But on close inspection of our data there is one obvious example of where this may not be true: the prices of energy goods. Our microdata sample of energy goods prices is dominated by petrol and diesel prices. The other major group of energy goods in CPI are gas and electricity utility prices, and these are centrally collected.

[4] Table A3 in the appendix presents a similar analysis for the data broken down by COICOP division (Classification of Individual Consumption by Purpose – this is an alternative decomposition of CPI which is presented in the CPI first release along with the breakdown by CPI component).

It is likely that petrol and utility prices behave differently; our prior would be that petrol prices change more often than utility prices. We interpret all our results with the caveat in mind that they may not be fully representative of CPI and that the energy goods component represents mainly petrol prices.

Helpfully, ONS price collectors mark whether a particular price in any given month is a 'sale' price – strictly, a temporarily discounted price, although we use the ONS's 'sale' shorthand throughout this paper – or not. So by using these identifiers, we can examine the role that sales play. We are also able to identify prices that are recovering from a sale or temporary promotion in the previous month. The CPI convention is that sale prices should only be recorded as such if they are available to everybody.[5] On average, approximately 5% of all price quotes are items that are marked as sale prices, and 2% are prices recovering from a sale. As an item can be 'on sale' for more than one month, a price being identified as being in the sale does not automatically imply that the price has been reduced on the month. Prices are said to be 'recovering from a sale' where the price increases and the item had been identified as being on sale in the previous month.[6] This partly explains why there are more 'sale' price quotes than 'recovering from sale' quotes. But, in addition, some items do not recover their price and they drop out of the sample after they have been in the sale, for example in the case of end of season reductions.[7]

Chart 2 shows how the share of prices associated with sales varies in different months of the year. January is the most popular month for sales followed by July. Unsurprisingly, February and August, the months immediately following January and July, are the months where the highest proportion of prices recover from sales. Sale prices are mainly associated with goods. Less than 1% of services price quotes are recorded as sale prices, whereas 11% of goods quotes are either sale prices or prices recovering from a sale. Within the goods component of CPI, sales are most prevalent among non-energy industrial goods, but very few energy goods price quotes are associated with sales.

All our results assume there is no measurement error in the collection of the underlying CPI microdata. The data were accessed and analysed using the ONS's Virtual Microdata Laboratory (VML). Ritchie (2008) describes the history of the VML, and the detailed terms and conditions that apply to users. The next section presents results from our analysis of the data.

3 Stylised facts on UK consumer prices

This section of the paper presents a set of stylised facts on UK consumer price changes. We start by analysing the frequency of price changes before moving on to look at the conditional probability of price changes (estimating hazard functions). We look at the magnitude of price changes and then the relationship between the frequency and magnitude of price changes.

[5] Money-off coupons, vouchers and loyalty card discounts are excluded in the CPI data. Prices associated with the purchase of another product, eg 3 for 2 offers, are ignored. And prices are not adjusted for temporary increases in quantities, eg 20% extra free offers.

[6] In the small number of cases where the sales identifier is no longer present when it had been in the previous month but where the price does not change, the price is not identified as recovering from a sale.

[7] Prices in closing down sales or where the goods are not of comparable quality to the previous month should not be included in the CPI aggregates.

3.1 Frequency of price changes

3.1.1 Aggregate frequency of price changes and comparisons with other countries

Across the sample as a whole, approximately one in five consumer prices changes each month. Table 1 shows that an average of 19% of UK consumer prices changed each month between 1996 and 2006. This is calculated as the total number of price changes over the total number of price quotes. These results suggest that prices are not fully flexible because around 80% of prices do not change each month, although our analysis cannot tell us about how many firms review their price each month and decide not to change it. Of these price changes, approximately 60% are increases and 40% are decreases – matching the results found by Dhyne *et al* (2006) and Nakamura and Steinsson (2008). As with other studies of UK microdata (Bunn and Ellis (2010)), these results suggest that downward nominal rigidities are not a prevalent feature of the behaviour of prices in the United Kingdom.

Table 1 also shows the impact of ONS-defined sales on the share of consumer prices changing each month.[8] Excluding observations identified as sale prices or prices recovering from a sale, an average of 15% of consumer prices change each month. Sale prices change far more frequently, and – unsurprisingly – the majority of prices changes associated with sales are downward. However, the impact of sales on the observed frequency of price changes is not as pronounced as in some other studies, such as Ellis (2009).

Table 1: Percentage of UK consumer prices that change each month

	All changes	Increases	Decreases
Including sales	18.8	11.1	7.7
Excluding sales	15.3	9.8	5.5
Sales only	67.0	29.2	37.9

UK consumer prices appear to be slightly more flexible than consumer prices in the euro area. Table 2 shows that only 15% of prices change in the euro area each month (Dhyne *et al* (2006)), compared with 19% in the United Kingdom (when including sales). However, some euro-area economies, such as France and Finland, have similar frequencies to the United Kingdom. UK consumer prices appear to be less flexible than US consumer prices, based on comparisons with results from Bils and Klenow (2004) and Nakamura and Steinsson (2008). However, with all these comparisons it must be borne in mind that the data in different studies cover different time periods and different methodologies, particularly with respect to the treatment of sale prices. These factors could account for at least some of the differences in results, although inflation was relatively low and stable over the sample periods used for most of the studies in Table 2.

Based on French and Austrian data, Dhyne *et al* (2006) estimate that sales can account for up to 3 percentage points of the overall share of prices changing each month in the euro area. Nakaruma

[8] The sale category includes all price quotes that are identified by the ONS as being sale prices and all prices that are recovering from a sale.

and Steinsson (2008) estimate that this figure is around 5 percentage points for the United States. At 4 percentage points, our estimate for the United Kingdom lies between the two.

Table 2: Percentage of consumer prices that change each month

	All changes	Period covered	Study
United Kingdom	19	1996-2006	
Euro area[9]	15		Dhyne *et al* (2006)
- Austria	15	1996-2003	Baumgartner *et al* (2005)
- Belgium	17	1989-2001	Aucremanne and Dhyne (2004)
- Finland	20	1997-2004	Kurri (2007)
- France	19	1994-2003	Baudry *et al* (2004)
- Germany	10	1998-2003	Hoffmann and Kurz-Kim (2006)
- Italy	9	1996-2003	Veronese *et al* (2005)
- Luxembourg	17	1999-2004	Lünnemann and Mathä (2005)
- Netherlands	17	1998-2003	Jonker *et al* (2004)
- Portugal	22	1997-2001	Dias *et al* (2004)
- Spain	15	1993-2001	Alvarez and Hernando (2004)
United States	26	1995-1997	Bils and Klenow (2004)
	27	1998-2005	Nakamura and Steinsson (2008)

3.1.2 Time between price changes and price changes per item

A simple way to estimate the average duration between price changes is to calculate the inverse of the share of prices changing each month. On this measure, the average time between changes for UK consumer prices is 5.3 months, rising to 6.5 months if we exclude sales. But not all items change price at this frequency, and there is a distribution in the frequency of price changes across different items. This distribution is shown in Chart 3.[10] 27% of items have an average time between price changes of less than 3 months, while 75% of items have an average duration between price changes of less than 12 months. The median number of months between price changes per item is 7.2 months. The mean is higher at 10.6. The average time between price changes is therefore longer taking this second approach. This is because a relatively small number of items account for many price changes which means that the average duration is longer when measured by the average for individual products rather than by taking the simple average across all prices. As noted in Bunn and Ellis (2010), due to Jensen's Inequality, the concavity of the duration-frequency relationship means that averaging price durations within and then across product groups yields a higher average price duration than calculating an average across the data sample as a whole. This difference arises because of the heterogeneity that is present in the data.[11]

[9] Weighted average. Based on a common sample of 50 products across euro-area countries.
[10] Items are weighted in the distribution according to their weight in the pooled sample across our sample period.
[11] See below for further discussion of this heterogeneity.

Chart 3: Distribution of number of months between price changes per item

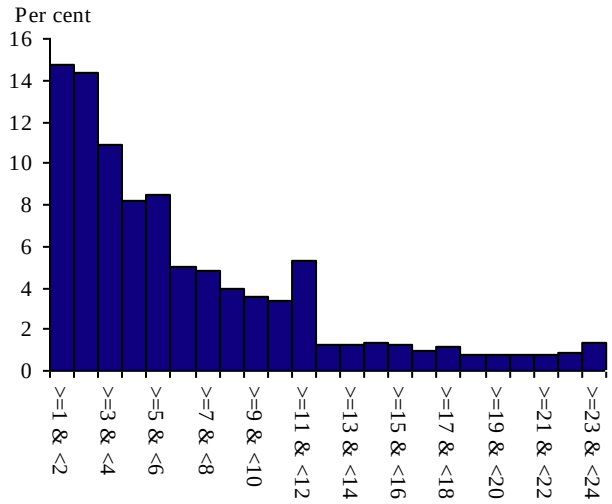

3.1.3 Frequency of price changes by product group

There is substantial variation in the frequency of UK consumer price changes between the different components of CPI, these are summarised in Table 3. Goods prices change more frequently than the prices of services: an average of 24% of goods prices change each month compared to only 9% for services. That is consistent with the results from a separate recent price-setting survey for the United Kingdom carried out by the Bank of England. Greenslade and Parker (2008) found that retailers change price once a month, which is more frequently than our implied average duration between price changes for CPI goods of around four months. The results for service sector firms are more comparable: the survey found that the average service sector company changes price once a year, which is consistent with our microdata based estimate of eleven months between price changes for service sector firms.

