BANK OF ENGLAND

Working Paper No. 394

How do individual UK producer prices behave?

Philip Bunn and Colin Ellis

July 2010

BANK OF ENGLAND

Working Paper No. 394
How do individual UK producer prices behave?

Philip Bunn[1] and Colin Ellis[2]

Abstract

This paper examines the behaviour of individual producer prices in the United Kingdom, and uncovers a number of stylised facts about pricing behaviour. First, on average 26% of producer prices change each month, although there is considerable heterogeneity between sectors and price changes occur less frequently when measured by the average for individual products. Second, the probability of price changes is not constant over time: prices are most likely to change one, four and twelve months after they were previously set. Third, the distribution of price changes is wide, although a significant number of changes are relatively small and close to zero. Fourth, prices that change more frequently tend to do so by less. And fifth, price changes are much less persistent at the disaggregated level than aggregate inflation data imply. We find that conventional pricing theories struggle to match these results, particularly the marked heterogeneity.

Key words: Producer 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

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 participants at the Bank of England for useful comments.

This paper was finalised on 29 March 2010.

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

UK monetary policy is concerned with keeping inflation on target at 2% a year. So it is important for policymakers to consider how prices behave. In particular, the degree of nominal rigidity in the economy will influence the short-term impact of monetary policy on real activity and hence inflation. This paper uses a large database of individual producer price quotes for the United Kingdom to examine the behaviour of prices. The aim of this work is to improve our understanding about how prices are set. The results may help to shed light on which pricing theories most closely reflect how prices are determined in the real world.

There have been recent euro-area and US studies that use very large databases of individual price quotes underlying published aggregate inflation series to examine pricing behaviour. Using data that has been made available by the Office for National Statistics, this paper examines the behaviour of individual UK manufacturing output prices between 2003 and 2007 using the price quotes underlying the published Producer Price Index.

This paper uncovers a number of stylised facts about pricing behaviour. First, on average 26% of producer prices change each month. The total number of price changes is concentrated among a relatively small number of items that change price very frequently. Because a small number of items account for many price changes this means that price changes occur less frequently when measured by the average for individual products than the simple average would suggest.

UK producer prices appear slightly more flexible than in the euro area and they display a similar degree of flexibility to producer prices in the United States. There is substantial variation in the frequency of UK producer price changes between different sectors and product groups. The prices of energy products change the most often, with an average of 87% of prices changing in any given month. In general, prices appear to change more often in industries where a relatively high proportion of manufacturers' costs are accounted for by basic commodities. The prices of textile and clothing products change the least often.

The probability of price changes is not constant over time. The average frequency of UK producer prices changing increased every year between 2003 and 2007, but there is also some evidence of a correlation between the share of prices changing each month and the aggregate inflation rate. January is the most popular month for prices to change, followed by April. December is the month in which the lowest proportion of prices change. Producer prices are most likely to change one, four and twelve months after they were previously set.

There is little evidence to suggest that downward nominal rigidities are important in UK product markets since 40% of all price changes are decreases and a large proportion of those price cuts are small changes. The distribution of price changes is wide, although a significant number of changes are relatively small and close to zero. Just under 30% of all price changes are between -1% and 1%, and around 45% are between -2% and 2%. The distribution of producer price changes in the United Kingdom appears to be a little wider than in the euro area. There is substantial variation in the distribution of price changes between different industries. For

periods of up to one year, the average size of price changes tends to be smaller for items that change price very frequently, although beyond one year there is little relationship between the frequency and magnitude of price changes.

UK producer price changes are less persistent at the disaggregated level than aggregate inflation data imply. Aggregate monthly inflation rates in UK producer prices are persistent, ie the change in prices in the current month is related to the change in the previous month. But we find no evidence of persistence in monthly inflation rates at the individual item level. Our results suggest that this persistence in aggregate producer price inflation rates may be a result of aggregation across heterogeneous products rather than that persistence in inflation rates at the individual item level is reflected in the aggregate data.

The notion of nominal rigidities is a feature of many economic models. A variety of mechanisms have been put forward to explain this assumption, which can have differing policy implications. The empirical evidence presented in this paper on UK producer prices is not consistent with any one pricing theory. There are pieces of evidence that can be used to both support and detract from different theories. Variation in the share of prices changing in different years and months, and differing probabilities of prices changing depending on the time since the previous price change, are not consistent with models that assume the probability of prices changing is constant over time. But the large number of small price changes that we see in the data are not consistent with models in which firms face small fixed costs to adjust their prices. Also, the significant number of large price changes we observe are not consistent with firms receiving disutility from making large price changes. The heterogeneity across industries and product groups implies that there may not be one theory that can explain pricing behaviour at the economy wide level. Different models may better explain pricing behaviour in different sectors. The clear heterogeneity in the data would argue against the use of 'representative agent' models.

1 Introduction

Nominal rigidities imply that prices cannot freely adjust. In particular, the degree of nominal rigidity in the economy will influence the short-term impact of monetary policy on real activity and the response of inflation to changes in policy. UK monetary policy is concerned with keeping inflation on target at 2% a year. It is therefore important for policymakers to consider how prices behave and to try and understand more about the nature of these rigidities in order to improve our understanding of the monetary transmission mechanism.

The notion of nominal rigidity is a feature of many economic models. A variety of mechanisms have been put forward to explain this assumption, which can have differing policy implications. These include time-dependent models in which the probability of a price change depends only on the time since the previous price change. A simple example proposed by Calvo (1983), assumes homogenous firms have a constant probability of changing their price in each period. Alternative time-dependent models include staggered contracts in which prices are fixed for the duration of a contract, but contracts overlap in that they do not all start and end at the same time (Taylor (1980)). State-dependent pricing models typically assume that firms face a cost to adjust their price and the decision to change price depends on the state of the economy and the market faced by firms. Examples of these adjustment costs include small fixed costs of changing price (Mankiw (1985)) or disutility associated with making large price changes (Rotemberg (1982)).

One popular pricing model that incorporates assumptions about the nature of price-setting behaviour is the so-called New Keynesian Philips Curve (NKPC). This relates current inflation to future expected inflation and the deviation of marginal cost from its steady-state value. One feature of these models is that, when estimated, they imply price durations – how long, on average, it takes for companies change their prices. Early estimates of the NKPC which assume a Calvo type fixed probability of price changes, implied that firms changed their prices every fifteen to eighteen months (Gali and Gertler (1999)), although some estimates have suggested that prices change once every two and a half years (Smets and Wouters (2003)).[1]

These timings are somewhat longer than evidence from direct surveys of companies' price-setting behaviour – Blinder *et al* (1998) and Fabiani *et al* (2005) both find that the median price changes once a year in the United States and the euro area, respectively. However, Amirault *et al* (2005) found that half of Canadian firms changes prices at least once every three months. The Bank of England has also recently carried out a similar survey of pricing behaviour in the United Kingdom (Greenslade and Parker (2008)). The results show that the median firm reviews prices twice a year, but only changes price once a year.

Alongside this survey work, other recent research outside the United Kingdom has focused on examining micro-level price data that underpin published aggregate inflation series. In particular, Bils and Klenow (2004) found that these individual consumer prices change on

[1] McAdam and Willmann (2007) find that the Calvo probability parameter defines the upper limit of price duration rather than the average, when NKPC is adapted to include a state-dependent price-resetting signal.

average every three to four months in the United States,[2] while Dhyne *et al* (2005) find a longer duration of four to five quarters using a subset of microdata from the euro area CPI. Vermeulen *et al* (2007) investigate monthly micro-level producer price data for the euro area: they find that around 21% of prices change each month, implying an average price duration of five months. Nakamura and Steinsson (2007) find that a quarter of finished producer price goods change price in the United States each month, an average duration of four months.

This paper adds to this literature by examining pricing behaviour in the United Kingdom. The only previous work using individual price quote data for the United Kingdom is a study of manufacturing prices in the late 1950s and early 1960s (Godley and Gillion (1965)). We use the microdata that underpins the UK Producer Price Index to gauge the degree of price flexibility in the manufacturing sector. This data has been made available by the Office for National Statistics (ONS) for the first time for use in research work. The data is only available over a relatively short time period, but several interesting features emerge. First, on average 26% of prices change each month, although there is considerable heterogeneity between sectors and price changes occur less frequently when measured by the average for individual products. Second, the probability of price changes is not constant over time: prices are most likely to change one, four and twelve months after they were previously set. Third, the distribution of price changes is wide, although a significant number of changes are relatively small and close to zero. Fourth, prices that change more frequently tend to do so by less. And fifth, price changes are much less persistent at the disaggregated level than aggregate inflation data imply.

These findings improve our understanding of how prices are set in the United Kingdom and they may be able to help us to gauge which pricing theories are most relevant for UK price-setting. The structure of the paper is as follows. We start by discussing the details of our data set. We then present the stylised facts from our analysis, starting with the frequency of price changes and hazard functions. This is followed by our results on the size of price changes and then the relationship between the frequency and magnitude of price changes. We look at price stickiness at the individual item level and draw out the implications of all of our results for pricing theory before concluding.

