

**PRICE FLEXIBILITY IN BRITISH SUPERMARKETS:
MODERATION AND RECESSION**

**HUW DIXON, JONATHAN S SEATON AND MICHAEL
WATERSON**

No 1041

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

PRICE FLEXIBILITY IN BRITISH SUPERMARKETS: MODERATION AND RECESSION

HUW DIXON, JONATHAN S SEATON AND MICHAEL WATERSON*¹

Draft: February 2014

ABSTRACT

This paper delivers a significantly different empirical perspective on micro pricing behaviour and its impact on macroeconomic processes than previous studies, largely resulting from the fact that our weekly price data for the three major British supermarkets spans a seven year period including the crisis years 2008-2010. We find that there is a large and significant change in the behaviour of prices from 2008 onwards: prices change more frequently and the average duration of price spells declines significantly. Several of our findings run strongly counter to established empirical regularities, in particular the high overall frequency of regular or reference price changes we uncover, the greater intensity of change in more turbulent times and the numerical dominance of price falls over rises. The pricing behaviour revealed also significantly challenges the implicit assumption that prices are tracking cost changes.

JEL numbers: E30; E31; L81.

Keywords: Micro pricing, price flexibility, regular prices, menu costs.

*Corresponding author: Department of Economics, University of Warwick, Coventry CV4 7AL, UK. Phone +442476523427; fax +442476523031; email: Michael.watson@warwick.ac.uk

¹We are grateful to Paul Dobson, Francesco Lippi, Michael McMahon, Thijs van Rens and Silvana Teneyro for helpful discussion of ideas expressed in this paper; in particular Paul initiated early collection of the data. Kun Tian provided some research assistance. We also thank attendees at the RES 2013 conference presentation and seminars at EIEF, Rome, October 2013 and at Brunel University for helpful comments. We would like to acknowledge support of the UK ESRC (Grant RES-062-23-1962) for this work. Michael Waterson is a member of the UK Competition Commission, but he was not involved in any of its inquiries into this industry nor has he had access to any data they collected that are not in the public domain. This paper should not be taken to reflect the views of the Competition Commission.

1. Introduction

A key question for macroeconomics and in particular monetary policy is how flexible prices are and to what extent prices at the micro level respond to macroeconomic phenomena such as inflation. It is a question on which there is growing microeconomic evidence, and some limited consensus. As an empirical question, the answer may differ from country to country and whilst most studies focus on the US (Bils and Klenow 2004, Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), there are also examinations of other economies, notably the Inflation Persistence Network (IPN) in the Euro area (Dhyne et al, 2005), and for the UK (Ellis, 2009; Bunn and Ellis, 2011, Dixon and Tian 2013) and recently Central and South America (Gagnon, 2009; Cavallo, 2012). The US findings are usefully presented as a series of “facts”, as in Nakamura and Steinsson (2008; 2010) and in Klenow and Malin (2010); of course the “facts” are particularly useful if they have a very broad application across countries and events. We take an important new British sample with some clear distinguishing features and report findings at considerable variance to many that have gone before. In particular, our data covers the period 2004-2010 and so includes the crisis years 2008-2010 and the transitional year 2007, enabling us to observe responses to a greater range of inflationary and deflationary pressures. Therefore they present new challenges both for modelling and the facts. To preview, we can say unequivocally that the prices we study become flexible in the crisis period. But quite what that implies is unclear.

We report several powerful findings that challenge a number of the empirical conclusions that have been drawn in previous studies of micro pricing behaviour, whilst reinforcing some others. Our work is related most directly to those studies that have used either “regular” or “reference” prices at the individual item level and it relates to a constant set of largely processed grocery products. “Regular” or “reference” prices are constructed so as to filter out short run posted price changes (for example due to sales) which are seen as better reflecting underlying price behaviour. Our work also makes significant use of a key institutional feature of supermarket pricing in Great Britain. As with Cavallo’s (2012) recent paper, our prices are drawn from websites and not scanner-based, representing in-store prices accurately because of the uniform pricing policies the key supermarkets adopt.

We can summarize our results in the form of five stylized facts.

Stylised fact 1. For the four year period 2004-7, the price data is fairly similar to the UK CPI data in terms of the frequency of price change and various measure of duration.

Our data contrasts somewhat with Klenow and Malin's stylised fact 1 that posted prices change at least once per year (Klenow and Malin 2010): our data indicates that there is more stickiness in posted prices in the pre-crisis period than would be consistent with this. In the period 2004-6 we find that there are many products (as high as 30%) that do not change posted price in a given calendar year, and over 20% of products have at least one posted price-spell lasting 100 weeks or more. However, our dataset is consistent with the equivalent UK CPI microdata (Bunn and Ellis (2012), Dixon and Tian (2013)).

Stylised fact 2. With the onset of the crisis, there is a dramatic change in the behaviour of prices: (a) the frequency of price changes increases, with a particularly large increase in the proportion of price cuts, (b) the average duration of prices falls by about 50% under all measures of duration, (c) almost all prices see an increase in the number of price-changes.

Stylised fact 2 (SF2) is not so surprising. However, what is new is that even regular and reference prices become short-lived: in table 3 below we will show that three popular measures of regular/reference prices each show mean price-spells of significantly less than 6 months. This is significantly different from Klenow and Malin's fact 3, that regular/reference prices last about one year. The fact that the frequency of price change increases in a recession is in line with Vavra (2013) who finds that the frequency is countercyclical in US CPI data. Most importantly, SF2 indicates that pricing behaviour responds to macroeconomic events, something that was not so apparent in studies based on the moderation period prior to 2008.

SF1-2 together indicate that pricing behaviour does respond to the macroeconomic environment. In that sense, time-dependent pricing models will not give a robust guide to price-stickiness and the state-dependent framework is supported. However, there is a strong caveat to this. *It is a very large shock indeed in the form of the crisis that gives rise to this effect.* As we know from previous studies, pricing has been pretty stable from the mid-90s up to 2008. Although we call this the moderation period, there were many substantial shocks that occurred in this period: the dotcom bubble bursting, large fluctuations in exchange rates (for example in the Euro-dollar), 9/11. These do not seem to have been large enough to substantially affect price-stickiness. Monetary policy in itself does not substantially affect the economy and is not likely to influence price-stickiness. Hence although we find evidence

for a move to state-dependent pricing, *it does not indicate* that a time-dependent approach is inappropriate for normal non-crisis times or for the analysis of monetary policy. Insofar as the effects of monetary policy are small, they will not affect the pricing behaviour of firms.

Stylised fact 3. In the crisis years, the frequency of downward price changes is much larger than the frequency of price increases. This is also true to a much lesser extent in all years.

This is a new finding and different from what we find in the CPI micro price-data for the UK and other countries - where price cuts are in general less common than price increase. It is possibly particular to the prices covered in this dataset. In our sample, there is a proliferation of 1 penny and 2 penny price cuts during the crisis. The average size of price-increases is much larger during the crisis than before, with a peak in 2008. The average size of price-cuts is more stable, with little change.

From SF3, we get a picture of “Edgeworth cycles” occurring in the crisis, sequences of small price cuts followed by large price increases. Certainly, some very stylised Edgeworth cycles are observable in the data (see Seaton and Waterson, 2013). We believe that this is a new finding and results not only from the fact that there was a crisis but also the oligopolistic structure of major grocery retailers in the UK. The three stores in our dataset are in direct competition with each other: undercutting their competitors by just 1 penny across a range of products enables them to market themselves as giving better value in a more competitive environment where consumers’ real incomes were squeezed post January 2008.

The distribution of price-growth has also been the object of recent attention by Midrigan (2011), Alvarez and Lippi (2013), Alvarez et al (2013), Vavra (2013). Our findings challenge some of the empirical results based on US and French CPI data.

Stylised Fact 4. The dispersion of price growth as measured by the Standard Deviation does not change significantly as a result of the crisis: dispersion as measured by the IQR falls if we include all price changes. Skewness is insignificant throughout the whole period. The distribution of price growth is leptokurtic and becomes more so in the crisis. Only part of this increase is excess Kurtosis explained by the proliferation of small price-cuts.

This contrasts with existing studies. In particular, Vavra (2013) finds in the US data that the standard deviation of price growth is significantly countercyclical and Kurtosis pro-cyclical. In our data we find evidence of neither: indeed, we find some evidence for the exact opposite

of what Vavra finds in the US data. Our finding is consistent with the standard S,s menu cost model with an aggregate shock and not with Vavra's extension to an uncertainty shock.

Turning from the distribution of price-growth to the distribution of price levels, we find that the level of disaggregation plays a crucial part in determining what we find:

Stylised fact 5. Looking across all prices, there was no significant change in price dispersion over the period. However, if we disaggregate to the product level, we see that there was a significant increase in price-dispersion from 2007 onwards, reflecting increased competition between the three retailers.

Stylised fact 5 reflects the fact that most price-dispersion is across products: the structure of prices across products does change, but it does not seem to be affected by macroeconomic factors. In our data, there was little within-product price dispersion across the three retailers: however, with the onset of the crisis this within-product dispersion increases.

Our results thus indicate that the crisis had a major effect of increasing the frequency of price changes and decrease in the duration of price-spells even if they are filtered to exclude short lived price changes. There was a big increase in small price cuts in the crisis which we believe resulted from the oligopolistic rivalry between the three stores, particularly Asda and Tesco. We believe that the existing models of pricing used in macroeconomics based on the monopolistic competition framework cannot explain this phenomenon.

Section 2 of the paper discusses existing data and filtering definitions and explains our approach to them. We then move on in section 3 to a description of the data at our disposal, including the institutional features that make our approach possible. Section 4 explores the various dimensions of pricing- frequency, magnitude, timing and direction. Section 5 and 6 explore the distributions of price growth and price levels respectively. Section 7 discusses the implications of our findings, and section 8 concludes.