Our results suggest that consumer goods in the CPI collection basket change price less frequently than estimates based on higher frequency supermarket scanner data: Ellis (2009) found that 40% of goods sold in UK supermarkets changed price each week. Part of the reconciliation between these two results is that our estimates only use price quotes collected once a month, which will therefore exclude intramonth changes, whereas the scanner data in Ellis (2009) is weekly. But even given this, our results still suggest that the prices of goods sold in supermarkets change price more frequently than the prices of goods sold in other retail outlets in the United Kingdom. At least in part this is because there is more evidence of temporary price promotions in the supermarket data and because the supermarket sample was predominately food and drink items, which change price more frequently than other non-energy goods in the CPI data (see Table 3).[12]

Of the price changes we observe in the microdata, 45% of the changes in goods prices are falls, while only 20% of the changes in services prices are falls.[13] So services prices are less likely to fall than goods prices. This could reflect the fact that services price inflation was higher than goods price

[12] It was not possible to identify supermarkets in the CPI microdata since the names of retailers were removed from the version of the data we had access to.

[13] These figures are broadly comparable with results for the euro area: see Dhyne *et al* (2006).

inflation over our sample period and that goods are much more likely to be in the sale (which would involve price cuts) than services prices.[14]

There is considerable heterogeneity in the frequency of price changes within the different CPI goods components. The prices of energy products are the most flexible, with an average of 64% of all prices changing in any given month. As noted earlier, this refers mainly to petrol and diesel prices and does not include the prices of utilities. Among the other goods components, the prices of non-energy industrial goods change less often than the prices of food, beverages and tobacco. Almost half of all the changes in the prices of non-energy industrial goods are accounted for by sales.

Table 3: Percentage of consumer prices that change each month by CPI component

	All changes	Increases	Decreases	All changes ex-sales
Food and Non-Alcoholic Beverages	21.8	11.9	10.0	17.8
Alcoholic Beverages and Tobacco	25.7	18.3	7.5	21.6
Energy Goods	64.3	37.5	26.8	64.3
Non-Energy Industrial Goods	18.4	9.1	9.3	10.5
Housing Services	8.4	5.7	2.7	8.4
Transport and Travel Services	8.2	6.4	1.8	8.1
Communication	20.0	7.2	12.8	14.4
Recreational and Personal Services	8.9	7.2	1.7	8.7
Miscellaneous Services	8.8	7.6	1.2	8.7
All goods	24.4	13.5	10.9	19.4
All services	8.7	6.9	1.9	8.6

There is less heterogeneity between the five services components of CPI; four have a frequency of price change of between 8% and 9%. Communication service prices change more frequently than this, although we only have limited data for this component. Around two thirds of the prices of communication services in our sample are price decreases, which is likely to be related to the consistently negative inflation rates for this component.

An alternative decomposition of the frequency of price changes is to look at the frequency of price change by COICOP division.[15] The results from this approach are shown in Table A4 of the appendix, and as with the split by component, they confirm that heterogeneity is present in the behaviour of UK consumer prices.

Using the data on the frequency of price change by component in Table 3 we perform a robustness check on the aggregate frequency of price change to see whether using the weights in our microdata sample that are slightly different to those in the published data (shown in Chart 1) appear to be

[14] The average CPI goods inflation rate over our sample period was close to zero, while service price inflation averaged approximately 4%.
[15] This is the international classification of individual consumption by purpose.

significantly affecting our results.[16] We weight together the frequency of price change for each component by the average weight in published CPI (rather than the weight in our microdata sample, which is cleaned and only covers locally collected data). When we do this, the average proportion of consumer prices changing each month is still 19% to one decimal place. If we perform a similar exercise to reweight the share of prices changing by COICOP division, we again get an overall figure of 19% of prices changing each month. Therefore, the fact the weights in our sample are not identical to those in published CPI does not appear to be significantly affecting our results. But we still have no way of testing whether the properties of locally collected prices within each component are similar to the centrally collected prices, and that remains a potential source of distortion in our results.

The final breakdown we report is a set of broad product categories that have been previously employed in euro-area studies, in particular Dhyne *et al* (2006).[17] This allows us to compare the flexibility of UK consumer prices at a product group level rather than just in the aggregate data (Table 4). While these results are less refined than those in Table 3, they suggest that UK consumer prices are more flexible than euro-area prices for all product groups other than energy goods, which is by far the most flexible group of prices for both economies.

Table 4: Percentage of UK and euro-area consumer prices changing each month

	UK	Euro area
Processed Food	20	14
Non-Processed Food	29	28
Energy Goods	64	78
Non-Energy Industrial Goods	18	9
Services	9	6

3.1.4 Price flexibility by year and calendar month

The frequency of price changes is not constant over time. The share of prices changing each month varies over our sample – Chart 4 plots the frequency of upward and downward changes against headline CPI inflation, and the results are summarised in Table 5. 2005 is the year when prices changed most frequently, with an average of 22% of prices changing each month. Prices changed least often in 2002, where an average of 16% of consumer prices changed each month. The patterns in the data are broadly similar if we exclude sales from the analysis (Chart 5).

[16] As discussed in Section 2, these differences result primarily from our sample only containing locally collected data.
[17] It is worth noting that Dhyne *et al* (2006) is based on a limited sample of 50 common products across different countries.

Chart 4: Percentage of UK consumer prices that change each month

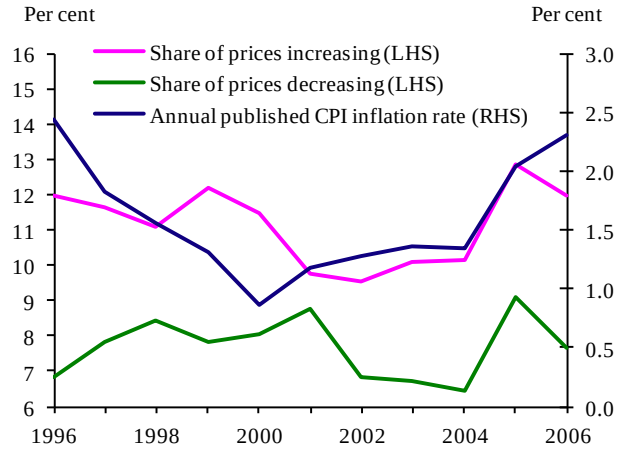

Chart 5: Percentage of UK consumer prices that change each month excluding sales

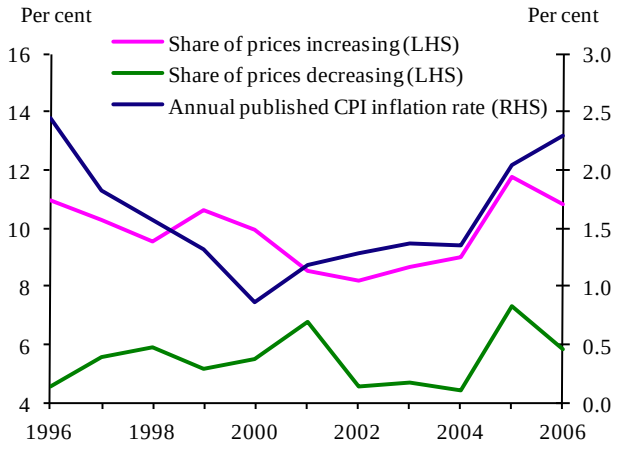

There appears to be some relationship between the share of prices increasing and the aggregate CPI inflation rate. The correlation coefficient between the two over our sample period is in the region of 0.6, rising to 0.7 if we exclude sales. The latter is statistically significant at the 5% level, although correlation coefficient is not quite significant if we use the total share of prices changing each year. There is little relationship between the average share of prices falling each month and the overall inflation rate.

Table 5: Percentage of UK consumer prices that change each month by year

	All changes	Increases	Decreases	All changes ex-sales
1996	18.8	12.0	6.8	15.4
1997	19.4	11.6	7.8	15.8
1998	19.5	11.1	8.4	15.4
1999	20.0	12.2	7.8	15.7
2000	19.5	11.5	8.0	15.4
2001	18.6	9.8	8.8	15.3
2002	16.4	9.5	6.8	12.7
2003	16.8	10.1	6.7	13.3
2004	16.6	10.2	6.4	13.4
2005	21.9	12.9	9.1	19.0
2006	19.6	12.0	7.6	16.6

It is well known that there is seasonal variation in prices, and that is clear in the microdata (Chart 6). When all changes are included, prices change most frequently in January. However, this reflects the impact of post-Christmas sales: when sales are excluded from the data, more prices change in April than in any other month (Chart 7). This could reflect the timing of duty changes and/or firms changing price to coincide with the start of the financial year. April is also the month when there are fewest price falls and the most price rises. Excluding sales, an average of 14% of prices rise in April,

but no more than 11% of prices increase in any other month of the year. Prices are least likely to change in November and December, and this is true whether we include sales or not.