2 Data

The data we use in this study are individual producer price quotes, collected by the ONS. These price quotes are the prices of individual products produced by individual firms in each particular month. The quotes are weighted and aggregated to form producer price indices (PPIs).[3] In this work, we focus on output prices – that is, the selling prices of manufacturing companies of products destined for sale in the United Kingdom. Our paper is the first to make use of this data and is the first UK study to use price microdata underlying recent aggregate inflation data to investigate the extent of pricing rigidities in the United Kingdom. We also have access to the consumer price microdata, but the PPI data were made available first and analysis of the consumer price data will be the subject of a follow-up paper. Although inflation measures based on consumer prices are the series typically targeted and monitored most closely by central banks

[2] Adjusting for sales, Nakamura and Steinsson (2007) find the median duration of retail prices is between eight and eleven months.
[3] A sample PPI release is available at www.statistics.gov.uk/pdfdir/ppi0310.pdf.

it is pricing decisions made by producers that are often modelled by economists in macroeconomic models. Producer prices are the prices charged by firms actually producing goods rather than the prices charged by retailers selling goods to consumers. Examining producer prices allows us to investigate the extent of pricing rigidities at an earlier point in the supply chain than would be the case using consumer price data.

It is worth noting that all of our results are conditional on the underlying data, where we are constrained by what was made available by the ONS. All the underlying data were accessed and analysed using the ONS' Virtual Microdata Laboratory (VML). Ritchie (2008) describes the history of the VML, and the detailed terms and conditions that apply to users.

In total, the final data set that our analysis is based on included approximately 430,000 individual producer price quotes, covering 18,000 products produced by 9,000 firms. A product is uniquely defined as being a principal output of the reporting firm(s), and as such the data are reliant on firms reporting on a consistent basis over time. Data are available at the individual product-firm level – ie price observations are supplied for each specified product at each firm in the sample.

Because firms and products enter and exit the sample on a frequent basis, the panel is not balanced and therefore price quotes are not available for every item in each period. This party reflects ONS policy to rotate survey samples, particularly for smaller firms. Unfortunately, the precise reason for an item leaving the sample is not available in the underlying data set. But we do know that the sample is updated annually to incorporate new products and changes in demand patterns for existing products, and that the methodology used for updating the PPI data set means that around a third of the sample should be rotated every year. As such, the relatively frequent entry and exit of individual products and firms is likely to reflect both regular sample rotation as well as non-response and other concerns. Overall, the high turnover rate means that around 10% of items are present in our data set for all 48 months – on average, an item is in the data for around 24 months, or two years.

ONS collect the underlying producer price data on a monthly frequency: our sample covers the period between January 2003 and December 2007.[4] For our purposes we have excluded imputed data and a negligible number of 'zero' price quotes, as we want to focus on actual price quotes.[5] Given that the price of each item is only collected once a month, we can only examine price changes at this frequency: so if a price changes within a given month that will not be captured in our data. This means that our estimates of how often prices change may be biased: we may overestimate the time between changes in manufacturing output prices if intramonth price changes are excluded from our analysis. We are not able to explicitly identify temporary price promotions or special offers in the data.

Unless otherwise stated, the results in the paper are presented on a weighted basis. Alongside the individual (unweighted) price quotes for each item, ONS also kindly made available the

[4] Unfortunately, weights were not available over a longer backrun of data prior to 2003.

[5] Approximately 3% of the raw sample is imputed. The ONS impute data where actual price quotes were not available, in most cases the imputed price is simply the price from the previous month carried forward.

appropriate weights for those items. The weights represent the individual weight of each particular item in each month in the aggregate producer price index published by the ONS. These weights are based on sales within the United Kingdom.

The individual items are identified in the microdata by a unique code, this allowed us to construct a time series of price quotes for each individual item. We used this time series to identify whether or not a price had changed by comparing to the price in the previous month. We include all observations where a price quote for the previous month is available. The results are then aggregated within each month using the weights described above.

All our results take the underlying PPI microdata as accurate. In practice, the data will be subject to both sampling and non-sampling error, as described in ONS (2008). One particular issue could be a specific form of non-sampling error: the underlying PPI survey asks respondents about their 'normal transaction price', which should be the price manufacturers achieve in a significant proportion of UK sales and representative of current output. If survey respondents find it difficult to report 'like with like' prices each month, this could introduce errors into the raw microdata. These prices are also 'average' prices over the month rather than the price on a single day. However, given the immense difficulties in identifying and compensating for these errors, we have taken the microdata as given. The next section presents results from our analysis of the data.

3 Stylised facts on UK producer prices

This section of the paper presents a set of stylised facts on UK producer price changes. We start by analysing the frequency of price changes before moving on to look at the conditional probability of price changes by estimating the hazard functions. The next section looks at the magnitude of price changes, and then we look at the relationship between the frequency and magnitude of price changes. Finally we examine price stickiness at the individual item level and compare this to persistence on the aggregate data.

3.1 Frequency of price changes

3.1.1 Aggregate frequency of price changes and comparisons with other countries

Approximately one in four UK producer prices change each month. Table 1 shows that an average of 26% of prices changed each month between 2003 and 2007. This is calculated as the total number of price changes over the total number of price quotes.[6] A large proportion of prices do not change every month, although for some firms it may be that they review their price each month and decide not to change it rather than that the presence of nominal rigidities prevents them from adjusting their price. Of the price changes we observe, approximately 60% are price increases and 40% are decreases.

[6] We drop quotes where there is no information on the price in the previous month because we are unable to measure whether the price has changed for these observations.

Our work on the PPI microdata implies that prices change more frequently than the results from the recent Bank of England pricing survey (Greenslade and Parker (2008)). The survey found that the median firm in the manufacturing sector only changes price once a year. [7] However, this result is consistent with recent work on the euro area and United States which has also found that microdata estimates imply that prices change more frequently than survey based estimates. Our results also suggest that prices change more frequently than the only other previous UK study using individual price data (Godley and Gillian (1965)), but that work covered a much earlier period. This study covered a sample of 470 manufactured products between 1955 and 1961 and found that the average interval between prices changes is around two years.

Table 1: Percentage of producer prices that change each month[(a)]

	All changes	Increases	Decreases	Period covered
United Kingdom	26.0	15.6	10.4	2003-2007
Euro-area weighted average	21	12	10	
- Belgium	24	13	11	2001-2005
- France	25	14	11	1994-2005
- Germany	22	12	10	1997-2003
- Italy	15	9	7	1997-2002
- Portugal	23	14	10	1995-2000
- Spain	21	12	9	1991-1999
United States				
- Finished goods	25	-	-	1998-2005
- Intermediate goods	27	-	-	1998-2005

(a) Euro-area data is taken from Vermeulen *et al* (2007). This paper summarises individual work from the six euro-area countries listed in Table 1. The US data comes from Nakamura and Steinsson (2007).

UK producer prices appear slightly more flexible than in the euro area. The euro-area data in Table 1 is taken from Vermeulen *et al* (2007), and this shows that only 21% of prices change in the euro area each month compared to 26% in the United Kingdom. Looking at the evidence from the individual euro-area countries, all have a frequency of price change that is lower than the United Kingdom. France comes the closest, where 25% of producer price changes change each month.

Producer prices in the United Kingdom appear to have a similar degree of flexibility to prices in the United States. The US evidence, from Nakamura and Steinsson (2007), does not report an aggregate frequency of price change of all producer prices. Instead they only report a mean share of prices changing each month for finished and intermediate goods. But at 25% and 27%

[7] The pricing survey looks at how often the median firm changes price, whereas we report the mean in Table 1. A measure of how often prices change that is perhaps more comparable to the survey results is the median number of months per price change for each item. This is seven months in our data set, closer to, but still less than the one year in the price-setting survey. Section 3.1.3 discusses these alternative measures of how often prices change in more detail.

respectively, both are very close to the UK figures. Without detailed micro-level data, it is difficult to make strong statements about the causality of price changes, as comparing micro adjustment frequencies to aggregate data is fraught with potential pitfalls – for example, we will discuss aggregation bias later on. But it is possible that greater degree of price flexibility in the United Kingdom compared with the euro area could reflect a lower incidence of contracts, implicit or otherwise, or possibly its role as a small open economy with its own currency.[8]

Comparisons between the flexibility of producer prices in the United Kingdom, the euro area and the United States must also be made with the caveat in mind that the calculations are based on different data sets and slightly different methodologies. The data cover different time periods, which may influence the comparisons of the flexibility of prices changes over time. The UK data set covers a more recent time period than the comparative work for other countries. The sample time frames used in all studies are generally characterised by relatively low and stable inflation, although producer prices did increase slightly faster in all regions of the world between 2003 and 2007 than in the sample periods of the US and euro-area studies. The rise in oil prices is one of the key factors that have contributed to increases in producer prices in recent years. If there is a correlation between overall inflation rates and the share of prices changing it is possible that the comparisons in Table 2 could overstate the true flexibility of UK producer prices relative to the euro area and the United States.

3.1.2 Frequency of price changes by product group and industry

There is substantial variation in the frequency of UK producer price changes between different sectors and product groups. Table 2 summarises the mean percentage of prices changing each month by product group. The prices of energy products (petrol and fuel in our sample) change the most frequently, with an average of 87% of all prices changing in any given month. The prices of consumer food products and intermediate goods change more frequently than the prices of consumer durables and consumer non-food non-durables. Around 40% of price changes are cuts for each of the product groups.

Table 2: Percentage of producer prices that change each month by product group[(a)]

	All changes (UK)	Increases (UK)	Decreases (UK)	All changes (euro area)
Energy products	86.9	56.3	30.6	72
Consumer food products	27.1	15.2	12.0	22
Consumer non-food non-durables	13.1	7.6	5.5	11
Consumer durables	15.6	9.5	6.1	10
Intermediate goods	25.4	15.2	10.2	22
Capital goods	18.8	11.0	7.8	9

(a) Euro-area data is taken from Vermeulen *et al* (2007).