2. Definitions of Prices and durations.

We make use of four basic definitions of prices. First, we have the raw data, *posted prices*, those that a consumer wishing to purchase one unit of the product in store or online

would face.² These show the most frequent movements in our sample and include temporary sales. The question of whether to include sales or not has been a matter of debate. For some products, sales represent a natural part of the product cycle (particularly seasonal goods such as clothing and footwear). There are also the traditional seasonal sales periods such as Christmas and the New Year. Whilst the sales might be temporary, they often are associated with a large volume of purchases (Griffith et al, 2009). Also, whilst some sales are calendar based, their exact timing might reflect macroeconomic factors: low demand in the run up to Christmas might result in the Christmas sales being brought forward or extended to last longer than usual. The argument for filtering sales out of the data is that the economic rationale for sales is largely microeconomic: sales enable price-discrimination between the informed and uninformed or between the patient and impatient (see for example Varian, 1980, Guimaraes and Sheedy 2011). Hence in addition to the raw data we consider three different filters that “purge” the raw data of temporary price changes in some way.

The first way of filtering the raw data is what Nakamura and Steinsson (2008) define as *regular prices*. As Nakamura (2008) graphically points out, certain products are subject to a form of seesaw price movement where there are frequent offers but based upon a largely unchanging regular price. Their algorithms correct for this “V” or “U” shaped price phenomenon,³ by removing from posted prices short-term reductions that are later reinstated, in part or fully. Here we use the “B” version of their algorithm that corrects only for price cuts where price returns to its previous level, using precisely the definition for weekly data used in Nakamura (2008), so counting a price as regular if the price falls below that level for six weeks or less before returning to that level.⁴ These prices we call *NSB regular prices*. The choice of 6 weeks is possibly controversial: it could well be argued that even if a price returns to its “regular value”, 6 weeks is not a short duration and represents something more

²This represents the price on the shelf and at the till, assuming single item purchase, no coupons, etc. Our data are *not* scanner data, so it is not average selling price. Thus it is nearer to a regular price than scanner data yielding average selling price would be.

³ Clearly, there are some complications, for example if the price “returns” to a different level, which they deal with through variant “A”. There are also choices to make regarding the length of the time interval.

⁴ We did attempt some experiments with the “A” version, but found sufficient ambiguities in working with our data that we do not adopt it here, choosing instead to use the Kehoe-Midrigan algorithm. The essential difficulty lies in unambiguously detecting a unique new regular price that is different from the current regular price when prices move in a variety of ways and seldom stay fixed for long.

than a temporary sale to attract the informed consumer. It means that seasonal sale periods like such as Christmas and the New Year might be viewed as “temporary”.⁵

The second filter is from Kehoe and Midrigan (2012) and Midrigan (2011), who also use the term *regular price* but adopt a somewhat different algorithm. The most obvious difference is that whereas the NSB filter removes only short-lived price cuts from the data, Kehoe and Midrigan also remove short-lived price rises. The resultant algorithm creates a modified version of a running mode: “The regular price is thus equal to the modal price in any given window surrounding a particular period, provided the modal price is used sufficiently often” (Kehoe and Midrigan, 2012). We adopt a quarterly window and call the resultant prices *KM13 regular*, to make the distinction from Nakamura and Steinsson clear.⁶ This series is in some senses intermediate between NSB regular prices and the fourth definition, due to Eichenbaum et al (2011).

Finally, we have the third filter of *reference prices*, as defined by Eichenbaum et al. (2011). These aim to strip out short term phenomena by replacing the posted price in any week with the most common price in the calendar quarter in which it lies, i.e. the quarterly mode.⁷ We develop three slight variants to the EJR reference prices as explained in section 4 below. In essence, the EJR reference price is simply the *modal price* in a particular period (e.g. calendar quarter): the variants we use differ in the way we choose the period in our weekly dataset. The modal price is perhaps a rather extreme concept of reference price: the mode might still be uncommon. In theory it is possible that a price that is charged for only two weeks might be the mode (i.e. the prices are all different in the remaining weeks of the quarter). With monthly data this is not so much of an issue, since there are only 3 observations in a quarter.

In our analysis below, we use all four definitions of price: posted, NSB regular, KM13 regular, and EJR reference prices. Posted prices will be most volatile and, save in unusual circumstances, reference prices least so. NSB regular prices will tend to have a higher mean duration than posted and reference prices since they include temporary price

⁵ Robustness experiments in Seaton and Waterson (2013) suggest this choice of 6 weeks is not a significant issue, however.

⁶ More detail on our construction of this is given in Appendix 1.

⁷ As is well known, the mode of a distribution need not be (uniquely) defined. This appears not to be a significant issue for Eichenbaum et al, although we observe occasions in our data where this arises in calculating KM regular prices. The issue is discussed in that context in Appendix 1.

risers but be less than KM13 which excludes all temporary prices. EJR is potentially the longest duration depending on the length of window chosen.

We will also look at the duration of prices from the perspective of two alternative *distributions*. First, following most studies, we look at the distribution of price-spell durations. This ignores the panel structure of the dataset and treats each price-spell as an individual element in the population. Our data set has 41,598 price spells. The distribution for price-spells simply tells us the proportion of price spells with length i in this total: for example, there are 7,625 price spells which last just one week, so that the share of one week spells is 18%. Conceptually, the distribution across spells treats each price spell as equal and does not depend on how long the spell lasts.

The second concept is the *cross-sectional* distribution of durations. This gives the distribution of price-spell durations observed at a representative or randomly chosen point in time. This is analogous to the cross-sectional distribution of *ages* observed at a point in time (as in the population census), but rather than using the incomplete duration (age), we instead measure the *completed* price-spell duration. The key point about the cross-sectional distribution is that price-spell durations are weighted by their duration: a longer price spell is more likely to be observed at a point in time than a short one. Let us take an example from our dataset (which we describe more fully in section 3): Paxo sage and onion stuffing (170g) sold at Tesco had 9 price spells over the 365 weeks, with price-spell durations (106, 82, 5, 5, 1, 144, 2, 14, 6). The first price spell (106 weeks) was left censored and may have started before we started observing the price, but we shall ignore this here and treat it as uncensored (the last spell of 6 weeks is uncensored). If we treat each price-spell equally we have the distribution of $2/9$ for 5 weeks and $1/9$ for the rest, with a mean of 41 weeks. However, for the cross-sectional mean, we weight each spell by the probability that it is observed: i.e. the duration of the spell divided by the total sample length (365). Thus for example, the one period spell is highly unlikely to be observed: its weight will be $(1/365)$. In contrast, the longest spell is much more likely to be observed $(144/365)$. Using the probability of being observed as a weight, the cross-section yields the distribution in the third row of Table 1:

Table 1 about here

Unless all price-spells are the same length, the cross-sectional mean duration is larger than the mean duration: in our example the mean price-spell is 41 weeks, whilst the cross-section mean is 107 weeks.

This notion of using durations as weights was first put forward by Baharad and Eden (2004) and was shown to be equivalent to the cross-sectional distribution in Dixon (2009). The cross sectional distribution has been used to calibrate the Generalized Taylor economy by Dixon and Kara (2010) and Dixon and Le Bihan (2012). The two distributions (price-spell durations and the corresponding cross-section) are different ways of looking at the same data. Indeed, if we are in steady-state there is a simple identity that links the two distributions (Dixon 2009). We believe that the cross-sectional distribution is a reliable guide to answer the question of how flexible prices are: it is a distribution that is taken across prices at a point in time and hence relates directly to the behaviour of the price-setter (in our case the grocery store). The mean price-spell might be low due to a proliferation of short-price spells generated by only a few prices which are of relatively little economic importance. *Taking the cross-section means that we treat each price as equal (when weighted appropriately) and so do not bias our attention towards more flexible prices.*

If we compare the cross-sectional distribution to the filters, we can see that like NSB and KM13, we weight short-lived price changes less. However, the cross-section uses duration based weights across all durations and does not focus on a particular class of spells thought of as “sales”. Unlike the three filters, however, the cross-section is based solely on the duration of price-spells and does not attempt to identify a *reference* or *regular* price.

3. Our Raw Data Sample, Institutional Features, and Macroeconomic backdrop

We have collected, week-by-week, the prices at three stores for individual units of 370 precisely defined products over seven years from late 2003 to late 2010 for the three largest players in the British supermarket industry⁸. This data is a panel with 1110 rows (the prices of the 370 products at the three stores) and 366 columns (calendar weeks) with a total of 405,150 price observations.⁹ The data is for prices only and includes no data on sales as found in the scanner data used by Eichenbaum et al (2011). A number of features of this industry make the prices set extremely useful for examinations such as this.

First, there is the supermarket sector’s importance and concentration. Verdict Research (2008), a market research organization, estimates that in 2007, food and grocery retailing accounted for around 42% of total UK retail spending. All three of our firms, Tesco,

⁸A fuller description of the underlying data is available in Chakraborty et al. (2011).

⁹ For most purposes, we work with 365 weeks, since we examine changes.

Asda and Sainsbury, had been growing market share gradually over our sample period and together make the majority of grocery sales in the UK, according to Verdict research.¹⁰ These three are major retailing companies. Tesco, the leading chain is UK-based but has presence in several other countries and is by some measures the world's number three retailer. Asda is the UK subsidiary of Walmart. Another key feature is that all three firms have, since at least the start of our period, set prices nationally across all their larger stores (Competition Commission, 2008). It is important to understand that this means wherever one happens to shop within Britain, in the north of Scotland, the south of England or the west of Wales, one will face the same price in a large store of a particular chain (e.g. Tesco). For our data set, it is important to note that the price is the same as the price seen on the internet, available through a home delivery service.¹¹ Asda operates almost exclusively larger stores, whilst the other two fascias also operate smaller stores that do not adhere precisely to this national price policy and are therefore not necessarily covered by our series. But the large stores are the place where most people would do their weekly or fortnightly major shopping expedition.¹²

To what extent does this data, which is based on UK wide internet prices, reflect sales? Here we have the very helpful study undertaken by the Competition Commission [hereafter, CC] (2008) which looked into competition in grocery supply. For the stores in the data set, pricing is determined nationally and in particular sales (product promotions) are uniform across the UK. As the CC observed: *"In addition to pricing, substantial other parts of the retail offer for grocery retailers are also set nationally on a uniform, or near uniform, basis. Asda, Sainsbury's and Tesco all have centrally managed product promotions that run in all their stores (with some variation according to whether stores stock the product in question"*¹³. The exception to this is in terms of vouchering, where vouchers are given out by particular stores for specific items: this element of "sales" is of course not captured by posted prices at all: however, the CC found that "we find that the local vouchering activities of most grocery retailers are not extensive"¹⁴. Hence, although the dataset gives nationwide coverage, we can be confident that whilst there is some promotional activity in terms of vouchering, sales that occurred over this sample will mostly be reflected in our price data.