Chart 6: Percentage of UK consumer prices that change each month

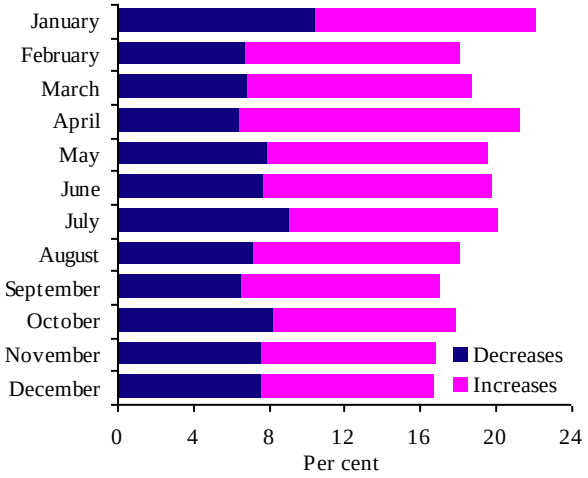

Chart 7: Percentage of UK consumer prices that change each month excluding sales

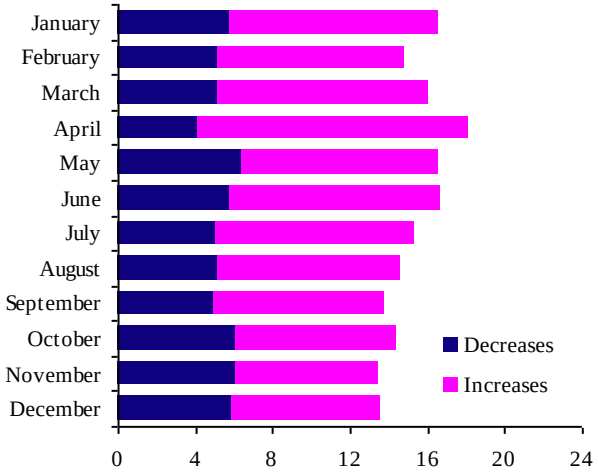

3.2 Hazard functions

So far, we have presented average results – the proportion of prices changing across our whole data sample. These averages can be interpreted as unconditional probabilities of price changes. However, conditional probabilities are also interesting. In particular, we are interested in the probability of a price change occurring if we know how long it has been since the last change in prices. These conditional probabilities are captured by hazard functions.

A hazard function, *h(t)*, measures the probability that a price will change in period *t* given that it has not changed in the last *t-1* periods (equation (**1**)). This is calculated as the share of firms adjusting their price in period *t*, *f(t)*, over the share of firms who have not changed their price in the last *t-1* periods, *s(t)*. The time profile for the share of firms that have not changed their price is known as the survivor function.

$$h(t) = \frac{f(t)}{s(t)} \qquad (1)$$

We only use items that have at least one price change in our estimation of the hazard functions. This is because we need to be certain how many months have elapsed since the previous price change. We only use each item once in the hazard function estimation – ie, we use the time between the first price change and the second (if the latter is present). This gives a more representative picture of the behaviour of the price of the average item compared to using all price spells for all items because it avoids the hazard functions being dominated by a small number of items whose prices change very frequently.

3.2.1 Aggregate hazard functions

Chart 8 plots estimated hazard functions for UK consumer prices on a weighted and unweighted basis.[18] Weighting only has a limited impact on the shape of the hazard functions. Both hazards decline over time, although there are visible spikes at the one-year horizon and to a lesser extent at two years. The hazard functions reported in the rest of the paper are weighted unless otherwise stated. Chart 9 shows the survivor function which represents the probability a price will not have changed given the time that has elapsed since the previous change. The probability of a price changing within three months of the previous price change is approximately 50%. The probability of a price change within twelve months is around 80%, and the chance of the price changing within three years of the previous change is 95%.

Chart 8: UK consumer price hazard function

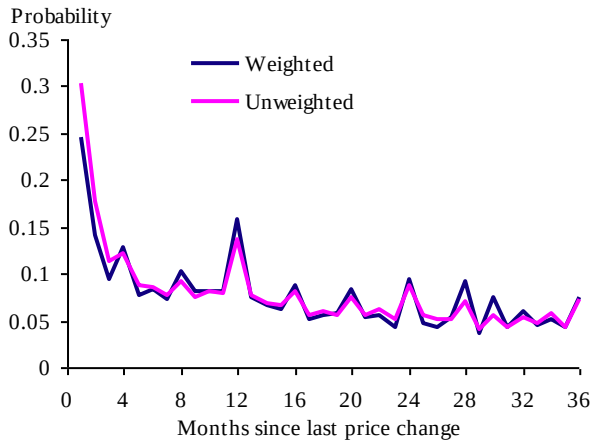

Chart 9: UK consumer price survivor function

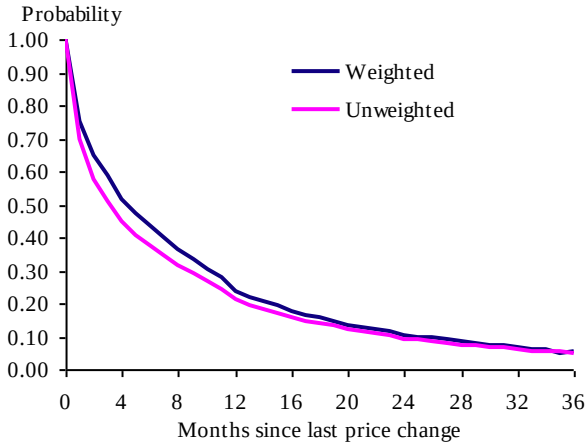

The hazard functions for UK consumer prices have similar properties to those estimated for other euro-area countries and the United States. Alvarez *et al* (2005) summarise these hazard functions as not being zero in any period, being downward sloping and having noticeable spikes at one and twelve months. Our estimated hazard functions clearly fit these stylised facts.

3.2.2 Hazard functions by product group

Perhaps more interesting than the aggregate hazard function is the clear distinction between the broad 'goods' and 'services' categories. Chart 10 plots weighted hazard functions for goods and services prices separately. The hazard function for goods prices has a spike at one month, showing that the probability of a price change is the highest in the month after the previous price change. This could be reflecting temporary price promotions. Beyond one month the goods price hazard function is downward sloping, broadly matching previous work on UK producer prices (Bunn and Ellis (2010)), although there are no clear spikes in the CPI goods hazard function beyond one month. This suggests that the conditional probability of changing prices declines over time.

In marked contrast, the hazard function for services prices is broadly flat with a large spike at twelve months and a smaller spike at 24. This profile is suggestive of annual price reviews and is more

[18] The weighted hazard functions are based on weights of items in the pooled sample.

consistent with strict time-dependent pricing models. These hazard functions suggest that retail goods and retail services firms are very different in terms of their price-setting behaviour. The heterogeneity that is evident in the data suggests that the standard 'representative agent' model will not accurately capture developments in the economy.

Chart 10: UK consumer price hazard functions for goods and services

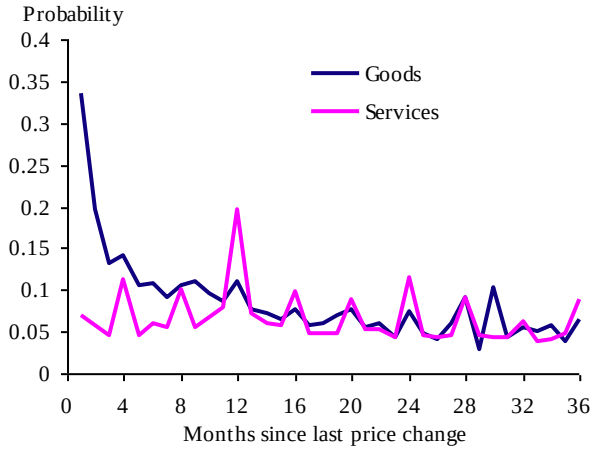

Chart 11 shows that there are differences in the hazard functions for the different goods components of UK consumer prices. The hazard function for energy goods has a very large spike at one month, and a steep downward slope. This is consistent with the high proportion of energy goods changing price each month.[19] The hazard functions for food and non-alcoholic beverages and non-energy industrial goods look quite similar. They have a smaller spike at one month, they decline to around four months and then the slope flattens so that it is only gently decreasing beyond four months.

Chart 11: UK consumer price hazard functions for goods components

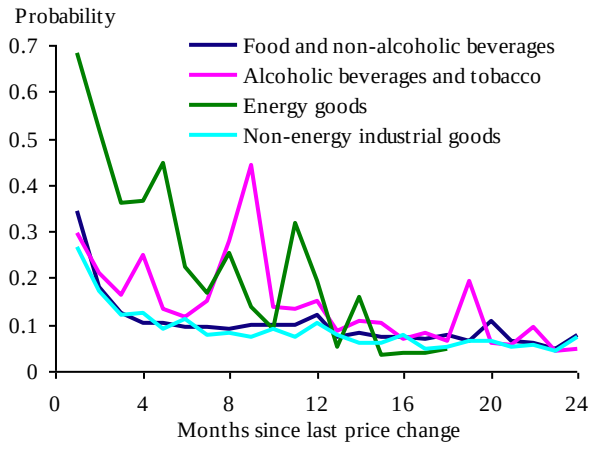

Chart 12: UK consumer price hazard functions for services components

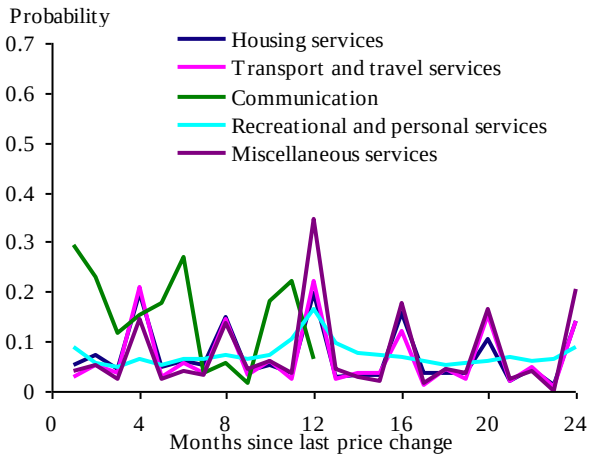

The hazard functions for the services components of CPI are shown in Chart 12. Housing services, transport and travel services and miscellaneous services all have very similar profiles. They are broadly flat with spikes every four months and a larger annual spike, the annual spike is largest for

[19] The hazard function for energy goods does not go beyond 18 months due to a lack of observations beyond this horizon.

miscellaneous services. Recreational and personal services has a slightly different profile, the probability of a price change is similar in each period except for a slightly smaller-than-average annual spike. The communication services hazard function looks very different to the other components, but it is based on a much smaller sample than the other hazard functions.