[8] For example, Buisan *et al* (2006) find evidence that is consistent with UK manufacturers (exporters) having a lesser degree of price-setting power than in the euro area, which could also manifest in more frequent price changes.

The product groupings used in Table 2 are the same as those employed for the euro area by Vermeulen *et al* (2007) and therefore a direct comparison can be made between the flexibility of producer prices in the United Kingdom and the euro area at the product group level. UK producer prices appear to be a little more flexible than in the euro area for all of the six categories in Table 2. The trends, in terms of the variations between the different groups, are the same in the United Kingdom as in the euro area. Energy products have by far the most flexible prices, followed by consumer food products and intermediate goods.

There is also substantial variation in the frequency of price changes at the two-digit industry level, this data is shown in Chart 1 and in Table A1 in the appendix. The finding that heterogeneity is important is consistent with previous work on producer price microdata for the euro area (Vermeulen *et al* (2007)) and the United States (Nakamura and Steinsson (2007)) and with survey evidence for the United Kingdom (Greenslade and Parker (2008)). Prices of textiles and clothing products change least often among all of the two-digit industries, followed by furniture prices. Prices change substantially more frequently than the average for all products (shown as the dotted line in Chart 1) for petrol and fuel, secondary raw materials, basic metals products and other non-metallic mineral products. These are all products where a relatively high proportion of manufacturers' costs are likely to be accounted for by basic commodities that are traded and whose price changes daily. Our sample period is characterised by significant increases in the prices of a number of commodities, particularly oil.

Chart 1: Percentage of UK producer prices that change each month by two-digit industry[(a)(b)]

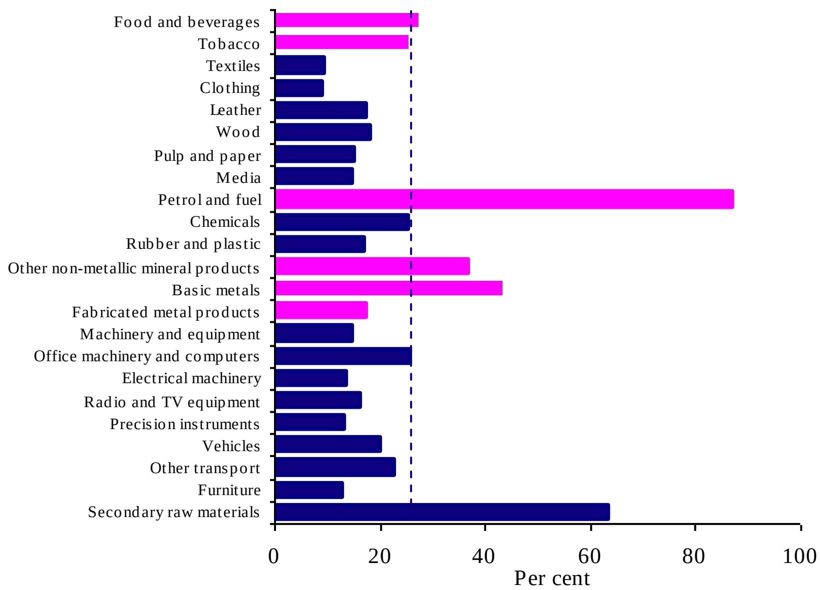

(a) The dotted line shows the average for the whole sample.
(b) The pink bars are industries which have more than 25% of inputs from agriculture, energy extraction and supply, iron, steel and non-ferrous metals.

To investigate further the possible relationship between the frequency of price changes and the type of inputs used in the production process we used the ONS Supply and Use tables to look at the proportion of inputs for each of our two-digit industries from primary industries (agriculture, energy extraction and energy supply), iron and steel and non-ferrous metals. There were six industries with more than a quarter of inputs from these sectors, indicated by the pink bars in

Chart 1. Four of these (food and beverages, petrol and fuel, other non-metallic mineral products and basic metals) are among the five industries with a higher than average share of prices changing each month. Tobacco has a very close to average share of prices changing each month and only fabricated metal products have a below average frequency of price change. Secondary raw materials is the other industry with a very high share of prices changing each month. This is essentially recycling and therefore the inputs come from across a range of industries, but output prices charged will be closely linked to prices in commodity markets.

We also considered whether there was any correlation between the frequency of price changes by industry and particular industry characteristics from the Bank of England Industry Database (BEID).[9] The motivation for this analysis was to try and better understand the industrial dispersion of the frequency of price changes. We considered the relationship between the share of prices changing by industry and profit margins and the capital to labour ratio. One hypothesis might be that firms with higher margins, or equivalently a lower labour share, may need to change price less frequently if they have more scope than firms with lower margins to accommodate changes in their costs. We typically think that labour costs change less frequently than some of the other costs faced by firms, so we also considered the relationship with the capital to labour ratio to see if firms who are more intensive users of labour changed price less often. Importantly, that second ratio should be free from possible endogeneity concerns that could affect the labour share, which itself is a function of price(s).

Chart 2: Frequency of UK producer price changes and profit margins

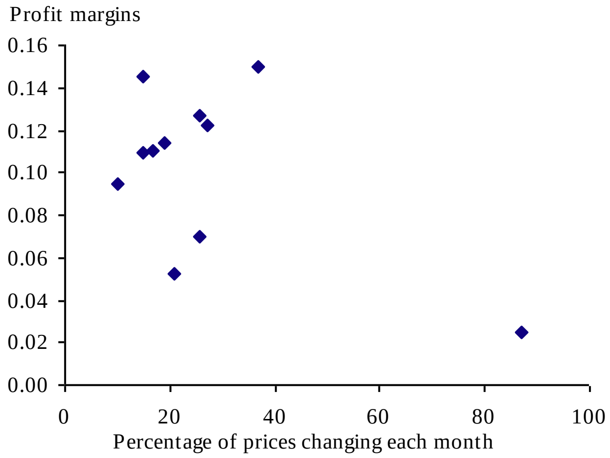

Chart 3: Frequency of UK producer price changes and capital-labour ratios

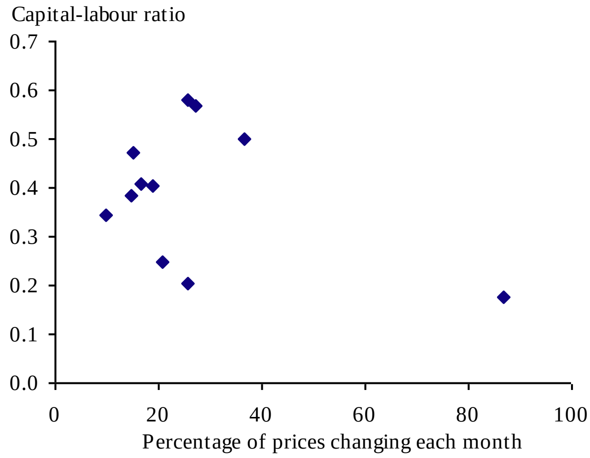

We aggregated some of the two-digit SIC industries together to match up to the ten BEID manufacturing industries so that we were able to make direct comparisons between our results and the BEID data. The mapping is given in the appendix. In summary, we found that there was no relationship between the frequency of price changes by industry and the level of profit margins in those industries. Similarly, there is no clear relationship between the capital to labour ratio and the share of prices changing each month. Chart 2 shows a scatter plot of the frequency of price changes and profit margins.[10] None of the correlations are statistically

[9] For more information on the BEID, see Groth *et al* (2004).
[10] Defined as gross operating surplus divided by gross output.

significant and the slope of the best-fit line changes from negative to positive if energy is excluded from the calculation.[11] Chart 3 shows a scatter plot between the share of prices changing each month and the capital to labour ratio.[12] Again the correlations are highly sensitive to the inclusion of energy, and none of them are statistically significant.

3.1.3 Price changes per item and time between price changes

The total number of price changes is dominated by a relatively small number of items which change price very frequently. On an unweighted basis, 10% of the 18,000 items in the data set account for almost 60% of all the price changes, although on a weighted basis these items only represent 27% of the sample. 30% of items never change price in the time they are in the data, but these items are on average in the sample for shorter periods of time. The mean number of months in the data is 14 for items whose price never changes and 29 for items that have at least one price change. The 30% of items whose price never changes only make up 13% of the pooled weighted sample. This reflects their shorter duration in the data, but also that these items have slightly below average weights, for example because this group does not include any energy products that tend to have a high average weight per item.

We examine the average duration of prices – or how much time passes, on average, between prices changes – by calculating two measures of duration. The first, and simplest, way of calculating the duration between price changes is to take the inverse of the mean share of prices that change each month. Using this simple measure, because roughly a quarter of prices change each month, the average time between prices changes is 3.9 months.[13] But this summary statistic masks a wide distribution of price frequencies – as described above, the total number of price changes is concentrated among a relatively small number of items that change price very frequently. And there are also a substantial proportion of items whose price changes less regularly than this, and have longer durations between price changes.

Our second method of calculating the duration was designed to exploit the variation across products. Rather than adding up all price changes across all the different products to calculate duration – as with our first measure – we also constructed a second measure of duration based on price changes for the average item. To do this we calculated the average number of months between price changes for each individual item (the number of months each item is in the data divided by the number of price changes), and then averaged these duration statistics across all items – in that regard, this approach is very similar to the hazard functions set out in Section 3.2. Because of the concavity of the duration-frequency relation, from Jensen's inequality, averaging after item-level inversion will yield higher duration estimates than averaging across items before inversion, unless all frequencies are identical. In addition, when we sample over products (examining the mean duration of a product-specific price change) rather than over all price changes, unless all product frequencies are equal, we will give less weight to the short durations and consequently measured average duration will be higher. Based on this second

[11] Energy is industry with by far the highest share of prices changing each month, it is the diamond on the far right of both Charts 2 and 3

[12] Capital services divided by quality-adjusted labour input.