¹⁰Thus Tesco, with a 30% share, accounts for more than 1/8th of total UK current consumer expenditure!

¹¹ This makes the practice in Britain different from that in some of the countries that Cavallo (2012) observes.

¹² The weekly shop is most common, so is an appropriate frequency at which to observe prices. British consumers also "top up" at other stores (CC, 2008).

¹³CC, paragraph 6.31).

¹⁴CC, paragraph 5.86.

Our sample starts when Tesco, the largest chain, started its “Tesco Pricecheck” website in late 2003. This was an independently collected large scale weekly comparison of precisely defined products across these three store chains plus first Safeway, then later Morrisons (which took over most Safeway stores). We supplement this with data, from 2008 to late 2010, downloaded from a website called mysupermarket.co.uk (who collected across Tesco, Asda and Sainsbury’s) to create the seven year sample. Like the “Pricecheck” sample, this reflects in-store prices.¹⁵ Given the significant overlap in time between the two samples we have, the greater part of 2008, we are able to check for ourselves the high degree of concordance in the prices generated by the two approaches. Thus we have consistent data for Tesco, Asda and Sainsbury’s over seven years.

Our 370 precisely defined products are those for which we are able to form a good quality weekly price series over the full seven year period.¹⁶ Most are branded products (for example, Nescafe Gold Blend Coffee 200g), others are store brand products (e.g. Own label fresh single cream, 568ml) which are essentially identical across the chains. Given the approach we have adopted, our sample is heavily weighted towards packaged goods, not fresh items,¹⁷ and it is of course biased towards products that remain unchanged over the entire period and are consistently stocked by all three firms. No product substitutions were allowed. Within this framework, our sample covers mostly food and drink: of the 370 products, 269 (75%) are in the COICOP classification of *Food and Non-Alcoholic Beverages* (FNAB). In other categories are alcoholic drinks 31 products (8.4%), cleaning products (bleach etc) 27 (7.3%), Petfood 19 (5.1%), Soap, toothpaste etc, 14 (1.8%). These data are the prices we work with in examining the various questions that have been posed regarding price flexibility. They contain many low cost items, and the basic descriptive statistics are in Table 2: over the period product prices range from 15 pence to £30.

Table 2 about here

¹⁵The Tesco Pricecheck website was restructured in a way that made it less useful in the later years, leading to our choice to change the source after some overlap. Morrisons had no online presence over our sample period.

¹⁶The sample is clearly not random. However, appropriately weighted, it tracks the official CPI well (see Chakraborty et al, 2011 and section 5 below). As with all such studies (Nakamura, 2008), there are occasional gaps in the price series for individual items, the most important of which are prices over the Christmas period in the early years. We resolve these by filling in with the minimum changes in price possible (so that for example, if price is the same before and after a 2 week gap then we fill in with the same price). In any case, as is apparent from the definition of regular and reference prices, small gaps would be filled in on all KM13 regular and on reference prices anyway, as well as most NSB regular prices.

¹⁷Fresh items were essentially absent from the earlier periods of the Tesco Pricecheck sample.

Since the sample of goods is closely related to the COICOP classification *Food and Non-Alcoholic beverages* FNAB, it is useful to examine what the CPI data tell us. If we look at quarterly inflation rates, FNAB is about average in terms of variance with almost no persistence (Dixon, Franklin and Millard, 2013, tables A and B). If we look at the CPI monthly data for FNAB, we find that the average duration of a price-spell is 4 months, whilst the cross-sectional mean is 15 months: across all COICOP sectors, the equivalent values are 4 and 11 months respectively¹⁸. FNAB is a very heterogeneous sector, with 50% of price-spells lasting no longer than 1 month, whilst 2% of prices last more than 5 years. As we show below, we can use our data set to compute an alternative CPI index for the items covered (using the prices in our data set and the CPI weights): the correlations between various alternative CPI indices and the corresponding official CPI index are all very high (usually a correlation coefficient in excess of 0.99 and never lower than 0.94).

The study with a sample most similar to ours is probably that in Ellis (2009): his sample also relates to UK supermarket data and is of similar product dimensions. However, there are several clear points of difference. First, his data cover a three year period (Feb 2005 to Feb 2008), which excludes much of the crisis period covered by our data. Secondly, Ellis uses scanner data, relating to average selling prices, whereas we have actual standard item prices (i.e. our sample does not incorporate temporary discounts for multi-buy, price after coupon redemption, price for damaged or date-end items). His sample covers also fresh products such as raw vegetables and meat that do not feature in our data.

Before starting our investigation proper, we first reprise the macroeconomic backdrop which includes the development of the Great Recession. Two key features are charted in figure 1. The upper panel shows the series for UK GDP (centred on 2005). Observe that after a long steady slow rise during the Great Moderation there was a significant fall in 2008 almost back to the 2005 level followed by a partial recovery in 2009 and 2010 (faltering in 2011). We also see inflation, which was increasing through most of the period prior to 2008: there was then a rapid fall roughly coinciding with the output fall in 2008 into 2009, with inflation thereafter returning to the higher pre-crisis period level.

The lower panel shows IMF indices for food and beverage and for energy commodity prices, i.e. the world market prices for key inputs into groceries such as wheat, rice, meat,

¹⁸ See Dixon and Tian (2013) data appendix, based on CPI data 1996-2007.

orange juice, fuel and production energy inputs, etc. These, particularly energy prices which inevitably permeate the production and retailing costs of all grocery products, fluctuate somewhat but also experience a substantial upward trend from 2000 which accelerates rapidly in the early part of 2008 followed by a very sharp fall later in 2008 and a partial recovery in 2009. It is clear that 2008 and 2009 are very turbulent years and that our sample includes both a sharp upturn and a sharp downturn in activity and input costs. These are very useful features when examining pricing reactions at the micro level. Of course, most processed goods are rather complex combinations of ingredients and it would be difficult to sort out precise cost drivers for individual products, but these world cost trends are factors even the largest grocery retailers with significant bargaining power cannot avoid.

Figure 1: Key macroeconomic factors underlying our framework, about here

4. Exploring Price Flexibility in its Various Dimensions

As Klenow and Malin (2010) note, there are several dimensions to price flexibility. Most obvious and most studied is the frequency with which prices change. We characterize this in various ways using the range of price definitions we have discussed, but focusing on the raw data of posted prices. We then turn to upward and downward magnitude of price changes, to the cross-product timing synchronicity of price changes, then finally to features of the distribution of price changes. In the subsequent sections we draw out some implications.

4.1 Frequency of price changes, the duration of price-spells and the cross-sectional duration.

We have calculated the proportion of products which change posted price in each of the 365 weeks in our sample: there are 370 products across 3 stores making a total of 1,110 prices which can either change or not. If they do change, we can subdivide changes into changes up and changes down. The weekly frequencies are depicted in Fig 2. If we take the whole sample, we find that in an average week, 10% of products change their price each week, which decomposes into 4% changing price upwards, and 6% changing price downwards. However, as we can see from Figure 2, there is a considerable difference pre and post January 2008. We can think of the period prior to 2008 as reflecting the Great Moderation: the year 2007 is something of a transition, but clearly it has more in common with previous years than the crisis. In January 2008 GDP fell rapidly and clearly this influenced price-setting behaviour in our dataset.

Figure 2: The weekly frequency of posted price changes, about here.

If we divide our data into pre-2008 and January 2008 onwards, the weekly mean frequency of posted-price changes more than triples from 5.4% to 16.8%, with price cuts increasing from 2.9% to 11.0%, price increases from 2.4% to 5.8%. This represents a significant increase, particularly in the frequency of price cuts.

The reciprocal of the weekly frequency gives us an estimate of the expected duration of a price spell related to when it starts: for posted prices pre-crisis this was almost 19 weeks, post crisis it fell to just 6 weeks, as we report in table 3. If we compare our data with the monthly CPI data for FNAB, our moderation figure is quite close to the Dixon and Tian (2012) estimate of 4 months pre-2007: the difference may reflect both the different frequency of observation and the fact that our sample excludes fresh items whose prices change more often. We calculated expected duration for NSB, KM13 and EJR prices in the same manner and append these to table 3.

From our data set we can also compute the actual mean posted price-spell directly from the population. There are some issues of censoring. Whilst censoring is a relatively small problem in our dataset, all of the first price-spells for our 1,110 prices are left censored: when we started observing our 1110 prices, we did not know for how long the price observed had already been in place. Likewise, at the end of the sample, unless the price happened to change in the last week of 17th November 2010 (which was true for 165 prices) we have right censored spells: we do not know how much longer they lasted. There are some very long spells in our data. For example, Asda's own label Dairy Cream Slices: we observe only one price change in week 301 (19th August 2009). We thus have two censored price-spells: a left censored spell that lasts 301 weeks and a right censored spell that lasts 64 weeks. Out of our 1112 prices, 245 had at least one price-spell in excess of 100 weeks, and 67 products a spell over 150 weeks. We deal with censoring by including left censored spells (i.e. assuming all left censored spells started in week 1) but excluding right censored spells. This is a compromise between excluding all censored spells (which tends to be biased against longer durations which are more likely to be censored) and treating all censored spells as complete (which will bias downwards durations). Dividing the data into two periods creates further issues of "between sub-period" censoring: how do we allocate price-spells that span both sides of January 2008? Again, this is particularly likely to affect longer price spells. We believe that there is no exact or correct answer to this issue and adopt the simple procedure of

allocating a price-spell to the sub-period in which it ended. For example, this places the 301 week Asda Dairy Cream Slice in the crisis sub-period.