3.3 Magnitude of price changes

3.3.1 Distribution of price changes by magnitude

The distribution of the size of consumer price changes (summarised in Table 6 and Chart 13) is wide, but there are also a significant number of price changes that are relatively small and close to zero. To illustrate the width of the distribution, around 95% of all price changes lie between -50% and 50%. 25% of all price cuts are bigger than 20% and 25% of all price increases are above 17%. There are spikes in the distribution of price changes at round percentage numbers such as 20% or 25%. These changes are large in relation to the overall aggregate inflation rate in our sample period. But there is some concentration of the distribution in smaller price changes. The interval between -10% and 10% covers around 60% of the distribution. The median price change is a 2% rise and the most popular 1 percentage point interval is a price increase between 1% and 2% (this covers around 8% of all price changes). There tend to be fewer very small price changes between -1% and 1% than there are more moderately small price changes between 1% and 2% or -1% and -2%.

Table 6: Distribution of the size of UK consumer price changes (percentage changes)

	All changes	Increases	Decreases	All changes ex-sales
5th percentile	-31.0	1.1	-44.0	-21.4
25th percentile	-6.3	2.7	-20.3	-2.6
Median	1.7	6.3	-10.0	2.1
75th percentile	8.3	17.5	-3.2	6.7
95th percentile	43.0	61.9	-0.9	30.8

Chart 13: Distribution of the size of UK consumer price changes[20]

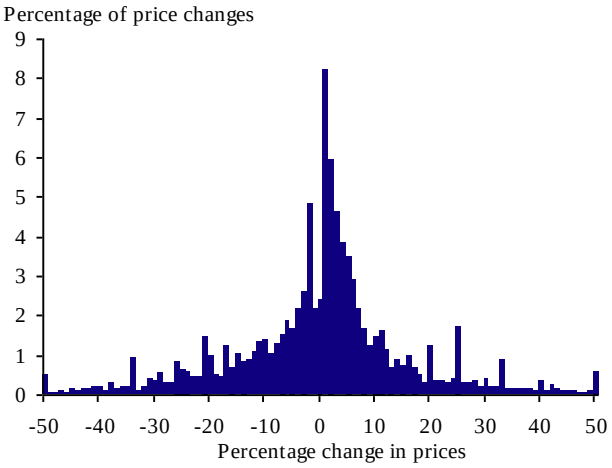

Chart 14: Distribution of the size of UK consumer price changes

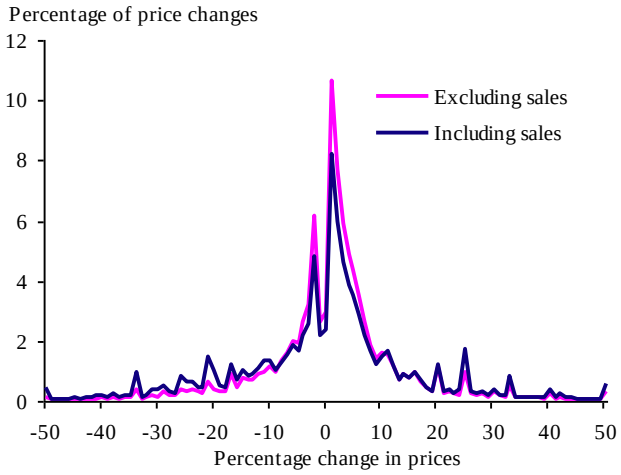

Chart 14 shows the extent to which price changes associated with sales shape the distribution of the size of price changes. As noted above, price changes associated with sales in our sample are more likely to be price cuts rather than prices recovering from sales. Excluding sales, there are fewer large price cuts (defined as price cuts that are larger than 10%), particularly at round percentages, which may be commonly used by firms when discounting their prices. A higher proportion of price changes in the distribution excluding sales are smaller price changes, particularly in the -1% to -2% and 1% to 2% intervals. Partly due to the importance of sales, the median size of a price cut of around 15% is larger than the corresponding average size of a price increase of 7%. However, if we exclude sales, the median price increase and the median decrease are both around 5%.

The distribution of the size of changes in UK consumer prices appears to be broadly similar to that for euro-area consumer prices. Dhyne *et al* (2006) show that the distribution of price changes in the euro area is also wide with some concentration around zero. At approximately 60%, the proportion of price changes in the euro area that lie between -10% and 10% is similar to the corresponding figure for the United Kingdom.[21] Dhyne *et al* (2006) also report that the average price increase in the euro area is slightly higher than the average for the United Kingdom (8% compared to 7%) and a lower proportion of price increases in the euro area are very small (below 2.5%). But the average size of a price fall in the euro area is only 10%, smaller than the 15% we observe in the United Kingdom. This implies that sales are perhaps more prevalent in the United Kingdom than in the euro area, although these differences are relatively modest, given the width of the distributions.

3.3.2 Distribution of price changes by product group

The earlier analysis highlighted some clear differences in how often goods and services prices change. And, looking at the magnitude of those changes, there are clear differences in the distributions as well. In particular, there are more small increases in services prices than there are in goods prices: 55% of all services price changes are between 0 and 10%, compared with just over 30% of goods price changes (Chart 15). There are also fewer falls in services prices relative to

[20] The bars in this chart represent 1 percentage point intervals, for example the 0 to 1 interval contains all price changes that are greater than or equal to zero and less than 1%. This definition in used for all the intervals referred to in the paper.
[21] These figures for the euro area refer to the period between 1996 and 2001.

goods. This may in part reflect the fact that services price inflation has been higher over our sample period than goods price inflation. It can also help to explain why the median price change for services in our sample is a 4% rise, compared with a median increase of only 1% for goods.

Chart 15: Distribution of the size of UK consumer price changes

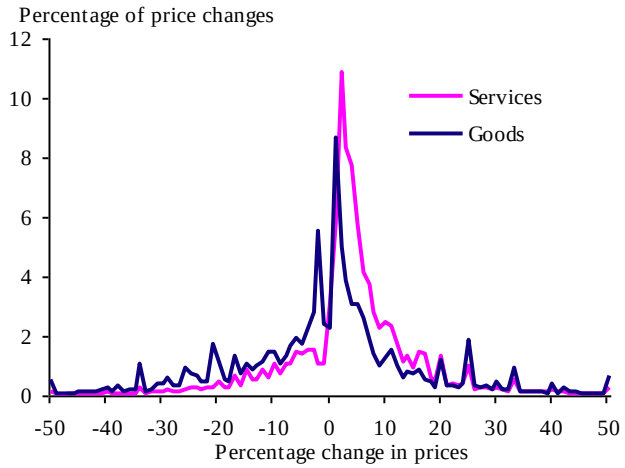

There are also clear differences between the distributions of price changes at the more disaggregated component level. Among the goods components, a large proportion of changes in the price of energy goods and alcoholic beverages and tobacco are relatively small (Chart 16). These small changes in the price of alcoholic beverages and tobacco are more likely to be positive changes, and they are particularly concentrated between 1% and 2% – in part, they may be associated with changes in duty on those products. Among energy goods, essentially petrol, few price changes are between -1% and 1%, but a large proportion lie between -1% and -2% and 1% and 2%. This is probably because petrol prices often change in units of one pence per litre, which over our sample period is usually a price change of between 1% and 2%. In contrast, the distribution of price changes in food and non-alcoholic beverage products and non-energy industrial goods looks very different to those distributions that we have highlighted above: these other distributions are very wide with little or no peak around zero. So the shape of the aggregate distribution of goods price changes, which has some large changes and some concentration in small price changes, results from the aggregation of these two different types of distribution at the component level: the wide distributions among food and non-alcoholic drink products and non-energy industrial goods, and the narrower distributions of energy goods and alcoholic beverage and tobacco products. Again this illustrates the heterogeneity in price-setting behaviour that is evident in the data.