[13] This calculation of the average duration implicitly assumes homotheticity; see Baudry *et al* (2007).

approach, the mean number of months between price changes per item is ten, significantly higher than the four months from the inverted mean frequency of price changes. The median number of months between price changes per item is seven months.

These two methods therefore yield different estimates of price duration, depending on whether we are interested in the average frequency of price changes as a whole, or on the average price duration for individual products. The difference between these two results is a function of the heterogeneity that is readily evident in the data set. Chart 4 shows the distribution of the average durations for the individual items.[14] 14% of items have an average time between price changes of between one and two months. This covers the items that change price very frequently and includes many of the energy and commodity based products. The second largest group covers items that have an average time between price changes of eleven to twelve months, this mainly includes products that tend to change price on an annual basis. For 75% of items, the average number of months between price changes is twelve months or less. A detailed breakdown of the average time between price changes by product group and industry is shown in Tables A3 and A4 in the appendix. As before, there is substantial variation between different industries and product groups.

Chart 4: Distribution of number of months between UK producer prices change per item

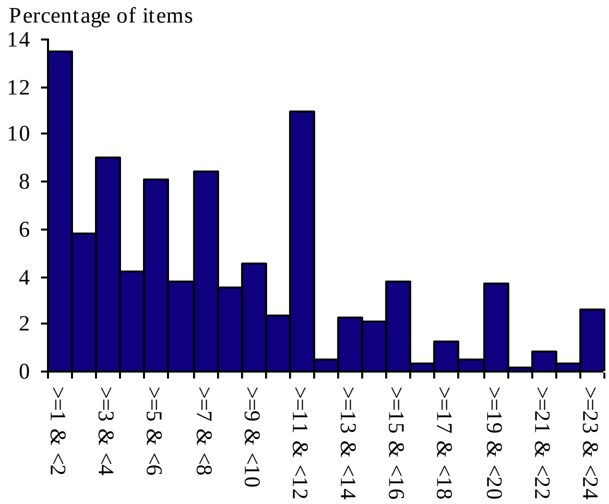

3.1.4 Changes in price flexibility over time

The frequency of price changes is not constant over time. The average proportion of prices changing each month increased every year between 2003 and 2007, rising from 24% in 2003 to 28% in 2007 (Table 3). This reflects a greater share of prices increasing each month, since the share of prices falling has been relatively stable.

[14] These are weighted together using that weight of that item in the pooled sample across the five years of the sample.

Table 3: Annual average percentage of UK producer prices that change each month

	All changes	Increases	Decreases
2003	24.2	13.5	10.7
2004	25.1	15.9	9.2
2005	25.9	15.4	10.5
2006	26.4	15.6	10.8
2007	28.0	17.4	10.6

Overall inflation rates can increase if either a higher proportion of prices rise each month (or if fewer prices fall) or if the prices that do rise increase by more (or if the prices that are reduced fall by less). Chart 5 shows that there is some correlation between the annual average share of prices increasing each month and the aggregate producer price inflation rate.[15] However, there is little relationship between the share of prices falling and overall inflation rates. The presence of a relationship between the frequency of price changes and aggregate inflation makes it difficult to draw firm conclusions about changes in underlying producer price flexibility. It may simply be that rising input costs (which could be particularly relevant for commodities in our data) forced more firms to increase their prices in the latter part of the sample.

Chart 5: Annual average percentage of UK producer prices that change each month and PPI inflation rates

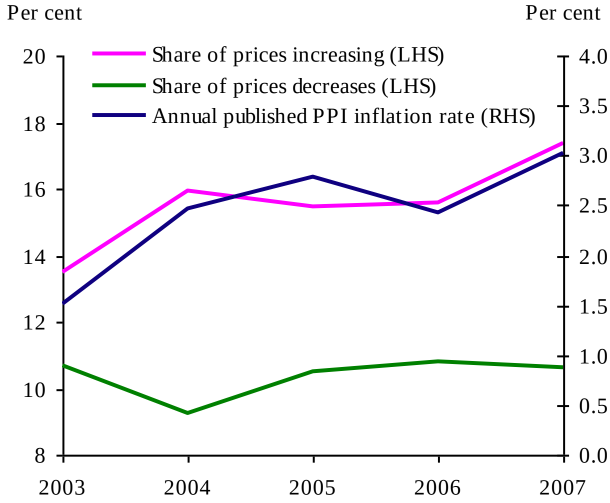

3.1.5 Frequency of price changes by calendar month

It is well known that there is seasonal variation in prices. Chart 6 shows the average share of prices changing varies across different calendar months. January is the most popular month for prices to change, followed by April. Prices are least likely to change in December. Table A5 in the appendix shows the precise numbers underlying Chart 6.

[15] Published PPI inflation rates used in this paper are taken from a vintage of data from before the (2005=100) rebasing exercise since the weights supplied to us were consistent with the previous base and weighting.

Chart 6: Percentage of UK producer prices that change each month by calendar month[(a)]

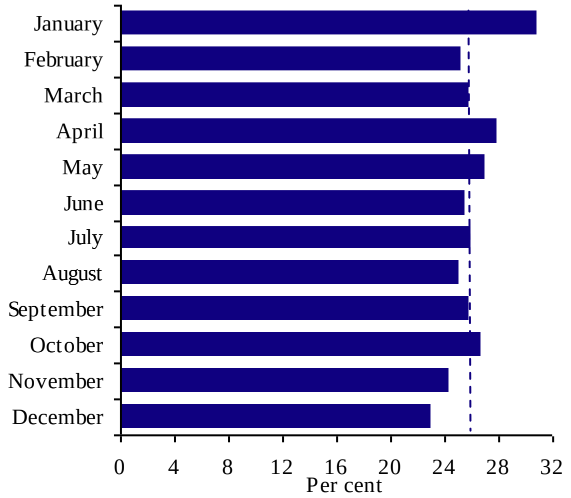

(a) The dotted line shows the average for the whole sample.

Chart 7: Frequency of UK producer price changes and average changes in aggregate PPI by calendar month

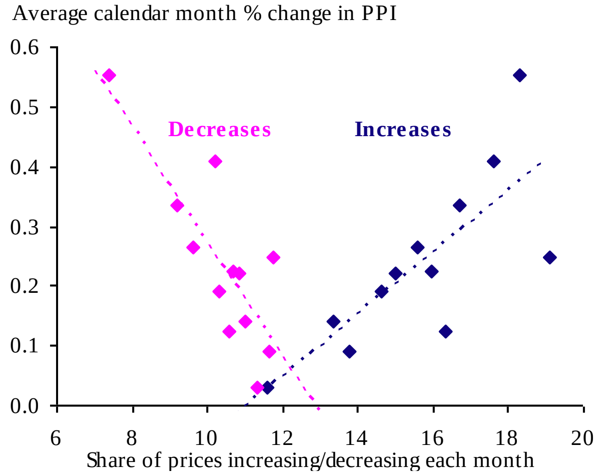

Chart 7 looks at whether there is any relationship between the size of price changes and share of prices changing in particular months. The published aggregate producer price index is not seasonally adjusted and therefore we might expect to see some seasonal variation in the month-on-month inflation rates. We find that there is a positive correlation between the share of prices changing in each calendar month and the aggregate monthly inflation rate in those months, and there is a corresponding negative relationship with the proportion of prices falling in each month. This result is consistent there being some relationship between inflation rates and the frequency of price changes.

3.2 Hazard functions

The analysis presented so far has concentrated on the average frequencies of prices changing, which can be interpreted as unconditional probabilities of price changes. In this section we look at conditional probabilities, which is the probability of a price change occurring given that we know the time elapsed since the previous change. We do this by estimating the hazard functions. The 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)*, which is known as the survivor function.

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

We only use items that have at least one price change in our estimation of the hazard functions. Items whose price never changes are excluded from this analysis. This is because we need to filter out left censored observations, ie 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, using the time between the first and the second price change (if there is one). Around 25% of items have their first observation after the first price change in 2003 H1, with the other items having their first observations spread relatively evenly across the rest of the sample.

3.2.1 Aggregate hazard functions

Chart 8 shows the estimated hazard functions for UK producer prices. The chart shows a simple unweighted version in which all items are given the same weight and a weighted version which is calculated by assigning weights to each item based on their weight in the pooled sample across the 2003 to 2007 period. Weighting makes relatively little difference to the hazard function.

The hazard function for producer prices has a large spike at one month; this implies producer prices are most likely to change in the month after they previously changed. There are also spikes at four and twelve months. The spike at twelve months suggests that some firms only adjust their prices on an annual basis. There is also a modest tick up at 24 months which could also be consistent with annual pricing reviews. The presence of these spikes in the hazard function is further evidence that the probability of a price change occurring is not constant over time. Other than the spikes identified above, the hazard function is relatively flat, although there is perhaps a very modest downwards slope. The hazard rate never falls significantly below 5% over the two-year window shown in Chart 8.