Table 3 about here

We present the different measures of duration in Table 3. In the first row we see that the mean duration of price-spells estimated from the frequency falls from 18.6 weeks in the moderation to 5.8 weeks during the crisis. Even though the crisis part of our sample is shorter than the moderation, since spells were substantially shorter in the crisis, it contains many more spells (27,713) than the moderation (12,440). Turning to the measured mean price spell in row 2, the large discrepancy between the crisis frequency estimate of 5.8 weeks and the measured mean of 9.6 weeks is almost entirely accounted for by the fact that 4% of price spells *ending* in the crisis period are over 100 weeks long (as opposed to 1% being over 100 weeks for the whole sample and 2% in the moderation). However, for the whole period and the moderation the estimated and measured mean are quite close, the differences reflecting censoring. The cross-section mean for the moderation is 48 weeks: this is quite close to the FNAB estimate of 14 months in Dixon and Tian (2012). Again the allocation of long-price spells makes a difference here: the fact that so many are allocated to the crisis will reduce the measured cross-section mean for the moderation. We also calculate the median price-spell: because of the fat tail of long-durations, the median is substantially less than the mean. There is a big drop in the median duration if we compare the two periods: from 7.5 (moderation) to 2.5 (crisis).

However we measure it, there is clearly a huge drop in the mean duration of posted price-spells. All three measures of the mean fall by about one-half when we compare the moderation with the crisis. The drop in the median is even more substantial. A similar story is told by the three filtered measures, which each show a similarly big difference between the periods, although the filters of course smooth away many of the price changes in the posted price data and therefore lengthen all the spells; hence our stylised facts 1 and 2.

Figure 3: The distribution of price-spell durations pre and post crisis.

Whilst the mean is a useful statistic, it is also informative to look at the whole distribution of durations. Firstly, in the upper panel of figure 3 we look at the distribution of posted price-spell durations in the two sub-periods, the moderation and the recession. On the horizontal axis is the duration in weeks (up to week 50) and the vertical axis gives the proportion of

price-spells lasting a particular number of weeks. We can see a dramatic change between the two distributions. Most noticeably, the crisis period sees a massive increase in one to three week durations. One week becomes the modal duration, with 23% of price-spells lasting just one week and 19% two weeks. In the moderation, just 8% of spells last one week and the modal duration was 4 weeks (9%). The two distributions are quite close for 4-6 weeks, but for durations of 7 weeks and above the crisis has significantly lower mass than the moderation. Thus in terms of the distribution of posted price spells, the crisis has the effect of moving much of the mass of the fat tail of long durations found in the moderation to durations lasting 1-3 weeks. This shows the importance in understanding the phenomena of having weekly data available- monthly data would fail to capture much of this movement. This point is brought into even sharper focus in the lower panel of figure 3, where we look at the distribution of NSB price-spells which if anything show a starker difference between periods.

Another way of looking at the frequency of price change is by looking at the number of products that experience a particular number of price changes in a given calendar year as described in Table 4 for NSB prices across the three stores. There are two significant changes when we look across the years. First, we see that in the years 2004-2006 there is a large proportion of prices which do not change in a given year: in 2004, we find that in Asda 37% of products experienced no NSB price change and likewise in Tesco 32%. The lowest proportion not changing was 18% in 2005 at Tesco, and the average over 2004-7 and all three stores is 25%. In contrast, in 2007-2010 there are few prices that remain unchanged in any given calendar year. On average across all 4 years and 3 stores only 6% of NSB prices remain unchanged in a given calendar year. The year 2008 stands out particularly, as we might expect: in Tesco only two NSB prices remained unchanged.

Table 4: Average number of products experiencing the number of price changes in the ranges shown, by supermarket and year, NSB prices.

Figure 4: Count of numbers of products (vertical axis) experiencing the number of price changes indicated per year, NSB prices

The second big change is the number of *price changes per product* which shows a remarkable increase between the moderation years and the crisis years 2008-2010, with 2007 as something of a transition, as we see in figure 4. For all years 2004-7, the modal number of price changes per year is the category 1-4 changes across all three stores: however, for 2008 the modal number of price changes is 13 or more for both Asda and Tesco. There is a huge

increase in the proportion of products changing NSB prices 5 or more times per year: in all three supermarkets, less than 4% of products changed price more 5 or more times in 2004. However, in Tesco and Asda, this increased to over 50% for all three crisis years; the increase in Sainsbury is less, but is still very high – with over 30% in 2008 to 9 and 25% in 2010 changing price 5 or more times.

This is a new finding in our sample: individual NSB regular prices across our 370 products and 365 weeks are far from sticky. If we examine median duration of NSB regular prices, as in figure 5, we see that for around 10% of Tesco and Asda products, median NSB price duration is only two (three) weeks! Beyond that, median durations increase and some products, particularly milk, stay more or less fixed over the whole period. But still, for around half the products in our sample, median NSB duration is six weeks or less in Tesco and Asda. *Hence it is not the case that rapid price movements are confined to a small subset of products, or to posted prices alone.* NSB prices in Sainsbury's are markedly less flexible, with half the products having median duration longer than 10 weeks on this definition, probably reflecting the fact that they have engaged less directly in the price promotion strategies of their rivals. These strategies commonly consisted of assertions regarding the number of product prices that had been reduced or the number of products cheaper in one chain rather than in a rival chain.¹⁹

Figure 5: Distribution of median duration of NSB regular prices across our sample of products

Of the other definitions of prices that we adopt, it is natural to consider also the EJR reference prices, since in principle these are the least likely to be flexible. Nevertheless we find even these become very flexible in practice during the crisis period. Based upon Eichenbaum et al.'s (2011) definition of a reference price as the modal price in a quarter, we can examine the behaviour of these in our sample across the 366 weeks of data at our disposal. We develop three slight variants using the basic definition: we use (i) our last 364 observations to examine 28 quarters (this is closest to Eichenbaum et al's 2009 approach as described), (ii) data based on "calendar quarters" commencing January 2004 and ending in

¹⁹ Consider the following snippet from Wikipedia, commenting on supermarket price competition: "...out of 7134 (compared to Asda) products, (Survey carried out between 9 July 2007 and 11 July 2007) Tesco is cheaper: 1835 (compared to 1251 the previous week), Tesco is more expensive: 975 (compared to 984 the previous week) and Tesco is the same price: 4324 (compared to 4996 the previous week)."

the third quarter of 2010 (27 quarters), (iii) data based on constructing reference prices after creating NSB regular prices (27 quarters). The results vary only very slightly as between these variants

Overall, EJR reference prices in our sample change far more frequently than annually, a point of considerable distinction relative to previous findings (Klenow and Malin 2010, fact 2). In fact, using our first definition, we find that 90% of our 370 products across the three firms change reference price more than seven times in our period. The mean number of reference price changes is 12.3 over seven years, with Sainsbury's products at 11.2 times and the other two just slightly less than 13. Since the maximum number of price changes on this methodology is 27, the average product reference price changes approximately every six months. This is far more frequent than Eichenbaum et al (2011) find to be the case.

Figure 6: Distribution of EJR reference price changes.

As we found with NSB prices, we find that EJR reference prices respond to the crisis and are much shorter than has been seen in previous work based on the moderation: it is 2008 and 2009 where reference prices change most often - across our sample in both these years, reference price changes on around 63.2% of the possible occasions across quarters.²⁰

Figure 6 contrasts the number of EJR reference price changes in 2005, a year of modest global price movements, with the numbers for 2008 and 2009, where these movements were much starker. The median number (one) of price changes across the four quarters for our products in 2005 is in line with Eichenbaum et al. (2011). However it contrasts sharply with the median of three in 2008 and 2009 and the differences in distributions are obvious without need for statistical test.²¹

²⁰The approach in this and the following paragraph involves definition (ii) above, since it relates to calendar quarters. The quarter to quarter change is measured (using 2005 as an example, as [(Q1 of 2005 – Q4 of 2004), (Q2 of 2005 – Q1 of 2005), (Q3 of 2005 – Q2 of 2005), (Q4 of 2005 – Q3 of 2005)]. Hence 2004 and 2010 observations relate only to three quarters each. We also note parenthetically that with definition (ii), it happens that for three goods for one firm, the modal price (i.e. Reference price) was undefined in one quarter, because the posted price changed literally every week within that quarter!

²¹For completeness, we also carry out the same calculations for KM13 (smoothed) regular prices. The overall pattern across chains is similar, with Asda and Tesco having a very similar distribution of product durations, and Sainsbury's somewhat above this. However, because the definition smoothes to a greater extent than NSB prices it excludes more movement than they do, so the distribution is markedly above that for NSB regular prices. Nevertheless, these values still imply prices, under the KM13 regular definition, that are on average quite flexible, significantly more so than has been found by Kehoe and Midrigan (2010) in their sample. See figure A4 in Appendix 2 for details.

In sum the previous US and Euro area findings of frequency being little affected by macroeconomic variables is contradicted by this study: the crisis has led to a substantial decrease in the duration of price-spells however they are measured and corresponding significant increase in the frequency of price changes. This appears to back Klenow and Malin's (2010) suggestion that the apparent stability of the frequency results from the fact that previous studies have been based on cross-section data from the moderation period. As in Gagnon's study of Mexico under high inflation (Gagnon 2009), we find that in more turbulent periods, the frequency of underlying price changes does increase, hence contradicting what has been viewed as a key empirical regularity regarding frequency. This is also consistent with theoretical menu-cost models of pricing (Sheshinski and Weiss 1977). What is significantly new about our results is that inflation appears not to be the key driver: arguably it is the crisis itself that gives rise to the big change in pricing behaviour, as we point out in section 7.3. What distinguishes the crisis period in the UK is the sustained drop in output. In the US and the Eurozone there was less inflation during the crisis period, yet we would expect similar changes in pricing behaviour to those found in our study.