Chart 16: Distribution of the size of UK consumer price changes by goods component

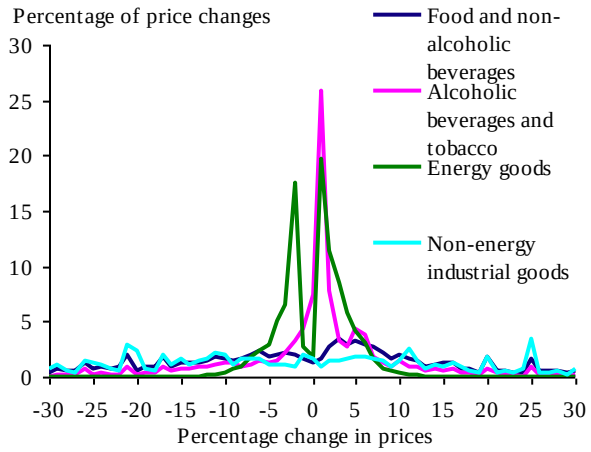

Chart 17: Distribution of the size of UK consumer price changes by services component

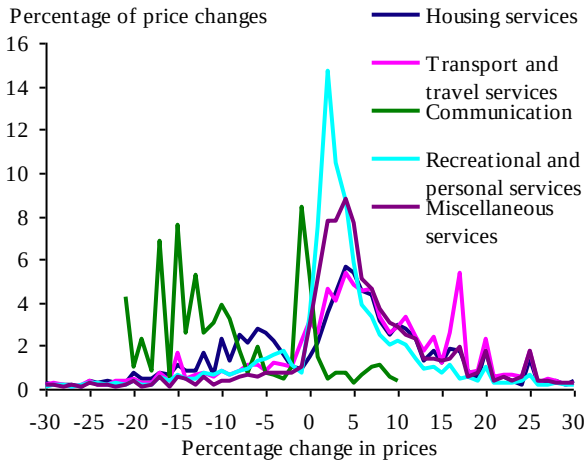

Excluding communications, where our sample size is small, there is less heterogeneity in the distribution of price changes between the services components of CPI (Chart 17). In fact, all other services components have some concentration of price changes between 0% and 10% and a relatively small proportion of price changes are price falls. Recreational and personal services has the largest proportion of small price changes, and the spike for price changes of between 2% and 3% for this component accounts for (most of) the corresponding spike in the overall services price distribution. Tables A7 and A8 in the appendix give more details of the distribution of the magnitude of price changes by CPI component and COICOP division respectively.

Overall, prices tend to rise in our sample, consistent with the positive inflation rate during the sample period, and we observe that more of the distribution of the size of price changes is above zero than below, particularly in the case of services. But all of the distributions of price changes at the CPI component level have a non-trivial share of prices changes that are price falls – and as such, there is little evidence to support the presence of downward nominal rigidities in product markets in the United Kingdom.

3.3.3 Distribution of price change by year and by calendar month

The distribution of price changes is relatively similar across different years in our sample, although there are still some differences from one year to the next. Chart 18 plots the distribution from three different individual years, which are chosen to illustrate the maximum extent of this variation: Table A9 in the appendix summarises the distribution for all of the individual years in our sample. In each year, the distribution of price changes is always relatively wide, with some concentration of price changes around zero. For price increases and decreases greater than 10% in size, the distributions look very similar in all years, as the biggest differences from one year to the next occur for price changes between -10% and +10%. 1% to 2% is generally the most common interval for price changes, but the height of this spike in the distribution varies. For instance, this spike was particularly high in 2004, but much lower in 1998. In 2005, there were more large price increases than average, but this was offset to some extent by there being more large price decreases as well. In 1998 there were a particularly high number of small price cuts between -1% and -2%.

Chart 18: Distribution of size of UK consumer price changes by year

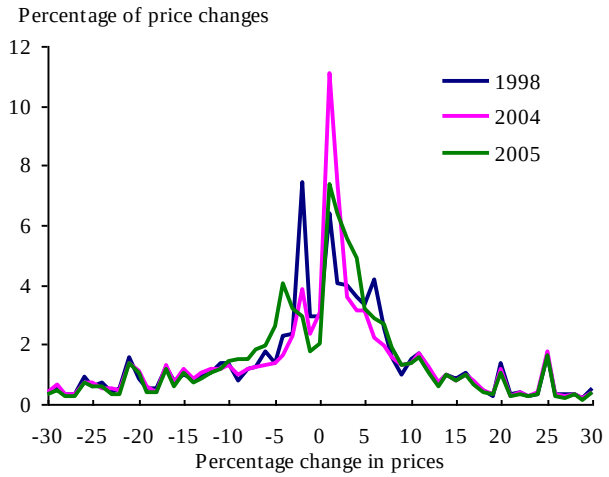

Chart 19: Median UK consumer price change by year

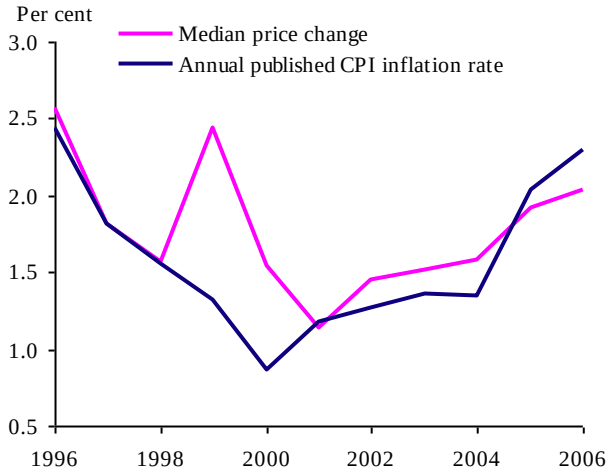

As well as variation in the distribution of price changes across different years, the average size of a price change also changes over time. Chart 19 shows that there is a correlation between the median price change and the overall CPI inflation rate.[22] Combined with the relationship between the share of prices changing and aggregate inflation that we identified earlier (see Section 3.1.4), this implies that periods of higher CPI inflation are characterised by both more firms changing their price, and by those firms that do change their price changing it by more.

Chart 20: Distribution of size of UK consumer price changes by calendar month

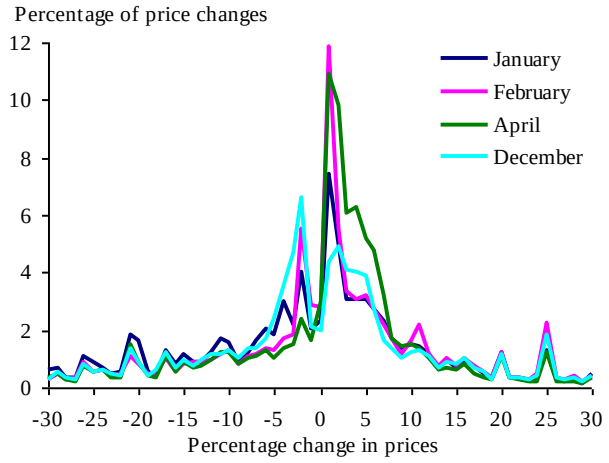

There is also some variability in the distribution of price changes in different calendar months of the year. Chart 20 shows some of the biggest differences among different months.[23] Chart 20 shows that there are more large price reductions than average in January, consistent with our earlier finding that January is the most popular month for sales, and a smaller proportion of price changes are modest price increases. These sale price reductions in January are often unwound in February, where there are more price increases of 10%, and indeed 25%, than normal. In April, more of the

[22] The correlation coefficient between the two series is approximately 0.7, and this is statistically significant.
[23] Table A10 in the appendix summarises the distributions for all months.

distribution of price changes is covered by increases between 0% and 10% relative to other months, and there are also fewer price falls. This could be partly related to duty changes which would typically come through in the April data, and this perhaps explains why April is the month in which the average monthly change in CPI tends to be largest. December is a month in which there are relatively few price changes. Of the price changes that there are, a lower proportion than average are small price increases between 0 and 3%, and a higher than average share are small price cuts, particularly in the -1% to -2% interval.

3.3.4 Price reversals

Because our sample has detailed figures in pounds and pence, we can identify precise price reversals in the data set – instances where one price change is the opposite size and sign to the previous change in the same price. These reversals can be used as another proxy for sales, as described in Ellis (2009). Twenty one per cent of price changes in our sample are direct reversals of the previous price change. Around 40% of these are price decreases where the price was previously increased, and 60% are price rises where the price had been previously cut – so price reversals do not just reflect sale discounts, but perhaps also retailers testing the sensitivity of demand to higher prices. Prices recovering from sales account for some of these price reversals, but even excluding sale prices, as identified by the ONS in their sampling, 17% of all price changes are direct reversals of the previous price change. The split between price rises followed by cuts, and decreases followed by increases, is roughly 50-50 once sales are stripped out. Price reversals are slightly more common among goods than they are for services, and they are particularly a feature of non-energy industrial goods, where 27% of price changes are direct reversals of the previous change.

The finding that a significant number of price changes are direct reversals of the previous price change, even once we exclude those that are identified as sales, suggests that there may be some temporary price promotions in our data that are not identified as sales. Interestingly, a number of firms appear to temporarily increase their price before reducing it, as well as cutting their price before increasing it again. One possible explanation for this pattern could be that firms repeatedly make the same temporary reductions in the prices of their products, which would mean that every time the promotion starts it is simply reversing the price increase made when the promotion last ended.

3.4 Correlation between the frequency and magnitude of price changes

From our previous results, we know that the distribution of price changes is wide, and that, overall, prices change fairly often. One hypothesis that the data allows us to explore is the idea that there may be a link between the frequency and magnitude of price changes – ie the longer it has been since the last price change, the more price-setters change their prices by when they next move them. If some constraint exists which encourages firms to set prices at infrequent intervals there is more scope for the actual price to differ from the optimal price as more time passes since the price last changed.

We can examine this by plotting the duration of a price since its previous change against the size of that price change. And when we do this across the distribution, there appears to be a clear

relationship between the two (Chart 21).[24] The median price change when that change occurs within three months of the previous price change is an increase of just over 1%. However, if a year has passed since the last price change the median price change is around 3%, and if it is two years since the last price change this rises to around 5%. Chart 22 splits out price changes into price increases and price decreases. Price increases tend to be larger for items that change price less frequently and price cuts are also larger when a longer time period has elapsed since the previous price change.