Chart 8: UK producer price hazard function

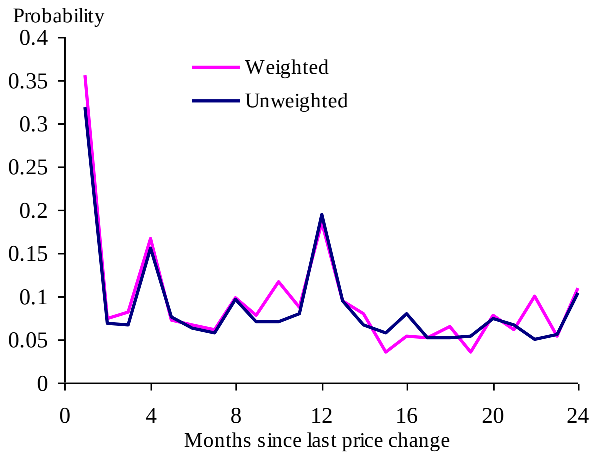

Chart 9: UK producer price survivor function

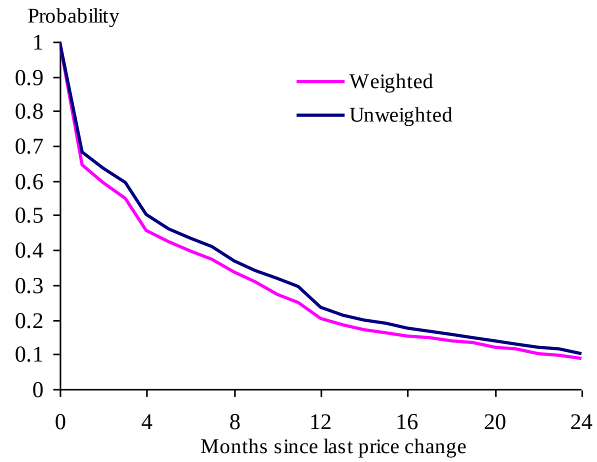

The shapes of these hazard functions are broadly consistent with those drawn for other countries in similar micro-price studies. Alvarez *et al* (2005) reports a set of stylised facts that are present in estimated hazard functions for a number of euro-area countries and for the United States using data on both consumer prices and producer prices. These facts are that the hazard rates are not zero in any period, there are annual spikes in the hazard functions, a spike at one month, and the hazard functions are downward sloping. Our hazard functions for UK producer prices clearly fit the first three facts and more debatably fit the fourth.

Chart 9 shows how the probability that a price change will not have occurred evolves over the two years since the previous change (the survivor function). By the time that three months have elapsed since the previous price change, the probability that another price change has not taken place is around 55%. By twelve months the chance that the price has not yet changed has fallen to around 20%, and by two years this probability is around 10%. These probabilities are from the weighted data, although the unweighted probabilities are only marginally higher.

3.2.2 Hazard functions by product group

Chart 10 shows the hazard functions calculated separately for each of the different product groups.[16] The hazard function for energy products is not shown because there are a relatively small number of items in this group and most change price within the first few months. The key result from the analysis of the product level hazard functions is that they all look relatively similar and have spikes in the same places. The spike at one month is largest for food products and intermediate goods, while the one-year spike is biggest for capital goods.

Chart 10: UK producer price hazard functions by product group (unweighted)

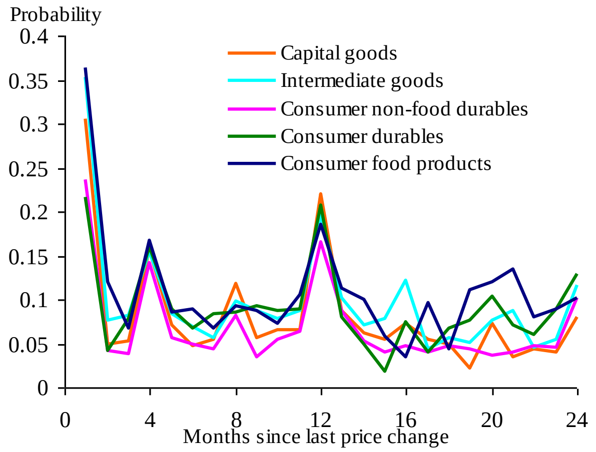

3.3 Magnitude of price changes

3.3.1 Distribution of price changes

The distribution of the size of price changes around the central estimates is wide, with a number of large price changes. But the distribution is not uniform, there is also a large proportion of price changes that are relatively small and close to zero. Just under 30% of all price changes are between 1% and -1%, about 45% are between 2% and -2%, 70% are between -5% and 5%, and 90% are between 15% and -15%. Chart 11 shows the distribution of the size of price changes, while Table 4 summarises some key percentiles in the distribution. The large share of price changes that are price falls and the large proportion of price cuts that are smaller than 5%, suggests there is limited evidence to support the presence of downward nominal rigidities in product markets in the United Kingdom.

[16] The unweighted hazard functions are shown because as the sample sizes get smaller towards the end of the two-year window the weighted versions start to become more volatile as they are dominated by small numbers of items with high weights.

Table 4: Distribution of percentage changes in UK producer prices

	All changes	Increases	Decreases
5th percentile	-12.1	0.1	-23.8
25th percentile	-1.3	1.0	-6.6
Median	0.6	2.9	-2.3
75th percentile	3.8	6.6	-0.7
95th percentile	14.3	19.4	-0.1

The broad shape of the distribution of the size of price changes appears similar to that from earlier work by Godley and Gillion (1965) in the non-engineering industries in the United Kingdom in the late 1950s and early 1960s. The distribution of producer price changes in the United Kingdom appears to be a little wider than in the euro area. Vermeulen *et al* (2007) show that the 75th and 90th percentiles in the distribution of price increases in the euro area are 5% and 13% respectively, which compares to 7% and 19% for the United Kingdom. Similarly for the distribution of price decreases, the distribution in the euro area is not as wide as in the United Kingdom. The 75th and 95th percentiles for the euro area are 5% and 14%, lower than the corresponding figures of 7% and 24% for the United Kingdom. Among the individual euro-area country results reported in Vermeulen *et al* (2007), only Portugal is found to have a wider distribution of the size of price changes than the United Kingdom.

Chart 11: Distribution of magnitude of UK producer price changes[(a)(b)]

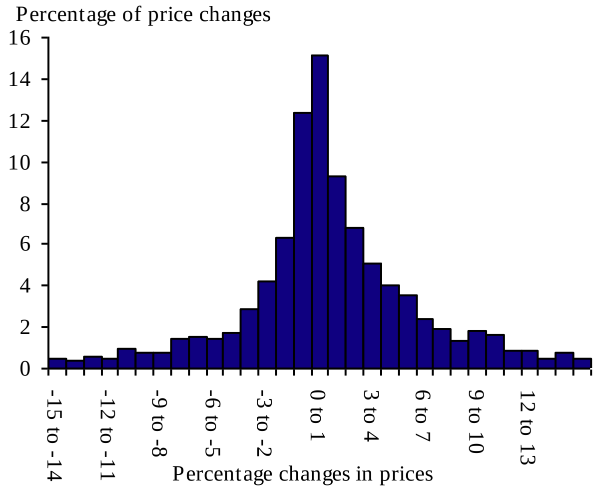

(a) Price changes are grouped into one percentage point intervals.
(b) The distribution only includes observations where the price changes.

There is significant variation across industries in the proportion of price changes that are relatively small. Charts 12 and 13 illustrate the share of prices changes that lie between -2% and 2%.[17] There tend to be fewer small price changes and hence more large changes in the prices of

[17] The differences across product groups and industries in these charts look similar if we choose a different measure of small price changes such as -1% to 1%.

energy goods than for other products. The larger changes in energy prices could be explained by the volatility in and the size of changes in oil prices over our sample period: oil prices roughly trebled between the start of 2003 and the end of 2007. In both Charts 12 and 13 the product groups and industries are sorted from top to bottom according to the share of prices that change each month. Excluding energy products, there seems to be no clear relationship between the share of price changes that are small in percentage terms and the frequency of price changes across industries.

Chart 12: Percentage of UK producer price changes between -2% and 2% by product group

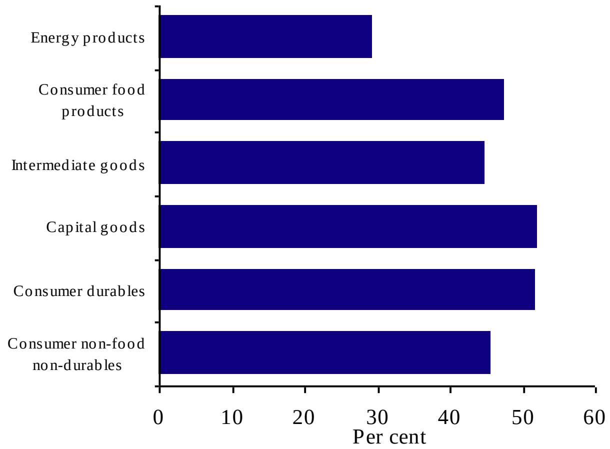

Chart 13: Percentage of UK producer price changes between -2% and 2% by two-digit industry

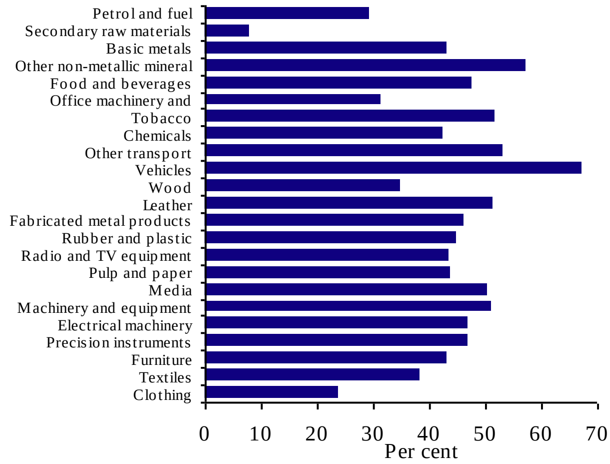

The distributions of the size of price changes by industry are summarised in Table A6 of the appendix. There is substantial variation in the width of the distribution between different industries. The distribution appears to be widest for clothing products; this applies to both the distribution of price increases and decreases. The distribution is narrowest for vehicles, at least in percentage terms. There is less variation in the width of the distributions between product groups, reflecting the fact that the groups are more highly aggregated than in the industry data. Consumer non-food non-durables have the widest distribution of price changes. One point that is consistent across all industry groups is the lack of obvious asymmetry in the distributions. Prices tend to rise on average, so more of the distribution is above zero than below – but that reflects positive inflation. As is the case for the aggregate data, there is little evidence of downward nominal rigidity in producer prices in our data in any particular industries.