5 The behaviour of price changes

5.1 Magnitude of price changes

It has been widely observed that the magnitudes of micro price changes generally exceed the change in aggregate inflation (Klenow and Malin's, 2010, "sixth fact"). On the one hand, if some prices are sticky and do not change, simple arithmetic dictates that since inflation is determined by the prices that do change, the changes must be larger than inflation. Also, the aggregate CPI inflation is a weighted mean of the different sectors in the CPI basket. Hence the variance of aggregate CPI inflation is a pooled variance of sectoral inflation variances. We would expect the prices in our sample to be more volatile than is reflected in the smooth behaviour of aggregate CPI. We indeed find that the magnitudes of price changes are large, although here our data reveal several unexpected surprises. In this section we again focus initial discussion on NSB regular prices, but also examine posted prices.

In line with previous studies, the average size of individual product price changes, both upwards and downwards, is large compared to inflation. In terms of NSB regular prices, the average rise (for goods where a change occurs) is over 10% of the item price, whereas price falls average just under 10% of item price. At the same time, by no means all products

experience price changes- around 41% of products experience no price falls in any given year, whilst even in the most inflationary years, there are some products that experience no price rise.

In Table 5 and 6 we show the average magnitude of price rises and falls for NSB and posted prices respectively. The magnitudes of NSB price increases can be seen to rise significantly in the years 2008-9: for example, Asda hovers around 11.5% in 2004-6, and rises to over 16% from 2008 onwards with the peak at over 17% in 2008 itself. Much the same happens for Tesco, but for Sainsbury's the increase is limited to 2008 (peaks at 13.5%) and by 2009 it is back to its 2004-6 value below 10%. The magnitude of NSB price falls is rather more stable: there is a clear downward trend in the absolute size of price cuts, but not such a pronounced 2008 effect (the 2004-7 mean is -10%, the 2008-10 mean is -8.2%). Since the NSB regular prices exclude temporary sales that revert to the "regular" price, the increase in the size of price increases cannot be due to recovery from large temporary price-cuts.

Table 5: Average percentage NSB price change across our products by year and chain.

Table 6: Average percentage posted price change across our products by year and chain.

In interpreting these data, we should also remember that many of our products are relatively low price consumer products. In fact, for 45% of our products the average price is one pound or less, and for 17%, 50p or less, hence the lowest possible change of 1p would constitute at least a 1-2% price change.

The corresponding data for posted prices is given in Table 6. The magnitude of average price rises is around 50% larger than for NSB prices, reflecting temporary price-cuts that are reversed, but exactly the same pattern occurs, if not as dramatic. The average posted price increase trends upwards from 14.6% in 2004 to a peak of 23.3% in 2008, after which it has fallen back slightly to 21.1% in 2010. There is some heterogeneity across stores: Tesco and Asda peak at 22-24% in 2009. Sainsbury's is again different: 2004 is a high value, it does peak at 2008 (24%) and then falls off rapidly. Posted price falls are remarkably stable in their magnitude, being between 8-10% and not showing any clear trend. Thus posted prices tell the same broad story as NSB prices. Since some falls are removed in the process of

generating NSB prices, this has a knock-on effect of removing the following rises when we filter out the V-shaped movements to obtain the NSB prices.

Overall, we can see that the crisis brought about a change in the magnitude of price increases. The size of NSB price increases rapidly increased through 2007 to a peak at 2008 and has remained at a substantially higher plateau through to 2010. The path of posted prices was smoother, with less of a 2008 peak. Hence we can see that not only was the frequency of price increases influenced by inflation and output, but also the size of increases. When we turn to price decreases, they are much more stable in magnitude: whilst there is a slight trend downwards in the size of NSB price-falls, it is in effect constant for posted price-falls.

5.2 Synchronization of price changes

In the analysis of prices at the micro level, another dimension that has attracted the interest of macroeconomists is the synchronization of these changes. The issue is analogous for example to that in examining wage setting to see whether it takes place at particular times of year. Other researchers have observed that this bunching phenomenon is not true of micro price changes (Klenow and Malin, 2010) and our findings match this, though with some novel features.

Figure 7: Number of weeks in which there are NSB regular price changes in the ranges shown

Figure 7 examines the distribution of NSB regular price changes across the weeks of our sample. It shows a rather marked divergence between price rises and price falls: price rises are spread widely across weeks, whilst falls are bunched somewhat more towards the upper end of the distribution compared with rises. There are literally only two weeks in the 365 in which there has not been at least one NSB regular price fall by one chain compared with the previous week. Most weeks see a flurry of price falls, with well over 10% of weeks seeing one hundred or more falls across the chains over the week, peaking at almost 1/3 of prices falling in one particular week, a remarkable downward degree of fluidity in prices, given that these are NSB regular not posted prices.

Even when we move to the substantially smoother KM13 regular prices, there are still two weeks with over 100 price falls overall, and there is still a slight bunching towards the right end of the distribution when compared with price rises, as Appendix table A1 illustrates. But the overall message is that price movements are well dispersed throughout our period.

5.3 Cuts and Hikes.

Drawing the previous two findings together, price rises are relatively large in magnitude, both compared with inflation and with price falls, whilst there are many weeks with a large number of price falls. This leads to questions on the nature, in particular, of the many price falls observed and concerning the overall net impact of rises and falls.

Here we come to the most remarkable finding. In our sample, the distribution of the sizes of price changes, particularly price falls, is unusually asymmetric, even more so than has been observed in previous studies such as Klenow and Kryvtsov (2008) and Midrigan (2011).

Figure 8 shows this graphically for NSB regular prices. We see the remarkable fact that there are, throughout our period, more NSB regular prices falling than rising in both Asda and Tesco, and indeed for most of the time at Sainsbury also. The finding of significant numbers of price falls is not new; for the UK Bunn and Ellis (2012) observe that around 40% of price movements are price falls in a period of general mild inflation. What is new in our dataset is the finding that price falls numerically exceed price rises on average throughout our period is novel. In particular in 2008 and 2009, there was a particularly dramatic excess, with up to around three times as many price falls as price rises. This is remarkable since it includes a period (2009 in particular) where inflation was rapid and cost changes were extensive.

Figure 8: The ratio of NSB regular price falls to price rises across the firms and time

The picture is just as dramatic when we look at posted prices. Returning briefly to Figure 2, we can see that the red line (representing frequency of price cuts) is well above the green line (proportion of price rises), particularly from January 2008. The overall picture for posted prices in figure 9 is similar to that in figure 8 for NSB prices. The additional feature that comes out very clearly is the sheer numbers of posted price falls observed that are penny price cuts. In Table 7 we see that around 1/3 of all Tesco and Asda price cuts in 2009 are penny cuts, as are around a quarter in 2008.²² Indeed, penny price cuts constitute one sixth of all price movements whether up or down in 2008 and a remarkable 23% of all price

²²This definition of “small” price cuts is not in line with previous work such as Klenow and Kryvtsov (2008) which uses percentages, but given the low unit price of many of our products, it seems sensible to take the smallest possible unit amount of change as an indicator, rather than a small percentage.

movements in 2009.²³ In addition to penny price cuts, there is also a large, though lesser, number of price cuts of 2p in the data, as table A2 in Appendix 2 shows.

Figure 9: The ratio of posted price falls to price rises across the firms and time

and Table 7 about here

Here we focused on posted prices, since these reveal the magnitude of individual price changes most clearly, and in particular highlight small price changes. Hence it is crucial to address the recent Eichenbaum et al. (2012) criticism regarding small price changes, specifically whether our data source is subject to this criticisms they have regarding identifying such price changes in data they and others commonly use. Since our data is of posted prices for single unit purchases, it is not subject to the Eichenbaum et al criticism. They refer to two source types, of which only one is potentially relevant to their critique, scanner data. As they point out, scanner data has the limitation that “price” is often based upon unit value indexes and so incorporates a potentially large number of promotional factors (coupons etc.).²⁴ Our data approximate most closely to the supplementary data set they have for 374 stores across four US states in 2004. These appear from their methodology to be the set of prices from which they would be most confident in drawing conclusions on small price changes. For completeness, we also note that there are no local sales taxes in the UK and that in fact, the majority of the goods in our sample do not attract sales taxes of any sort;²⁵ for the minority that does, the tax change is once per year. Tax differences are not giving rise to price effects in our data.

In sum, we appear to have uncovered a degree of price lowering of a different order of magnitude to that previous studies have found, with well over twice as many price falls as rises in certain periods. In fact, the overall numbers are instructive, in that they highlight the overwhelming dominance of falls. We observe 19,571 falls and 10,358 rises in NSB prices across the whole sample. There are 24,891 posted price falls and only 15,262 rises across the whole dataset. The difference is even more marked from 2008 onwards, with 18,104 price

²³ This welter of small value price cuts is the focus of our companion paper, Chakraborty et al. (2011)

²⁴ By the same token, our data are poor at catching promotions.

²⁵ Ambient and frozen food is zero rated for VAT in the UK, as are soft drinks. Alcoholic drinks are taxed, and the rate change usually takes place in April and are always increases in the period considered.

falls and 9,609 rises.²⁶ Prior to the crisis, the number of falls was also higher than the number of increases, with 6,920 falls against 5,748 rises.

Stylised Fact 3: There are more price-cuts than price hikes in the sample, for both NSB and posted prices. The ratio of price cuts to price-hikes increases significantly during the crisis.