Chart 21: Size of UK consumer price changes and duration since previous change

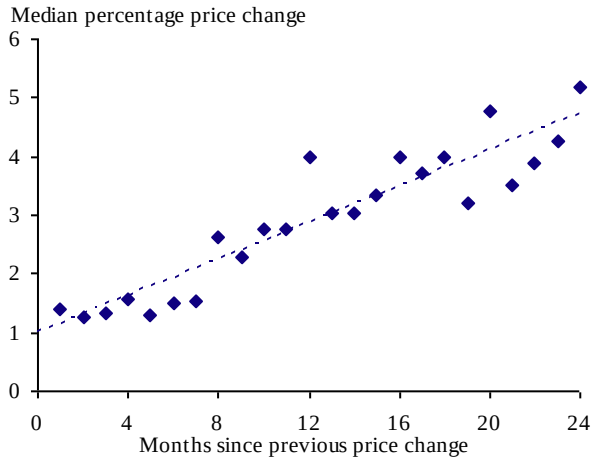

Chart 22: Size of absolute UK consumer price changes and duration since previous change

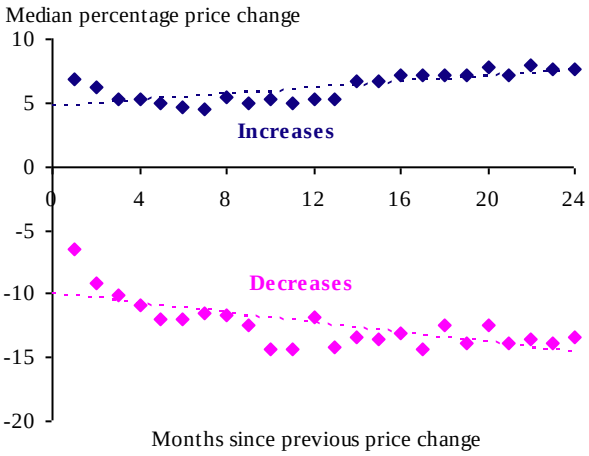

When we examine this relationship individually for goods and services, it is clear that the relationship is strongest for services prices (Chart 23). In fact, the correlation coefficient between the median price change and the number of months since the previous change for services is around 0.95. Initially, there appears to be little correlation between the frequency and magnitude of price changes for goods (Chart 24). But this relationship is distorted by sales. Once we exclude sales, a stronger relationship emerges between the size and the frequency of goods price changes (Chart 25). The correlation coefficient between the two is around 0.85, only a little below the corresponding figure for services. Sales are not as important for services, so excluding them makes little difference for that sector. In aggregate, the overall relationship between the frequency and magnitude of price changes is a little closer than is shown in Chart 21 if we exclude sale prices.

[24] This analysis is pooled across the sample and observations are weighted by their weight within the pooled sample.

Chart 23: Size of UK consumer price changes and duration since previous change for services

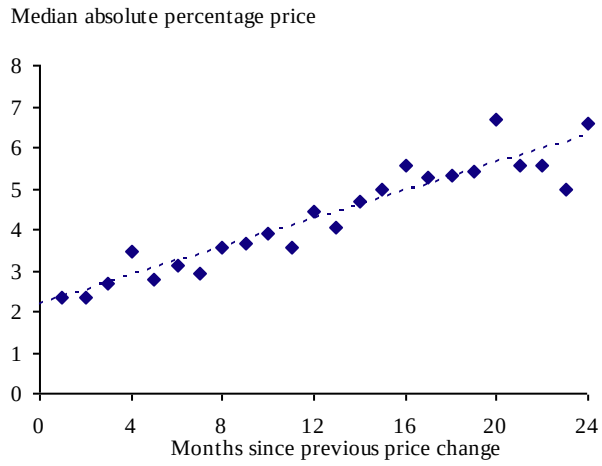

Chart 24: Size of UK consumer price changes and duration since previous change for goods

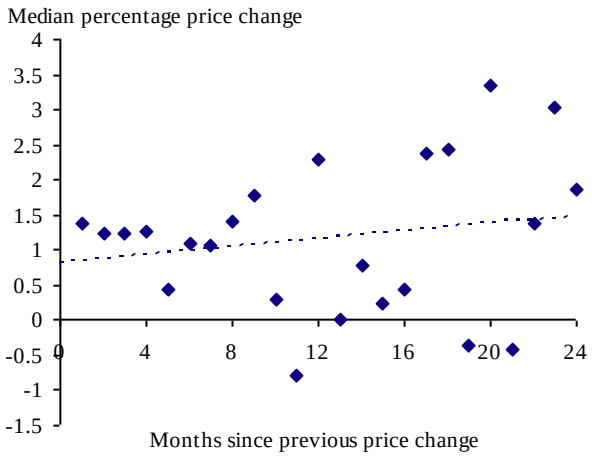

Chart 25: Size of UK consumer price changes and duration since previous change for goods excluding sales

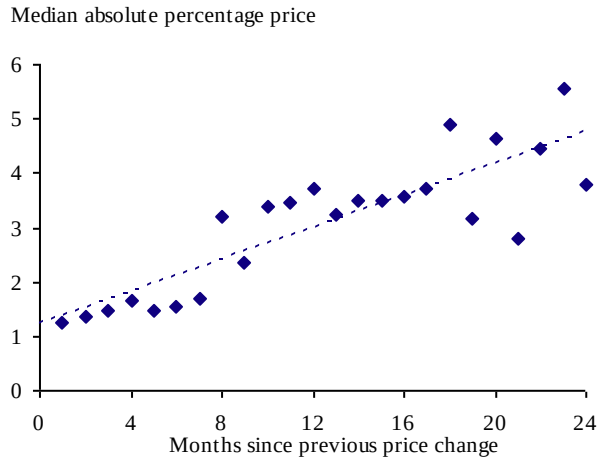

3.5 Comparison with producer price results

In a companion paper, Bunn and Ellis (2010) examine pricing behaviour using the microdata underlying the Producer Price Index in the United Kingdom. They follow a similar methodology to this paper, but there are a number of differences between the two studies. The CPI study examines how prices behave further down the supply chain than the PPI work, and it covers the behaviour of services prices as well as goods prices. Another obvious difference is the sample – while our CPI sample covers eleven years, the PPI study covered just five years. That said, we can compare results from the two studies over a comparable data sample, 2003-06. Unfortunately, we only have locally collected data that are used to construct the CPI, while the PPI study covers the vast majority of the underlying microdata.

A first observation is that consumer prices appear to be less flexible than producer prices. Over a comparable time period between 2003 and 2006, 25% of producer prices change each month (compared to 26% in the full PPI sample which also includes 2007), whereas only 19% of consumer

prices change per month, or 15% if sales are excluded. However, this is partly because the CPI includes services prices that are not included in the PPI; and our results indicate that CPI services prices behave differently to CPI goods prices. If we restrict the CPI sample to goods only and compare both statistics over the same 2003 to 2006 time period, we find that 25% of CPI goods prices change each month, which is the same as the share of PPI prices changing to the nearest percentage point (Table 7).[25] Another common result across the two sets of prices is the marked heterogeneity in the frequency of price-setting.

Table 7: Percentage of prices changing each month 2003 to 2006

	All changes	Increases	Decreases
All CPI	18.7	11.3	7.5
CPI goods	24.9	14.1	10.8
PPI goods	25.4	15.1	10.3

When we compare the hazard functions for CPI and PPI prices, we again find a number of similarities between the two (Chart 26). Both curves decline with duration, indicating that the conditional probabilities of price changes fall as the time since the previous price change increases. The CPI hazard function exhibits the spikes at one and twelve months that are evident in the PPI hazard function, but in CPI it is services prices that generate the annual spike in the hazard function. Once we plot the CPI goods hazard function this still slopes downwards, but it does not have this obvious spike.

Chart 26: CPI and PPI hazard functions

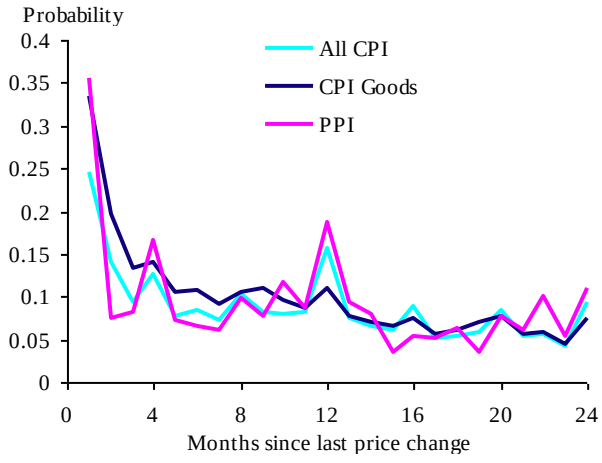

Chart 27: Distribution of the size of CPI goods and PPI price changes 2003-06

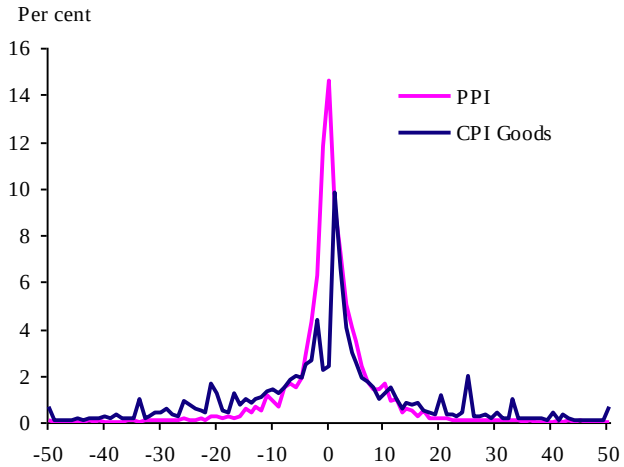

Chart 27 compares the distribution of the size of CPI goods and PPI price changes. Both have a wide distribution, but both also have a significant number of price changes that are relatively small and close to zero. The distribution of CPI goods price changes is a little wider than for producer prices, while the distribution of the size of producer price changes has a higher proportion of price changes

[25] Limiting the sample periods to make them directly comparable does not make substantial differences to the results discussed earlier in the paper.

that are relatively small. The CPI goods distribution has some spikes at particular round numbers, but these spikes are not a feature of the PPI distribution. And in both cases, there is a positive relationship between the size of price increases and the number of months since the previous price change.