3.3.2 Price reversals

In absolute terms, 2% of price changes are precise reversals of the previous price change. Of those price changes that are reversals about half are reversing a previous increase and half reverse a previous price fall. There is little difference in the duration between price reversals and price changes that are not reversals. We are not able to identify which of the price changes within our data set are due to temporary sales and price promotions, but this evidence suggests that such practices are relatively uncommon in producer prices and do not significantly affect our results.

3.3.3 Distribution of price changes by year and by calendar month

Table 5 shows that the distribution of the size of price changes by year is relatively similar in each of the five years of our sample. This is perhaps not surprising given that our sample period is characterised by relatively low and stable rates of inflation. There is a positive correlation between the median price change and aggregate inflation rate, and between the percentiles of the distribution shown in Table 5 and the growth in published PPI. Chart 14 shows the relationship between the median price change and aggregate PPI.[18] Combined with the results on the frequency of price change over time reported in Section 3.1.4, this implies that periods of higher aggregate inflation rates are characterised by both a higher proportion of prices changing and by those prices that do change increasing by more than in periods of lower inflation.

Table 5: Distribution of percentage changes in UK producer prices by year

	5th percentile	25th percentile	Median	75th percentile	95th percentile
2003	-13.3	-1.7	0.4	3.3	13.1
2004	-11.7	-1.1	0.9	4.3	13.9
2005	-11.7	-1.4	0.5	4.1	15.0
2006	-12.6	-1.6	0.5	3.0	12.6
2007	-10.0	-0.9	0.6	4.1	14.3

Although the different points in the distribution are correlated with the aggregate inflation rate, there is not a correlation with the width of the distribution and overall inflation. The width of the distribution of price changes is quite similar over time; the interquartile range is always between 4.6 and 5.5 percentage points in each year of our sample. This interquartile range only has a correlation coefficient of 0.2 with the aggregate inflation rate in PPI and the range between the 5th and 95th percentiles is negatively correlated with aggregate inflation. The distribution of the size of price changes by calendar month is also relatively similar, although not identical, across the different months of the year. This is summarised in Table A8 of the appendix.

[18] We report the median rather than the mean price change because the median is more representative of the average price change, movements in the mean price change from the microdata appear to be influenced to some extent by large price changes by a small number of observations in the tails of the distribution.

Chart 14: Median UK producer price change by year

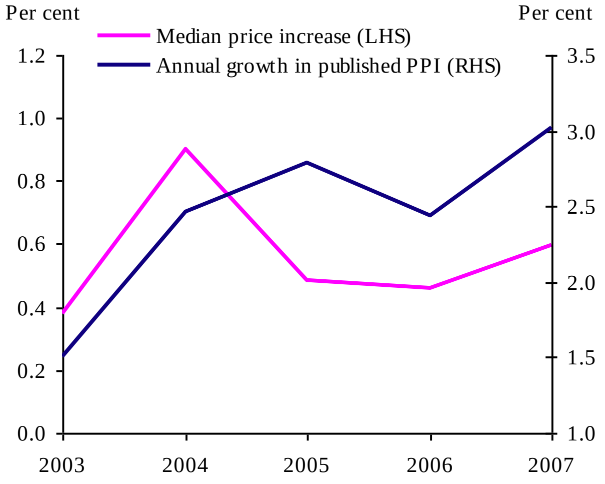

3.4 Correlation between the frequency and magnitude of price changes

This section draws together the results of the analysis of the frequency of price changes and the sizes of price changes. The analysis is pooled across the sample, and observations are weighted by their weight in the pooled sample. We have already noted that there is some relationship between the share of prices changing and the aggregate inflation rate and between the size of price changes and overall PPI inflation. We find that at the micro level, the average size of price changes is smaller for items that change price very frequently. If prices can be changed without cost in any period, there is no reason why price changes should be larger the longer it is since the previous price change. The finding that price changes tend to be smaller for items that change price frequently would be consistent with either some costs of adjustment or the presence of constraints which only allow or incentivise firms to changes prices at infrequent intervals.

Chart 15 shows a scatter plot of the size of the average price change against the number of months since the previous change. For periods of up to one year there is a strong positive correlation between the average size of price changes and the time since the previous change. For price changes that take place no more than three months since the previous change, the median price change is around 0.5%, but for price changes that take place one year since the previous change, the median price change is approximately 3%. Beyond one year since the previous price change, the correlation between the frequency and magnitude of price changes appears to break down, although the sample size for this group is substantially smaller. This is especially true as we get towards two years, as relatively few prices in the data set are unchanged for more than 18 months.

Chart 15: Median percentage UK producer price change by number of months since previous price change

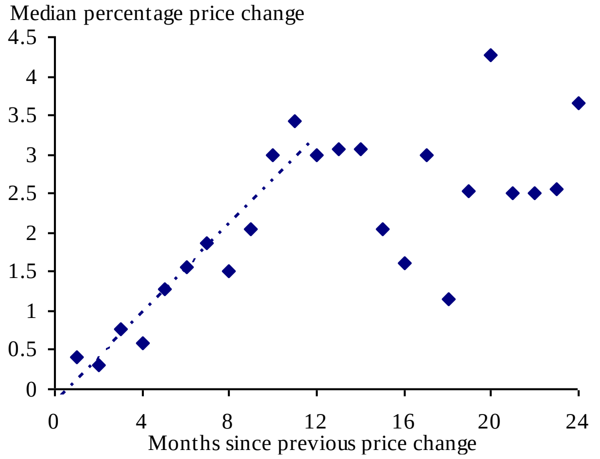

Chart 16: Median percentage UK producer price increase/decrease by number of months since previous price change

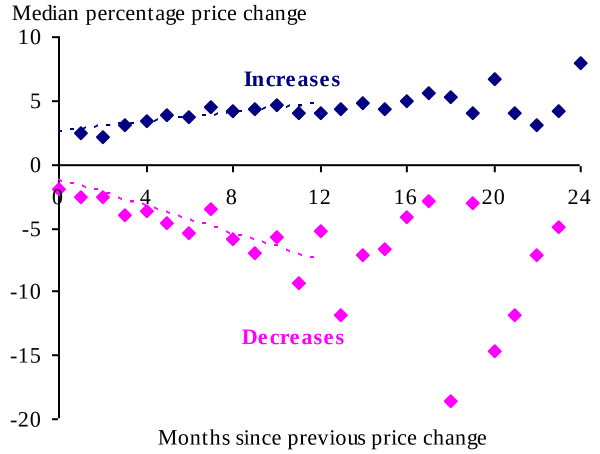

The scatter plot is separated into price increases and price decreases in Chart 16. Price increases are larger for price changes that are infrequent, and price decreases are also larger (or more negative) where the duration since the previous price change is longer. Again these relationships are reasonably good where prices change within one year of the previous change (the best-fit lines are shown as the dotted lines in the charts), but they are less robust beyond one year. One point of note is that the best-fit line for price decreases has a slightly steeper slope than the line for price increases – which could indicate some potential non-linearity in price adjustment. That could be consistent with firms being more able to pass on price cuts than recover rises in costs through higher prices.[19]

3.5 Price stickiness at the individual item level

Price changes are less persistent at a disaggregated level than they are in aggregate. A stylised fact about aggregate inflation rates is that they tend to be persistent, ie the change in prices in the current period is related to the change in the previous period.[20] To be precise, we are interested in persistence in prices in non-overlapping periods – not persistence in annual inflation rates from one month to the next, where eleven of the twelve monthly changes would remain in a calculation of annual inflation rate. By construction, the latter approach would generate persistence, and as such we have focused on the shortest-frequency changes possible, namely the month-on-month change. Some may regard this approach as containing too much 'noise' – however, if this high-frequency variation is present in our observed individual price quotes, we do not want to ignore that variation by using an approach that introduces an element of 'smoothing' into the data.

To test this hypothesis, and look for evidence of persistence in monthly inflation rates at the individual item level, we ran AR(1) regressions for the 11,761 items that have at least one price

[19] The best-fit lines in Chart 16 do not go through the origin and nor should we expect them to. The fact that the best-fit line in Chart 15 almost goes through zero is by chance rather by construction.