The obvious question is why it happened that so many price falls took place in the period and across the retailers covered by our data in a period when, on balance, costs and prices overall were rising, as was as basket prices within our data (see section 5.2 below). We reiterate here that these are the major firms in their market, not some niche players. However, they are also in a situation of close oligopolistic rivalry. This is a completely different market structure to the Dixit-Stiglitz version of Chamberlin-Robinson monopolistic competition assumed by most current macroeconomic models. Over parts of the period we are examining, this oligopolistic rivalry commonly took the form of claiming that in a particular chain (be it Tesco or Asda), more than x hundred prices fell over a particular week, or more than y hundred prices were cheaper at one rather than the other. To achieve this, small value price cuts clearly became a core method of competition.²⁷ This became particularly prominent in the crisis period when consumers were stretched in their budgets and focussing on value for money. In effect, one can think of this phenomenon as a form of “epsilon-undercutting” by the retailers, as in Bertrand competition where prices are restricted to integer values in terms of pence. Whilst of course we do not envisage customers switching stores to buy the cheapest biscuits as in the classic homogenous good Bertrand model, the grocery stores can use the fact that they are the cheapest across a range of products as part of their marketing strategy to indicate value for money. The least expensive way to undercut your rival is to be a penny cheaper. If your competitor is undercutting you by one penny, the best way to undercut them is to lower your price by 2p. There is thus a sort of “Edgeworth cycle” set up as the grocery stores compete to show that they are cheaper across a range of products.

Another possible explanation of the proliferation of small price-cuts is also based on the argument of Bennett and LaManna (2001). They argue against the “Keynesian” convention that prices are more rigid downwards than upwards: in fact they argue to

²⁶ It might be queried why the differences between rise and fall numbers are not exactly the same across NSB regular and posted prices. The answer is that sometimes, the V shapes excluded in the NSB algorithm incorporate two different lower prices before the return to the same price.

²⁷ See footnote 19 above.

opposite, that this “Keynesian asymmetry” be reversed. Their argument focuses on the market structure, adopting a Bertrand setting with free entry. However, since this argument relies on free-entry in a form of a contestable market a la Baumol et al (1982), it is perhaps not applicable to the large grocery stores we are considering.

Figure 10: The dispersion of NSB price growth including and excluding small price changes.

5.4 The Distribution of price-growth

If we look at the overall distribution of price growth (excluding the zeros), we find that the higher moments reflect the proliferation of small price changes in the crisis period. Turning first to the dispersion of price growth, we find that the two classic measures diverge if we include small price changes. In Figure 10a, we show the smoothed²⁸ weekly Standard Deviation (SD) and Interquartile range (IQR). Until 2008, they move together. From 2008 onwards, the IQR drops significantly below the SD. This drop is entirely due to the small price cuts in the data post 2008. In Figure 10b, we show the same measures excluding not only zero price changes, but also all absolute price changes of 2p or less. The IQR and SD are now similar throughout the period. There is considerable short-run fluctuation in the dispersion of prices as measured by the SD, but it is stationary around a slightly rising trend. There is no clear crisis effect. The IQR tells the same story if we exclude the small absolute price changes. Note that defining small price changes in terms proportionate to price is very different from defining it in *absolute* terms as we have done. Many of the goods in our sample are cheap. If a good costs 50p, then a 2p price change represents a 4% change. The fact that price is not a continuous variable matters for understanding price competition in the data we are considering.

Vavra (2013) finds that the dispersion of price-growth is counter-cyclical in the US CPI data he covers²⁹. This would imply that we would expect dispersion to increase when there is a fall in output as occurred in the crisis. We do not find little evidence of this in our data. Pre-crisis the SD of price growth had a weekly mean of 0.14 and during the crisis 0.16, a very modest and insubstantial increase given the behaviour of output. The IQR fell from 0.13 to 0.9. This suggests that Vavra’s empirical findings for the US may not be general.

²⁸ A 7 week centred moving average is used for smoothing unless otherwise stated.

²⁹ “The cross-sectional standard deviation of price changes is strongly countercyclical: price changes become substantially more disperse during recessions.” Vavra (2013).

Vavra also finds that “*The standard deviation of price changes comoves strongly with the frequency of price adjustment in the economy*”. Again this is not found in our data. Whilst the frequency of price-adjustment rose substantially during the crisis (Stylized fact 2), this was unaccompanied by a similar significant increase in the dispersion of price-growth. If we include small price changes, the exact opposite happened.

Skewness changes from being close to zero (-0.05) pre-crisis to being positive (0.27), which again reflects the increase in small price cuts. Excluding price cuts of 2p or less leads to a pre-crisis value of -0.16 and a crisis value of 0.02: a small increase in small numbers. By Bulmer’s rule of thumb, any measure of skewness between -0.5 and 0.5 is “approximately symmetric”: hence skewness is not significant in our data³⁰. This indicates that the overall distribution of price-growth is fairly symmetric, despite the proliferation of price-cuts.

Figure 11: The excess Kurtosis of NSB price growth, including and excluding small price changes.

If we turn to estimated excess Kurtosis, we find a large increase in the crisis period if we include all price changes in a given week³¹. The average weekly Kurtosis increases from 2.7 pre-crisis to 7.6 post crisis³². There is an increase in the mass of price changes close to zero. However, if we exclude price changes of 1p and 2p, we find that this increase is much reduced especially in the post-crisis period - the average weekly Kurtosis pre and post crisis are 1.9 and 3.2 respectively. The increase in Kurtosis we observe in the crisis is thus only partly explained by the proliferation in small price cuts that we have already described. The increase in Kurtosis of price-growth during the crisis is the *opposite* of what Vavra found in the US data: he found the level of Kurtosis to be pro-cyclical, so that you would expect Kurtosis to fall in a recession. Thus, within our data we find that the distribution of price-growth is leptokurtic as compared to the normal distribution (which has excess Kurtosis of 0). This implies that there are heavier tails and a higher peak than with the normal distribution.

³⁰ Smoothed skewness is depicted in the appendix figure A5. Bulmer (1979) indicates that Skewness over 1 in absolute value is highly skewed, between 1 and 0.5 “moderately skewed”. From figure A5, whilst there is some evidence of moderate and high skewness in the crisis period this arises from the small price cuts: when these are excluded, skewness remains almost always in the “approximately symmetric” zone. However, when looking at smoothed weekly values the small sample test would indicate even larger deviations from zero are needed to indicate skewness (Bulmer’s rule of thumb applies to population data). The values in the data do not indicate any significant skewness and asymmetry.

³¹ There are on average 83 changes in NSB prices per week across the whole sample. We use the sample estimate provided by the KURT operator in Excel for each week and smooth across 7 weeks.

³² As an alternative calculation, if we take Kurtosis across all weeks and prices pre-crisis the figure is 7.2 and post-crisis 8.0.

Note that we have not standardised the price changes, as in Alvarez et al (2013). They standardise price changes by expressing the price growth relative to the mean and standard deviation of a class of items over the whole period. Our Kurtosis measure is taken as a weekly cross-section across all products and retailers in our data. There are two arguments against standardizing the price growth distribution in our dataset as suggested by Alvarez et al (2013). Firstly, the practical one of sample size: on average, only 83 NSB prices change each week. If we look at smaller subsets, or even individual items, then there are simply too few observations to obtain a meaningful distribution. However, the second reason is that the data is telling us that the behaviour changes due to the crisis. It makes no sense to assume that there is a stable distribution across the period with which to standardise. We could of course split our sample and standardise over the two periods separately: however, this would then run even more into the small numbers problem, particularly in the moderation period when prices changed much less frequently.

This leaves us open to the criticism of Alvarez et al (2013) that it is misleading to aggregate over heterogeneous distributions and can lead to excessive Kurtosis. However, we would argue that our data set is far more homogeneous than the French CPI data set they are using. The un-standardized Kurtosis levels we find are much lower than the standardized one in the French CPI data³³. This is partly because the outlets are very similar (the CPI data will include small stores and a range of outlets) and the products much less diverse than the CPI basket. Lastly, we would argue that there is considerable heterogeneity over time of Kurtosis. Standardising over time will lead to excessive Kurtosis for much the same reason as aggregating over cross-sectional data. Where we do agree strongly with Alvarez et al is in noting how sensitive the measurement of kurtosis is to relatively small changes in definitions, in our case the de minimis rule one chooses to adopt.

We summarise the findings of this section as

Stylised Fact 4. The dispersion of price growth as measured by the Standard Deviation does not change significantly as a result of the crisis: dispersion as measured by the IQR falls if we include all price changes. Skewness is insignificant throughout the whole period. The

³³ For example, Alvarez et al (2013), Table 1 states that Kurtosis for the non-standardised French data excluding sales is 20.9. After standardisation it is reduced to 10.4. In our sample, without standardisation, the weekly average for NSB prices over the whole period of excess Kurtosis is 4.7 (equivalent to Alvarez et al Kurtosis 7.7).

distribution of price growth is leptokurtic and became more so in the crisis. Only part of this increase is excess Kurtosis explained by the proliferation of small price-cuts.

6. Price Dispersion.

In this section we examine the behaviour of price level dispersion in terms primarily of NSB prices and to a lesser extent posted prices. We measure this using the coefficient of variation CV, the standard deviation of prices divided by the mean. We normalise the standard deviation since there is significant inflation during the period considered and the standard deviation is an absolute measure that will tend to increase if all prices are drifting up. Dividing by the mean price effectively corrects for the effects of general inflation.

The path of the CV for posted prices across all 1,110 prices is depicted in Figure 12. The global CV provides a general description of what is happening over time.

Figure 12. Posted price dispersion as measured by the all-price CV

The time-series for the all-price CV does not show any great break in 2008. Indeed, the main variations are seasonal and occur in the run up to Christmas and the New Year when there is a considerable drop in the all-price CV and a smaller drop around June. These represent times when there are seasonal sales of certain goods (for example alcohol, barbeque items). Visually, there is a slight downward trend in the CV, but nothing dramatic. The average CV over the whole sample is 1.67: for the moderation it is 1.68, and for the crisis period it is 1.65.