In our work on producer prices we highlighted the finding that producer prices are much less persistent at the disaggregated level than aggregate inflation data imply. One difference between the CPI and PPI data over our sample period is that the aggregate CPI inflation rate is much less persistent than the corresponding PPI measure of inflation. We have therefore not looked in detail at persistence at the item level in the CPI data as there is no persistence to explain in the aggregate data. However, a preliminary investigation suggests that there is no evidence of persistence at the item level in CPI, [26] which is consistent with our findings on producer prices.

Throughout this paper, we have also highlighted the differences between the behaviour of individual goods and services prices within CPI. CPI goods are a more relevant set of prices to compare to the producer prices than CPI prices as a whole. All told, the results on the properties of CPI goods prices bear remarkable similarity to the results from producer price data. This similarity suggests that there are few pricing frictions that exist between the production and retail sectors in the United Kingdom.

4 Implications for pricing theory

The primary reason for undertaking the work described in this paper is to improve our understanding of how prices are set in the economy, and to learn more about the nature of the nominal rigidities that exist. Monetary policy will have no effect on real output if prices are fully flexible, and so nominal rigidities are an important element of many economic models because they provide a mechanism through which changes in monetary policy can affect real output. Learning more about these rigidities can help to improve our understanding of the effectiveness of monetary policy and the monetary transmission mechanism.

Nominal rigidities take a number of different forms in monetary policy models. Depending on the assumptions made about the structure of these rigidities, different models can have varying implications for policy. These pricing models can be categorised under two main headings: time-dependent and state-dependent models. In a time-dependent model the probability of a price change depends only on the time since the previous change. The model developed by Calvo (1983), in which homogenous firms have a fixed probability of changing their price in each period, is one of the most popular specifications. Alternatives include staggered contracts in which prices are fixed for the duration of the contract (Taylor (1980)). In a state-dependent model the decision to change prices is dependent on the current state of the economy and the market facing the firm. These models often work by incorporating some cost to adjusting prices; examples include quadratic adjustment cost models in which firms receive disutility from making large price changes (Rotemberg (1982)) and menu cost models (Mankiw (1985)). In this type of model, prices change intermittently as firms react to shocks and developments in the economy.

[26] The coefficient on the lagged dependent variable in a simple AR(1) regression of CPI month on month inflation rates is not statistically significant.

It is clear from our results that a degree of nominal rigidity is present in the economy. Prices do not adjust continually – we find that only a fifth of prices change in any given month, although some firms may review their price each month but decide that they do not want to change it. The empirical evidence is not consistent with any one pricing theory which can explain the form of these nominal rigidities. There are pieces of evidence that can be used to argue both for and against different models, and different theories appear to better fit the behaviour of some groups of prices than others. For example, the strict Calvo price-setting model, which implies a constant probability of price changes in each period, is not consistent with the variation in the share of prices changing that we see in different years and in different calendar months of the year. It is also not consistent with the hazard functions that we observe, which contain a downward slope (for goods) and annual spikes (for services). However, the hazard function for services prices is relatively flat once the spikes are excluded, and could be broadly consistent with other time-dependent type models such as staggered contracts which only allow prices to change annually.

If 'menu cost' state-dependent models were able to fully explain the nominal rigidities we see in the data, we might expect to see relatively few small price changes. But we find that a significant proportion of price changes are small, which would not follow from fixed menu costs. However these costs of changing price could still be important for some firms. For example, there are very few small changes in the price of food and non-alcoholic beverage products or non-energy industrial goods. At the same time, the fact that we observe some very large price changes argues against firms being unwilling to make large price changes for fear of upsetting customers (Rotemberg (1982)), where the cost of changing price increases more than one-for-one with the size of the price change. A single, aggregate adjustment cost model cannot fully explain our results, although it could be that these adjustment costs are still important for some firms and we are merely observing the effects of such heterogeneity.

The finding that no one theory can explain how firms set their prices is consistent with the recent Bank of England price-setting survey (Greenslade and Parker (2008)). The survey found that some UK firms use mainly time-dependent pricing rules (44%), some use state-dependent pricing rules (15%), and the remainder use a combination of the two.

The heterogeneity that we find in pricing behaviour across different industries and product groups is one of the most interesting results from our study, and is consistent with very similar observations from the other microdata studies. Given this heterogeneity, it is likely that particular theories can better explain pricing behaviour in some sectors than in others and therefore it may be difficult to find any one theory that can explain pricing behaviour at the economy-wide level. For example, almost 65% of energy product prices change each month and therefore it could be argued that nominal rigidities are not particularly important in this sector. But less than 10% of services prices change each month, and therefore a different model may be needed to explain the nominal rigidities in this sector. The differing properties of goods and services prices suggest that there may not be one pricing theory that can easily explain the behaviour of both types of prices.

5 Conclusions

This paper has analysed the behaviour of individual prices that are used in the construction of the UK Consumer Price Index. There is no previous similar work using UK consumer prices, so our paper adds to the micro-pricing literature by being the first to make use of this data. It complements existing work on UK producer prices and studies on consumer prices in other countries.

Our study has uncovered several interesting results. First, on average 19% of prices change each month, although this falls to 15% if sales are excluded. A small number of items account for many price changes, which implies that price changes occur less frequently when measured by the average for individual products. Second, the probability of price changes is not constant over time; there is variation between the different years in our sample and between different months of the year. The probability of a price change occurring also varies depending on the time elapsed since the previous price change. Third, the probability of prices changing varies significantly between the different components of CPI. In particular, goods prices change more frequently than services prices. Fourth, the distribution of price changes is wide, although a significant number of changes are relatively small and close to zero. Fifth, prices that change more frequently tend to do so by less.

These results are consistent with the conclusions of many other micro-pricing studies in that they show that the probability of prices changing is not the same in all periods and they emphasise the importance of heterogeneity between the behaviour of prices of different groups of items. The results are not consistent with any one price-setting theory and the heterogeneity we observe suggests that different pricing models may be able to better explain price-setting behaviour in different sectors. This heterogeneity is often ignored in typical 'representative agent' models, and micro-studies also tend to find that prices change more frequently than is implied by macro models. If we want to use micro-founded macro models that match the stylised facts that we observe in these types of micro-studies, the challenge is to develop a new theory of price-setting behaviour that is consistent with these facts while also fitting the properties of aggregate data.

Appendix: Detailed tables of results

Table A1: Number of observations by year

	Number of observations
1996	1,072,975
1997	1,131,659
1998	1,040,234
1999	982,722
2000	983,323
2001	996,054
2002	988,986
2003	956,047
2004	981,015
2005	944,488
2006	989,333
Total	11,066,836

Table A2: Sample coverage by CPI component

	Number of observations	Number of items	Average CPI weight (%)	Weight in microdata (%)	Percentage of CPI covered by microdata
Food and Non-Alcoholic Beverages	2,699,578	124841	12.4	17.6	82.9
Alcoholic Beverages and Tobacco	454,052	18288	5.6	7.3	78.2
Energy Goods	272,706	11025	6.9	6.0	52.1
Non-Energy Industrial Goods	4,902,859	290956	33.3	33.4	58.1
Housing Services	348,356	16653	7.7	5.8	44.3
Transport and Travel Services	238,244	10593	6.0	4.6	45.4
Communication	22,605	1496	2.4	0.2	4.1
Recreational and Personal Services	1,942,387	87519	20.1	21.5	61.2
Miscellaneous Services	186,049	9365	5.6	3.6	38.5
All items	11,066,836	570,736	100	100	57.3

Table A3: Sample coverage by COICOP division

	Number of observations	Number of items	Average CPI weight (%)	Weight in microdata (%)	Percentage of CPI covered by microdata
Food and Non-Alcoholic Beverages	2,699,578	124841	12.4	17.6	82.9
Alcoholic Beverages and Tobacco	454,052	18288	5.6	7.3	78.0
Clothing and Footwear	1,598,830	126667	6.6	7.6	66.3
Housing and Utilities	526,819	24566	11.8	6.7	35.2
Furniture and Home Maintenance	1,511,308	74797	7.9	11.3	86.1
Health	187,038	7273	1.7	1.9	64.4
Transport	641,444	27168	15.4	11.4	43.0
Communications	22,605	1496	2.4	0.2	4.2
Recreation and Culture	987,796	53606	14.6	9.2	34.6
Education	0	0	1.4	0	0
Restaurants and Hotels	1,593,636	72855	12.8	18.3	80.7
Miscellaneous Goods and Services	843,730	39179	7.6	8.5	63.0
All items	11,066,836	570,736	100	100	57.3

Table A4: Percentage of consumer prices that change each month by COICOP division