[20] Altissimo *et al* (2007) explore this point, finding that the aggregation process explains a fair amount of aggregate inflation persistence.

change and at least five consecutive months in the data.[21] We found that only 5% of these have a coefficient on the lagged dependent variable that is statistically significant at the 5% level. Chart 17 shows the (unweighted) distribution of the coefficients from the AR(1) regressions, and they are heavily centred around zero. By itself, the fact the coefficients are slightly centred on the negative side of zero could be consistent with lumpy item-level adjustment, which could raise concerns that the implied adjustment speed may be downwardly biased (see Caballero and Engel (2003)). However, any potential concerns about lumpy price adjustment must be considered alongside our previous findings from the microdata about the frequency and magnitude of price changes, which suggested in particular that many price changes are fairly frequent and often small in nature.

Chart 17: Distribution of AR(1) coefficients from regressions for individual items

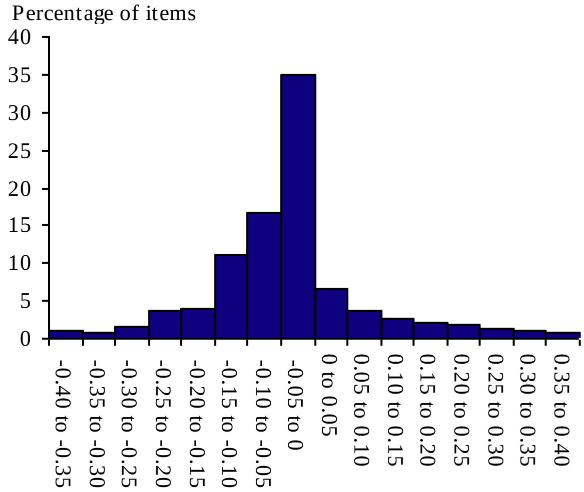

There is no evidence of persistence in month-on-month price inflation at individual item level. To confirm that we were able to find persistence in the aggregate index that our micro-price data underlies we also ran a simple AR(1) regression on monthly inflation rates from the manufacturing producer price output series over the same time period that is covered by our micro data set. The results are reported below in equation **(2)**, which shows t-ratios in parentheses:

$$\Delta P_t = \underset{(3.4)}{0.16} + \underset{(2.7)}{0.33}\Delta P_{t-1} \qquad \textbf{(2)}$$

The coefficient on our lagged dependent variable in equation **(2)** is positive and statistically significant at the 1% level. Persistence in inflation rates in the aggregate data does not appear to be a result of persistence at the individual item level being reflected in the aggregate level data. An alternative explanation is that the persistent in the aggregate data is largely a consequence of the aggregation of heterogeneous disaggregated data.[22] This result that there is less evidence of

[21] We used a univariate autoregressive process for simplicity – but thankfully there were no signs of serial correlation in the error terms, which could indicate our models were misspecified.

[22] If persistence varies across different industries or product groups then imposing a single aggregate autoregressive coefficient gives rise to a correlation between the regressors and the residuals. Imbs *et al* (2005) prove that the presence of this correlation leads to a

rigidity in disaggregated data than is apparent in macro data is not particularly new, or unique to pricing. Imbs *et al* (2005) establish it in the context of purchasing power parity.

4 Implications for pricing theory

One of the key motivations for analysing micro-price data is to understand more about how prices are set, and in particular about the degree of nominal rigidity in the economy. Nominal rigidities imply that prices cannot be freely adjusted. They are an important element of many economic models that are used in the analysis of monetary policy because they allow changes in monetary policy to affect real output, rather than changes in nominal interest rates immediately being offset by changes in nominal variables such as inflation. With fully flexible prices, monetary policy will have no effect on real output. The degree of nominal rigidity in the economy is therefore a crucial part of the monetary policy transmission mechanism and understanding more about these rigidities is a key part of understanding the effectiveness of monetary policy.

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 very different policy implications. Sticky price mechanisms can be categorised under two main headings: time-dependent and state-dependent pricing 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 firms are typically assumed to face some cost to adjust their price. The decision to change prices would then depend on the size of the gap between the optimal price and the actual price. Firms change their price when the extra profit they expect to receive from the price change outweighs the cost of price adjustment, so prices only change intermittently. Examples of these models using these adjustment costs include quadratic adjustment cost models in which firms face disutility from making large price changes (Rotemberg (1982)) and menu cost models in which firms incur small fixed costs when they change prices (Mankiw (1985)).

In the data, UK producer prices do not adjust continually – we find that only a quarter of prices change in any given month, and price changes occur less frequently when measured by individual products, reflecting the heterogeneity in the data. Some firms may review their price each month and decide not to change it, but nevertheless three quarters of firms not changing price each month would be consistent with the presence of some type of nominal rigidity in product markets.

The empirical evidence presented here is not consistent with any one pricing theory that can explain the form of those rigidities. There are pieces of evidence that both support and detract from the different models. For example, the strict Calvo price-setting model which implies a

potentially substantial bias in the estimated degree of aggregate persistence. This bias is always positive if prices are positively autocorrelated and it increases with the degree of heterogeneity.

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 spikes that we observe in the hazard functions, which show that the probability of a price changing varies depending on the duration since the previous price change. However the hazard functions are relatively flat apart from the spikes at one, four and twelve months and could still be consistent with other time-dependent type models such as staggered contracts and it is possible that downward sloping hazard functions could be a result of aggregating across heterogeneous price-setters (Alvarez *et al* (2005)).

If state-dependent models with small fixed costs of price adjustment were able to fully explain the nominal rigidities we see in the data, we might expect to see relatively few small price changes in the data. But we find that almost half of all price changes are between -2% and +2%, which we might regard as small. Such adjustment costs may still be important for the other half of price changes that are larger, or it could be that menu costs are heterogeneous, with some sectors having very small costs of adjustment but other sectors facing much larger costs. But a single aggregate model in which 'menu costs' of price adjustment explain why prices change is not consistent with our results.

Similarly, a quadratic adjustment cost model, as set out in Rotemberg (1982), also fails to match the data. Rotemberg's model, where firms minimise deviations from their optimal price subject to (quadratic) costs of changing output, suggests that firms adjust prices continuously – prices move slowly from their previous level to the new optimal level. Our analysis rejects this result – we do find evidence of infrequent price adjustment and we find that there are a number of large price changes in the data that are not consistent with gradual adjustment towards an optimal price. And, as with other theoretical models, by itself a single Rotemberg model cannot account for the observed heterogeneity in the frequency of price adjustment.

The heterogeneity that we find in pricing behaviour across different industries and product groups is perhaps the most interesting result from our study, and chimes with similar observations from 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 90% 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 clothing products change price each month, and therefore a different model may be needed to explain the nominal rigidities in this sector. This heterogeneity would argue against the use of 'representative agent' type models.

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 in pricing behaviour across different sectors is also clear in the results of the pricing survey.

Our analysis of the UK producer price microdata suggests that there is less evidence of nominal rigidity at the firm level than in the aggregate data and that aggregation of heterogeneous disaggregated data may overstate the true degree of price stickiness in the economy. This result is relevant to all classes of models, whatever the source of the nominal rigidity.

5 Conclusions

This paper has examined pricing behaviour at the individual item level for manufacturing firms in the United Kingdom. In doing so, we have added to the growing literature of micro-pricing studies, providing the first set of recent UK results using data underlying official inflation statistics.

Using the data that underpins the UK Producer Price Index, we have uncovered several interesting features about the behaviour of those prices. First, on average 26% of prices change each month, although there is considerable heterogeneity between sectors and product groups. 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: prices are most likely to change one, four and twelve months after they were previously set. Third, the distribution of price changes is wide, although a large number of changes are relatively small and close to zero. Fourth, prices that change more frequently tend to do so by less. And fifth, price changes are much less persistent at the disaggregated level than aggregate inflation data imply.

These results suggest that none of the conventional theories for price stickiness are borne out by the data. In particular, the marked degree of heterogeneity in the behaviour of prices is often just ignored, for example in the typical ‘representative agent’ models. Furthermore, the notable inflation persistence we observe at the aggregate price level is simply not present at the micro level – or in other words, inflation persistence appears to just reflect aggregation bias. Together these results imply that, if we really want to understand and model prices with any degree of accuracy, we need to find a way of capturing the richness of the heterogeneity that is present in the data, while matching the time-series properties of inflation that arise from aggregation. One option here could be to further pursue the so-called ‘factor augmented vector autoregression’ models set out by Boivin *et al* (2007) and Mumtaz *et al* (2009). An alternative avenue would be to explore some of the dynamic programming analysis (eg Miranda and Fackler (2002)) that is typically more prevalent in other fields, such as consumption theory. But, whatever direction future work takes, if we want to use genuinely micro-founded models – ie models that match the heterogeneity in the microdata – we may need to develop the pricing models that are currently employed.