This result may appear puzzling. However, in order to understand the behaviour of price-dispersion better, we need to disaggregate the prices. Our first step is to remove the effects of temporary sales by using the NSB prices. Secondly, we look at prices at two levels of disaggregation. At its finest, we take each product and look at the dispersion of the three prices across sellers. As an intermediate, we divide the 370 products into 12 CPI sub-categories and measure price-dispersion within each category. The all-price CV for NSB prices behaves pretty much like the all-price CV for posted prices: the average across the whole period is 1.68. Hence removing sales does not affect all-price CV by much. Most of the variation in prices is due to variation across products, so that there is a big reduction when we disaggregate: the mean CV across the CPI categories is 0.61, and when we look at the product level it is only 0.04. Hence, 97.6% of total price-dispersion is between products, and 2.4% is within products.

Figure 13: The normalised CV for NSB prices at different levels of aggregation.

In Figure 13, we show how the CV behaves at these different levels normalised relative to its mean (taken over the whole period). Here we see that when we look at the all-price CV and the CPI groups, there is little variation around their mean. Most of the variation is generated by inter-product differentials which are stable (except for seasonal sales, particularly in alcohol, which includes many very expensive items such as large bottles of spirits). However, when we look at the product level, the CV is much more variable. The CV in 2005-6 is less than 80% of its mean value. It increases throughout 2007 to a value 2008 and beyond of over 120% of its mean. This represents a growth in the product level CV of over one half.

Figure 14: Relative NSB prices across the three retailers (Asda the numeraire).

Since there are only three prices per product, we can get an understanding of what leads to the increase in CV over this period by comparing the prices across the three retailers. A higher CV means in effect that the spread of prices between the three retailers has increased. However, closer inspection reveals a particular story. In 2007 onwards there was a price-war between Tesco and Asda. If we take Asda as the numeraire, in Figure 14 we show the average % price differential across all 370 products relative to Sainsbury and Tesco. Sainsbury is generally the most expensive, except for a brief episode over the Christmas/new-year period 2009-10 when all three retailers were an average about the same. Tesco was slightly more expensive prior to 2007. However, from 2007 onwards there are many weeks when Tesco is cheaper than Asda (the red line lies below zero). In figure 15a we depict the number of products for which Asda is cheaper than Tesco and vice-versa. In the early years of the sample, the vast majority of products had exactly the same price at the two outlets. However, from 2007 onwards we find a significant increase in the number of products which are cheaper in one of the two outlets. Whilst most of the time Asda has many more products which are cheaper than Tesco does, in the first half of 2008 the intense price-competition is reflected in the two lines being close together. It is important to note that these are NSB prices: we are not seeing the effect of temporary sales, but of changes in *regular* prices. In Figure 15b, we show the importance of small price-differentials. We depict the number of products in Asda and Tesco which have identical prices and which have an absolute difference of 1p and 2p. We can see that there is a reduction in the number of products for which prices are identical: on average over the period prior to June 2007 over 300 products

were identical, with the average in 2010 being less than 200. However, there was also a big increase in prices which were within 1p or 2p across the two stores. Prior to 2008, only a handful of products were within this range. However, in 2009-2010, 50 products were in this range. In mid-2008, around 100 products were within this range, indicating intense price-rivalry.

Figure 15: Price competition between Asda and Tesco (NSB prices).

The story of price-dispersion depends on what level of disaggregation is used. All-price dispersion and dispersion within CPI categories are fairly stable over time, being dominated by dispersion across products and there being no evidence of an effect of the crisis. However, if we look at the product level there is a different story. There is a significant increase in price-dispersion, reflecting the increased price-competition between the three retailers, in particular Tesco and Asda. Enhanced price rivalry took the form of the Tesco and Asda reducing the number of goods for which they charged the same price, resulting in more price differentials and increasing the price-dispersion as measured by the product-level CV.

Stylised fact 5. Looking across all prices, there was no significant change in price dispersion over the period. The variations in all-price CV reflect seasonal affects as certain groups of products were discounted (e.g. Alcohol). However, if we dis-aggregate to the product level, we see that there was a significant increase in price-dispersion from 2007 onwards, reflecting increased competition between the three retailers across a range of products.

7. Implications of our Findings

7.1 Menu Costs and state dependent pricing.

The implications of the finding of this paper for models of nominal rigidity are that given a big shock, pricing behaviour is clearly state-dependent, or at least contains a state-dependent element. What are the implications of our model for models of pricing used in macroeconomic models?

The standard model of menu costs as a framework for explaining price-rigidity was initially developed in the context of monopolistic competition (Barro 1972, Sheshinski and Weiss 1977, Akerlof and Yellen 1985, Mankiw 1985). The essence of the Chamberlin-Robinson model of monopolistic competition is that there is no direct interaction between

firms in terms of prices or demand: the individual firm's demand is determined solely by the aggregate price index, aggregate demand and its own price. In the Dixit-Stiglitz version of this model with CES preferences, the demand for the good of monopolist i depends solely on the price of monopolist i relative to the CES price-index and aggregate expenditure. This can be thought of as a world in which each monopoly is small and limited and there is no local interaction between monopolists. Adding menu-costs to the monopolistic setting implies that the decision to adjust price is not affected directly by the prices of competitors, but only indirectly through the aggregate price index. Since competitors are effectively atomistic, no one competitor has a discernable effect on the aggregate price.

This leaves out two important dimensions: first oligopolistic interaction, whereby the price of a specific competitor matters, and second, the multiproduct dimension that firms are competing across a broad range of products (at least 370 in our case).

Current research on menu-costs has developed the theory somewhat in these directions. For example, Alvarez and Lippi (2013) consider menu-costs in the context of a multi-product monopolist, but do not capture the "oligopoly" or the "bundling" element in the story. Bennett and La Manna (2001) develop the oligopoly side of the menu-cost model, but not in a fully dynamic setting and using the extreme case of a Bertrand market with free-entry.

The multi-product dimension of pricing has been developed within the monopolistic menu-cost framework with S,s pricing (Alvarez and Lippi, 2013 and Alvarez et al, 2013). One implication of the single product menu-cost model is that if a firm will only change its price if the price is far away from the flexible price optimum (far enough to take it to the edge of its "range of inaction"). Hence, with a single product monopolist you will observe few small price changes. The multi-product case is different: it can generate small price changes. Alvarez and Lippi assume that there is a single menu cost for adjusting all prices: that is, the same menu cost is incurred whether you adjust one or more prices. Hence, when the firm decides to change prices, it will adjust all of them: not only those prices that are far from the optimum, but also those that are near to the optimum. Since the marginal cost of adjusting a single price is zero (the same lump-sum menu cost applies), even small price changes become worth it. Hence, they find that the multi-product dimension is able to generate the high frequency of small price-changes that the single product models such as Midrigan (2011) cannot, leading to a bell shaped distribution of magnitudes for price changes when there are

more than 6 products. More importantly, the Alvarez-Lippi model predicts that an increase in uncertainty (the Brownian motion variance) will increase the frequency of price-changes (Proposition 4 combined with equation 11), as we find in the data during the more turbulent period of the crisis.

Our data has three multi-product retailers. To what extent do our retailers bunch their price changes, as in the Alvarez-Lippi model? Since ours is weekly data, we cannot observe any bunching that happens within the week (e.g. change more prices on Sundays when the shops are required by law to be open for 6 hours only). If we look at the weeks when there were no price-changes, there were 6 for Asda, 5 for Sainsbury and only 1 for Tesco. The coefficient of variation for the weekly frequency of price changes is 0.97 for Asda, 0.74 for Sainsbury and 0.92 for Tesco, all indicating that the standard deviation is large relative to the mean. We show the time-series for the weekly frequency of price-changes across the three stores in the Appendix Figure A5. There does not appear to be bunching of price-changes by retailer. There does appear to be fairly high correlation across retailers, particularly Asda and Tesco: the weekly frequencies have a correlation of 0.87. The weekly frequency of Asda and Tesco correlate less with Sainsbury: 0.57 and 0.59 respectively. The fact that price changes do not appear to be bunched over time, but rather correlated across retailers, suggests that the retailers are not acting as multi-product monopolists, but rather as oligopolistic competitors.

There is a further dimension from the demand side: supermarkets are amongst the set of retail outlets (other examples include DIY stores and restaurants) where the consumer normally shops for a basket of items.³⁴ Thus the sales of a particular item depend not only on its own price, but the range of prices charged in the store. A consumer might buy an item that is more expensive if it is purchased alongside a basket of goods that represent good value relative to the cost of the basket in other stores. If my rival as a firm cuts the price of a particular item- be it sliced white bread or whole chickens or whisky, then they may capture some of my consumers. To prevent this, I may change some prices even if that would not have been worthwhile in the absence of the competitive action. I am selling a basket and consumer demand for any one item will be finite.³⁵ Ultimately, as a multiproduct retailer, my

³⁴See Lach and Tsiddon (1996), Midrigan (2011) and Alvarez and Lippi (2013) for work in this area.

³⁵In the late 1990s, there was a supermarket “price war” in Britain on baked beans (Manez, 1999), in which the price of a can of beans fell to such a level that, reportedly, revenue did not even cover the cost of the can, let alone the contents. Nevertheless, even the British public’s desire for baked beans is finite, so whilst some people

earnings come through revenue from the baskets sold. It will not be optimal in any case to charge a uniform mark-up across my product lines (Bliss, 1988). Hence with careful selection I can provide tempting special offers whilst at the same time making money on the other things consumers purchase at my store.³⁶ The effect of bundling in grocery stores is also important. Clearly, the consumer is engaged in a two tier decision problem: which store to frequent and what to buy at that store. Whilst many loyal consumers shop repeatedly at the same store for their major shopping expedition, there are other consumers who switch if they think that there is something to be gained. Since consumers do not observe the full range of prices until they shop, marketing in the form of advertising a range of products as on sale or cheaper than at competitors can be seen as targeting those consumers thinking of changing which store to frequent.