	All changes	Increases	Decreases	All changes ex-sales
Food and Non-Alcoholic Beverages	21.8	11.9	10.0	17.8
Alcoholic Beverages and Tobacco	25.7	18.3	7.5	21.6
Clothing and Footwear	23.8	11.0	12.8	11.6
Housing and Utilities	11.1	6.8	4.3	10.5
Furniture and Home Maintenance	19.9	10.3	9.6	9.7
Health	8.4	6.3	2.2	7.6
Transport	36.7	22.1	14.6	36.6
Communications	20.0	7.2	12.8	14.4
Recreation and Culture	14.5	7.2	7.3	11.0
Restaurants and Hotels	9.1	7.3	1.7	8.9
Miscellaneous Goods and Services	11.0	7.2	3.9	8.9

Table A5: Number of months between price changes by CPI component

	Inverted frequency of change	Median months per change per item	Mean months per change per item	% of items that never change price (weighted)	% of items that never change price (unweighted)
Food and Non-Alcoholic Beverages	4.6	5.7	8.2	6.1	17.3
Alcoholic Beverages and Tobacco	3.9	3.9	5.3	3.2	13.2
Energy Goods	1.6	1.5	1.9	0.4	7.5
Non-Energy Industrial Goods	5.4	6.8	10.3	12.1	23.4
Housing Services	11.9	11.3	15.1	18.3	36.2
Transport and Travel Services	12.2	11.5	15.8	18.0	31.9
Communication	5.0	5.0	6.6	10.5	28.7
Recreational and Personal Services	11.3	11.3	15.2	15.5	31.5
Miscellaneous Services	11.4	11.3	14.8	14.3	30.0
All goods	4.1	5.0	8.3	8.4	20.9
All services	11.4	11.3	15.2	16.1	32.0
All items	3.9	7.2	10.6	11.1	23.3

Table A6: Number of months between price changes by COICOP division

	Inverted frequency of change	Median months per change per item	Mean months per change per item	% of items that never change price (weighted)	% of items that never change price (unweighted)
Food and Non-Alcoholic Beverages	4.6	5.7	8.2	6.1	17.3
Alcoholic Beverages and Tobacco	3.9	3.9	5.3	3.2	13.2
Clothing and Footwear	4.2	4.8	7.7	10.9	21.7
Housing and Utilities	9.0	10.7	13.9	15.7	32.0
Furniture and Home Maintenance	5.0	6.3	10.1	11.8	21.8
Health	11.9	11.5	18.6	15.4	31.2
Transport	2.7	2.3	8.0	9.0	22.2
Communications	5.0	5.0	6.6	10.5	28.7
Recreation and Culture	6.9	8.5	11.9	15.3	28.0
Restaurants and Hotels	11.0	11.0	14.7	15.0	31.1
Miscellaneous Goods and Services	9.1	11.0	13.8	13.4	28.5
Total	3.9	7.2	10.6	13.3	26.6

Table A7: Distribution of percentage changes in consumer prices by CPI component

	5th percentile	25th percentile	Median	75th percentile	95th percentile
Food and Non-Alcoholic Beverages	-33.4	-11.2	2.6	12.7	50.0
Alcoholic Beverages and Tobacco	-16.7	-0.9	1.5	4.7	20.0
Energy Goods	-5.7	-1.5	1.3	3.1	6.5
Non-Energy Industrial Goods	-42.2	-17.9	-0.1	17.6	73.7
Housing Services	-20.2	-4.8	5.0	11.8	38.3
Transport and Travel Services	-20.0	0.9	6.7	15.3	34.8
Communication	-39.6	-16.7	-8.3	11.1	60.0
Recreational and Personal Services	-16.0	1.3	3.4	7.3	25.2
Miscellaneous Services	-16.7	2.5	5.5	11.4	36.4
All goods	-33.3	-7.9	1.3	7.8	46.7
All services	-17.7	1.1	4.0	9.3	30.8

Table A8: Distribution of percentage changes in consumer prices by COICOP division

	5th percentile	25th percentile	Median	75th percentile	95th percentile
Food and Non-Alcoholic Beverages	-33.4	-11.2	2.6	12.7	50.0
Alcoholic Beverages and Tobacco	-16.7	-0.9	1.5	4.7	20.0
Clothing and Footwear	-50.0	-25.0	-8.3	25.0	100.0
Housing and Utilities	-23.1	-5.7	3.7	10.2	35.2
Furniture and Home Maintenance	-35.2	-13.9	0.8	16.1	59.3
Health	-33.5	-0.8	4.3	9.5	50.3
Transport	-6.4	-1.5	1.3	3.5	9.4
Communications	-39.6	-16.7	-8.3	11.1	60.0
Recreation and Culture	-38.8	-14.0	0.0	13.3	60.0
Restaurants and Hotels	-14.3	1.2	3.2	6.5	22.7
Miscellaneous Goods and Services	-33.4	-6.7	4.6	13.3	52.3

Table A9: Distribution of percentage changes in consumer prices by year

	5th percentile	25th percentile	Median	75th percentile	95th percentile
1996	-26.2	-5.1	2.6	7.6	33.5
1997	-28.6	-5.6	1.8	8.0	39.3
1998	-30.6	-6.1	1.6	8.3	43.9
1999	-30.1	-6.7	2.4	8.3	43.2
2000	-33.4	-7.1	1.5	8.3	49.7
2001	-33.1	-6.3	1.1	7.9	45.0
2002	-33.4	-7.9	1.5	10.0	50.0
2003	-31.8	-6.4	1.5	9.4	48.2
2004	-33.3	-6.7	1.6	8.4	43.2
2005	-30.0	-5.9	1.9	7.5	40.0
2006	-31.3	-6.3	2.0	8.3	44.0

Table A10: Distribution of percentage changes in consumer prices by calendar month

	5th percentile	25th percentile	Median	75th percentile	95th percentile
January	-39.2	-11.1	1.1	7.3	37.2
February	-30.0	-4.4	1.5	9.8	43.3
March	-30.0	-4.3	2.7	9.1	50.0
April	-27.5	-3.1	2.5	6.6	30.7
May	-25.0	-5.2	1.8	8.0	40.8
June	-28.6	-4.5	2.3	8.6	42.9
July	-37.5	-11.6	1.3	6.1	37.2
August	-33.3	-6.0	1.6	8.5	42.6
September	-28.6	-5.9	1.9	10.0	44.4
October	-30.0	-7.3	1.3	9.0	50.0
November	-28.8	-6.7	1.4	11.1	53.3
December	-28.7	-5.4	1.6	9.6	57.6

References

Alvarez, L and Hernando, I (2004), 'Price setting behaviour in Spain. Stylised facts using consumer price micro data', *ECB Working Paper No. 416.*

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

Aucremanne, L and Dhyne, E (2004), 'How frequently do prices change? Evidence based on the micro data underlying the Belgian CPI', *ECB Working Paper No. 331*.

Baudry, L, Le Bihan, H, Sevestre, P and Tarrieu, S (2004), 'Price rigidity in France - evidence from consumer price micro-data', *ECB Working Paper No. 384.*

Baumgartner, J, Glatzer, E, Rumler, F and Stiglbauer, A (2005), 'How frequently do consumer prices change in Austria? Evidence from micro CPI data', *ECB Working Paper No. 523*.

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.

Bunn, P and Ellis, C (2010), 'How do individual UK producer prices behave?', *Bank of England Working Paper No. 394.*

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. 12824.*

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*.

Dias, M, Dias, D and Neves, P (2004), 'Stylised features of price setting behaviour in Portugal: 1992-2001', *ECB Working Paper No 332*.

Dhyne, E, Álvarez, L, Le Bihan, H, Veronese, G, Dias, D, Hoffman, J, Jonker, N, Lünnemann, P, Rumler, F and Vilmunen, J (2006), 'Price changes in the euro area and the United States: some facts from individual consumer price data', *Journal of Economic Perspectives*, Vol. 20, No. 2, pages 171-92.

Ellis, C (2009), 'Do supermarket prices change from week to week?', *Bank of England Working Paper No. 378.*

Fabiani, S, Druant, M, Hernando, I, Kwapil, C, Landau, B, Loupias, C, Martins, F, Mathä, T Y, Sabbatini, R, Stahl, H, and Stokman, A C J (2006), 'What firms' surveys tell us about price-setting behaviour in the euro area', *International Journal of Central Banking*, Vol. 2, No. 3.

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 Kingdom, *Bank of England Quarterly Bulletin*, Vol. 48, No. 4, pages 404-15.

Hoffmann, J and Kurz-Kim, J-R (2006), 'Consumer price adjustment under the microscope: Germany in a period of low inflation', *ECB Working Paper No. 652*.

Jonker, N, Blijenberg, H and Folkertsma, C (2004), 'Empirical analysis of price setting behaviour in the Netherlands in the period 1998-2003 using micro data', *ECB Working Paper No. 413*.

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

Kurri, S (2007), 'Price changes in Finland: some evidence from micro CPI data', *ECB Working Paper No. 728*.

Lünnemann, P and Mathä, T (2005), 'Consumer price behaviour in Luxembourg: evidence from micro CPI data', *ECB Working Paper No. 541*.

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.

Nakamura, E and Steinsson, J (2008), 'Five facts about prices: a re-evaluation of menu-cost models', *Quarterly Journal of Economics*, Vol. 123, pages 1,373-414.

ONS (2007), *Consumer price indices technical manual, 2007 edition.*

Ritchie, F (2008), 'Secure access to confidential microdata: four years of the Virtual Microdata Laboratory', *Economic and Labour Market Review*, Vol. 2, No. 5, May.

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.

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*.