Appendix: Detailed tables of results

Table A1: Percentage of UK producer prices that change each month by two-digit industry

	All changes	Increases	Decreases
Food and beverages	27.4	15.4	12.0
Tobacco	25.6	17.7	7.8
Textiles	9.4	5.8	3.5
Clothing	9.0	4.9	4.1
Leather	17.7	9.6	8.0
Wood	18.1	11.7	6.4
Pulp and paper	15.3	8.1	7.2
Media	14.9	8.8	6.1
Petrol and fuel	86.9	56.3	30.6
Chemicals	25.5	15.7	9.8
Rubber and plastic	16.9	10.0	6.9
Other non-metallic mineral products	36.7	20.0	16.7
Basic metals	43.3	27.3	16.1
Fabricated metal products	17.4	11.5	5.9
Machinery and equipment	14.8	10.1	4.7
Office machinery and computers	25.8	9.8	16.0
Electrical machinery	13.6	8.3	5.2
Radio and TV equipment	16.2	6.6	9.6
Precision instruments	13.3	7.8	5.5
Vehicles	20.0	12.0	8.0
Other transport	22.7	13.8	8.9
Furniture	13.0	7.9	5.1
Secondary raw materials	63.4	35.4	28.0

Table A2: Mapping from two-digit SIC industries to BEID industries and percentage of prices changing each month

BEID industry	SIC industries	All changes	Average PPI weight
Manufactured fuel	Petrol and fuel	86.9	7.2
Chemicals and pharmaceuticals	Chemicals	25.5	7.4
Non-metallic mineral products	Other non-metallic mineral products	36.7	3.4
Basic metals and metal goods	Basic metals Fabricated metal products	25.8	9.4
Mechanical engineering	Machinery and equipment	14.8	6.4
Electrical engineering and electronics	Office machinery and computers Electrical machinery Radio and TV equipment Precision instruments	16.7	9.9
Vehicles	Vehicles Other transport	20.9	11.4
Food, drink and tobacco	Food and beverages Tobacco	27.3	17.4
Textiles, clothing and leather	Textiles Clothing Leather	10.0	3.4
Paper, printing and publishing	Pulp and paper Media	15.0	13.1
Other manufacturing	Wood Rubber and plastic Furniture Secondary raw material	18.9	10.9

Table A3: Number of months between price changes by product group

	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)
Energy products	1.2	1.0	3.0	0.0	0.0
Consumer food products	3.7	6.5	9.2	7.3	23.6
Consumer non-food non-durables	7.7	11.8	14.8	21.1	40.0
Consumer durables	6.4	7.9	10.1	8.1	29.0
Intermediate goods	3.9	6.7	9.8	12.9	27.0
Capital goods	5.3	7.9	12.0	16.4	31.7
Total	3.9	7.3	10.3	12.6	29.8

Table A4: Number of months between price changes by two-digit industry

	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 beverages	3.7	7.0	9.6	7.7	22.8
Tobacco	3.9	3.7	3.7	0.0	22.2
Textiles	10.7	11.0	13.3	25.9	42.3
Clothing	11.1	13.5	19.2	30.1	55.5
Leather	5.7	10.8	12.3	23.0	39.2
Wood	5.5	5.9	9.6	12.7	30.9
Pulp and paper	6.5	8.5	11.2	17.8	30.0
Media	6.7	11.8	13.0	20.5	42.2
Petrol and fuel	1.2	1.0	3.0	0.0	0.0
Chemicals	3.9	8.4	12.2	12.0	25.8
Rubber and plastic	5.9	8.6	12.1	13.9	27.5
Other non-metallic mineral products	2.7	5.4	8.3	4.6	18.8
Basic metals	2.3	2.8	5.8	3.0	14.0
Fabricated metal products	5.7	7.3	10.0	19.7	33.3
Machinery and equipment	6.8	10.3	13.9	13.9	32.6
Office machinery and computers	3.9	3.9	8.5	9.2	34.1
Electrical machinery	7.4	11.0	13.4	22.8	39.6
Radio and TV equipment	6.2	7.7	9.8	35.6	32.6
Precision instruments	7.5	11.8	16.1	16.8	34.4
Vehicles	5.0	6.8	9.3	6.4	26.8
Other transport	4.4	8.8	11.5	10.6	32.3
Furniture	7.7	10.8	14.2	13.0	26.8
Secondary raw materials	1.6	1.4	1.9	0.0	0.0
Total	3.9	7.3	10.3	12.6	29.8

Table A5: Percentage of UK producer prices that change each month by calendar month

	All changes	Increases	Decreases	Mean monthly % change in aggregate PPI
January	30.8	19.1	11.8	0.25
February	25.2	15.6	9.6	0.26
March	25.7	18.3	7.4	0.55
April	27.8	17.6	10.2	0.41
May	26.9	16.3	10.6	0.13
June	25.4	13.8	11.7	0.09
July	25.9	16.7	9.2	0.34
August	25.0	14.6	10.3	0.19
September	25.8	15.0	10.8	0.22
October	26.6	16.0	10.7	0.22
November	24.4	13.3	11.0	0.14
December	22.9	11.6	11.3	0.03

Table A6: Distribution of percentage changes in UK producer prices by two-digit industry

	5th percentile	25th percentile	Median	75th percentile	95th percentile
Food and beverages	-13.5	-1.3	0.3	3.2	14.4
Tobacco	-1.7	-0.2	0.5	4.7	9.2
Textiles	-16.8	-1.5	0.8	3.6	16.8
Clothing	-42.4	-5.1	0.6	6.8	73.6
Leather	-9.4	-1.4	0.2	3.0	12.0
Wood	-8.6	-1.0	1.3	5.0	21.1
Pulp and paper	-15.0	-2.1	0.1	3.4	12.2
Media	-20.8	-0.9	0.3	3.4	22.3
Petrol and fuel	-10.2	-1.8	1.5	5.7	12.1
Chemicals	-14.3	-1.4	0.7	3.5	15.8
Rubber and plastic	-13.8	-1.3	0.5	3.5	16.1
Other non-metallic mineral products	-7.6	-1.0	0.1	2.3	10.6
Basic metals	-6.7	-1.0	1.0	3.8	13.4
Fabricated metal products	-11.0	-0.6	0.8	3.9	16.6
Machinery and equipment	-11.9	-0.4	0.7	3.1	18.0
Office machinery and computers	-28.5	-9.3	-1.3	1.7	37.1
Electrical machinery	-13.7	-0.9	0.5	3.5	25.4
Radio and TV equipment	-23.1	-4.8	-0.5	1.0	16.5
Precision instruments	-17.1	-1.7	0.4	2.7	15.0
Vehicles	-8.0	-0.5	0.2	1.6	8.5
Other transport	-7.5	-1.0	0.5	2.1	14.1
Furniture	-16.2	-1.2	0.5	3.6	21.2
Secondary raw materials	-19.3	-4.8	3.3	10.3	19.3

Table A7: Distribution of percentage changes in UK producer prices by product group

	5th percentile	25th percentile	Median	75th percentile	95th percentile
Energy products	-10.2	-1.8	1.5	5.7	12.1
Consumer food products	-13.7	-1.2	0.3	3.3	14.4
Consumer non-food non-durables	-24.0	-1.2	0.3	3.7	25.4
Consumer durables	-14.3	-0.9	0.4	2.4	14.8
Intermediate goods	-10.1	-1.3	0.6	3.5	14.5
Capital goods	-17.6	-1.3	0.3	2.3	17.1

Table A8: Distribution of percentage changes in UK producer prices by calendar month

	5th percentile	25th percentile	Median	75th percentile	95th percentile
January	-11.4	-1.5	0.8	4.0	15.2
February	-13.2	-1.5	0.9	3.6	13.5
March	-10.0	-0.2	1.7	5.8	17.2
April	-15.7	-1.1	0.8	5.0	14.8
May	-13.3	-2.0	0.5	4.1	12.2
June	-10.3	-2.0	0.2	3.2	14.3
July	-7.9	-0.7	0.7	3.7	15.2
August	-10.0	-1.2	0.4	3.3	12.9
September	-14.2	-1.3	0.5	3.6	12.1
October	-12.2	-1.4	0.5	3.1	14.1
November	-10.4	-1.9	0.2	3.2	14.4
December	-12.1	-2.3	0.0	2.2	13.2

References

Altissimo, F, Mojon, B and Zaffaroni, P (2007), 'Fast micro and slow macro: can aggregation explain the persistence of inflation?', *European Central Bank Working Paper No. 729*.

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

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

Baudry, L, Le Bihan, H, Sevestre, P and Tarrieu, S (2007), 'What do thirteen million price records have to say about consumer price rigidity?' *Oxford Bulletin of Economics and Statistics*, 69(2), pages 139-83.

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

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

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

Buisan, A, Learmonth, D and Sebastiá-Barriel, M (2006), 'UK export performance by industry', *Bank of England Quarterly Bulletin*, Vol. 46, No. 3, pages 308-16.

Caballero, R and Engel, E (2003), 'Adjustment is much slower than you think', *NBER Working Paper Series No. 9898*.

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

Dhyne, E, Álvarez, L, Le Bihan, H, Veronese, G, Dias, D, Hoffman, J, Jonker, N, Lünnemann, P, Rumler, F and Vilmunen, J (2005), 'Price setting in the euro area: some stylized facts from individual consumer price data', *European Central Bank Working Paper No. 524*.

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

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

Godley, W and Gillion, C (1965), 'Pricing behaviour in the manufacturing industry', *National Institute Economic Review*, Vol. 33, pages 43-47.

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.

Groth, C, Gutierrez-Domenech, M and Srinivasan, S (2004), 'Measuring total factor productivity for the United Kingdom', *Bank of England Quarterly Bulletin*, Vol. 44, No. 1, pages 63-73.

Imbs, J, Mumtaz, H, Ravn, M and Rey, H (2005), 'PPP strikes back: aggregation and the real exchange rate', *Quarterly Journal of Economics*, Vol. 120, pages 1-43.

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

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

Miranda, M and Fackler, P (2002), *Applied computational economics and finance*, MIT press.

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

Nakamura, E and Steinsson, J (2007), 'Five facts about prices: a re-evaluation of menu-cost models', Columbia University, *mimeo*.

ONS (2008), *Summary quality report for Producer Price Indices*.

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.

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.

Vermeulen, P, Dias, D, Dossche, M, Gautier, E, Hernando, I, Sabbatini, R and Stahl, H (2007) 'Price setting in the euro area: some stylised facts from individual producer price data', *European Central Bank Working Paper No. 727*.