Putting together a general theory of menu-costs in a dynamic oligopoly model is beyond the scope of this paper. However, the incentives to change price will clearly depend on what competitors are doing in addition to the industry demand and cost conditions. Hence one might expect the optimal flex-price of a firm (the price it would wish to set if there were no menu costs) to depend on the prices of the other firms as well as cost and demand. This is in complete contrast to the monopolistic competition models of current macro-modelling where the flex price depends only on general cost and demand parameters.

7.2 Basket prices in practice

Our findings have a number of implications. The first relates to the consumer experience. The overall impression coming from our analysis is that prices became more flexible in every dimension with the onset of the crisis: frequency, magnitude and timing. This is true not only of posted prices but also for NSB regular prices and, where relevant, KM13 regular prices and EJR reference prices. Since these are supermarket prices, consumers will, on the whole, be buying a basket of products. Experiments that we have carried out using baskets of goods lead to findings we might expect from our results on timing taken together with those on frequency. Each week a consumer seeking to buy the

no doubt stocked up to some extent, with limited storage space at home the potential damage to profits was limited.

³⁶ Bliss (1988) sets out the firm's problem more formally, albeit ignoring direct competition between supermarkets. The firm faces fixed costs of staff, heating and lighting, equipment and so on. It needs to cover these costs through mark-ups across goods. Optimally, these mark-ups vary across products, dependent upon demand characteristics (roughly speaking, elasticities). Formally the problem is equivalent to a Ramsey optimisation problem.

same basket of products will find the overall bill changed. Thus to the extent to which it is relevant to consider the basket price, rather than the prices of individual items, basket prices are not sticky at all.

The large number of price cuts observed does not mean that consumer basket prices are falling! Indeed, on any sensible definition of a basket, they are rising.³⁷ Clearly there is a composition effect at work within the magnitudes of price changes. There are many more prices falling than rising, but the falls are much smaller in percentage terms than the rises. Whilst this may create an illusion that prices generally are falling, in fact average prices and consumer basket purchases can and do rise. This is demonstrated in two related exercises we carry out below.

In the first exercise, illustrated in figure 16 below, we show weighted basket prices calculated from our data sample of 370 products, using weights equivalent to those used in the CPI. In other words, taking the NSB regular prices of our products, we allocate the products to the relevant component category of the UK CPI and construct the subcategory index using geometric means, then generate the arithmetic weighted mean index across product categories, in such a manner as to imitate the construction of the official CPI for the UK. It is clear from the figure that the general trend of prices is upwards over our period, though not monotonically so, and that the CPI basket ends the sample period substantially more expensive than it starts, in each supermarket. As we would expect given commodity price movements, the most rapid rise in our constructed version of the CPI for these chains is in 2008. We also see that, in common with other evidence, Sainsbury's takes a somewhat different path from Asda or Tesco, with somewhat higher pricing and a slightly looser relationship to the other two.

The second exercise compares this series of prices directly with two relevant CPI indices that we call CPI1, which is the index for the food, drink and tobacco group of products, and CPI2, a narrower index covering processed food and non-alcoholic drink only. As can be seen in table 8, the overall movements in the official CPI indices are mirrored very closely by the movements in our constructed indices for each of the chains. Thus, although more individual prices fall than rise, because the price rises are larger in magnitude than the

³⁷ This can also be seen crudely to be the case from table 2 above, as well as the slightly more sophisticated analysis below.

price falls, the typical basket price rises, and that roughly in line with general inflationary trends in the industry.

Figure 16: Price indices calculated using CPI weights from our sample of NSB regular prices

and Table 8 about here

This analysis which has drawn together the various dimensions of price movements has a further very potent implication. We have shown that prices in the latter period, by whatever criterion, are very flexible. But much of the flexibility appears spurious, in that it has the apparent aim of suggesting that supermarket prices are falling (which they are, in terms of a raw price count) when the shopper is in fact most likely to pay more for a basket of goods. The staggering number of price falls we observe at various points strongly suggests this.³⁸

Moreover, whilst we do not have cost data at individual product level, examination of individual price movements also strongly suggests that many of these movements do not relate to underlying movements in input costs. There are a number of reasons we say this. One is that there are some clearly documented sequences (“Edgeworth cycles” and similar phenomena; see Seaton and Waterson, 2013) where the price of an individual non-perishable product changes every week, going up one week, then falling by a penny a week over the next few weeks, before rising again. Another reason is because, to take an example, a price fall of one penny on a bottle of whisky priced at over ten pounds is not at all likely to have as its origin a fall in costs! This is not an isolated example in our data. The posted price graphs for exemplar products in figures A2 and A3 in Appendix 1 strongly suggest posted prices are not closely tracking costs. In fact, it would probably be reasonable to discount all but a very small fraction of the price falls we observe which relate to one, two or another small number of pence, as being a result of cost influences.

7.3 Crisis or changed behaviour?

We have identified the pricing patterns observed, in particular the change in behaviour as between the early part of our period of observation and the later period of crisis, as being caused by the crisis. Clearly, this assertion is potentially sensitive to the criticism

³⁸See our companion paper, Chakraborty et al (2011) for more on this point.

that what we are observing is a gradual change in behaviour by supermarkets, occasioned for example by technological change reducing menu costs, or a response to increased inflation. In order to test these possibilities, we have used our dataset to generate values for calendar months which can then be related to macroeconomic variables. We focus on three dependent variables: Fr being the proportion of prices which experienced at least one change in a particular calendar month; Up which is the proportion of prices which experienced at least one price rise in a month and Down the proportion of prices experiencing at least one price-cut. Note that with monthly intervals, a single price may move up and down over the weeks within the month, so that the sum of Up and Down may exceed FR.

Our macroeconomic explanatory variables are industrial output growth IO (a proxy for general output) and inflation in the IMF index of world food and beverage prices in sterling terms FB (general food and beverage inflation and the “costs” of our retailer’s inputs). IO and FB are included both for the current month and also over the last 12 months (the first capturing any short run influence, the latter a longer-run effect). These variables are all (trend) stationary. We also have a time trend (representing changes in behaviour over time) and a dummy for the crisis period which takes the value 1 from January 2008 onwards and we also include monthly dummies (not reported). Note we do not include a dummy for VAT changes, since most of the products are either exempt from VAT³⁹ or have an excise duty (Alcoholic beverages). Since alcohol duties usually change every April (the month of the UK budget), this is captured by the monthly dummy. We have 82 monthly observations available (May 2004 to October 2010), and we employ two methods of estimation - OLS and IV (with lagged values for instruments). The results are listed in Table 9.

Table 9: Monthly time-series analysis of Frequency data (posted prices).

Turning first to the OLS regressions, we can see that the overall fit is good as evidenced by the adjusted R-squared, given that we are not including any lagged dependent variables. There is no evidence of serial correlation but the monthly dummies are collectively significant, reflecting the clear seasonality in the data. Both the trend and crisis dummy are positive and significant for all three dependent variables. As for the macroeconomic variables, annual FB inflation is positive and significant for the overall frequency and for the frequency of price-cuts: the current monthly growth in industrial output is negatively related

³⁹ In the UK, most food and non-alcoholic beverages are exempt from VAT.

to the frequency of price increases. Overall, *whilst we can find some evidence that macroeconomic variables had an impact on the frequencies, with an upward trend, there is still a clear and significant “crisis effect” over and above these other factors.*

Whilst there is little question that the macroeconomic variables are exogenous in relation to our dependent variables, there might be measurement error in the sense that both macroeconomic variables might be poor proxies. Industrial output is only a small part of GDP, and is quite distantly related to the products sold in our sample. World food prices are only a part of the costs of the stores.

This suggests that an IV approach might be appropriate (2SLS). In stage 1 we use 4 lags on the exogenous variables, plus trend, crisis and monthly dummies to obtain the fitted values for stage 2. When we do this, we find that none of the macroeconomic variables is significant (although the signs are mostly unchanged). However, the trend is still positive and significant for all three dependent variables. Whilst the overall equation is satisfactory, no individual explanatory variable is significant for price rises except for the trend with IV. However, the crisis dummy is still positive and significant for the overall frequency and for the frequency of price-cuts.

In sum, our time-series results for the monthly data suggest that there is an upward trend in the frequency of price changes (certainly overall and for price-cuts). But more importantly, there is also a clear impact of the crisis on pricing behaviour. Price falls become more common in the crisis, and the impact is both statistically and economically significant even when we control for macroeconomic effects and different estimation methods.

8. Concluding Remarks

Why have our data come up with prices that are so much more flexible than previously observed? Several possibilities can be discounted. We work from posted, not scanner, prices so that excess volatility that may be present in the latter is not an issue. The three companies are major food retailers, not idiosyncratic small players- together their sales amount to perhaps $\frac{1}{4}$ of current consumer expenditure in the UK. We can discount differences in methodology, because where relevant we have used established methodologies

that smooth short term fluctuations to calculate our values. We are not working with fresh products where the market price naturally fluctuates.⁴⁰

We are however working with data that includes the significantly more turbulent macroeconomic crisis period 2008-2010, and this has led to several differences. Indeed, the results in the later period are more in line with those found in more turbulent economies (e.g. Gagnon, 2009). We are also working with data on oligopolistic companies where there is clear price rivalry, possibly intensified by their national presence and the national nature of their pricing structure. So the results are real, albeit that they challenge previous findings significantly.

At the same time, the findings we have documented relating to the staggering flexibility of pricing in British supermarkets since 2008 leave open several questions of macroeconomic interest. We have shown that, in an important category of consumer expenditure, prices are far from sticky, particularly in response to a crisis. But there remains the question of whether they respond to cost shocks quickly and flexibly. They may instead simply be responding to marketing pressures, which might actually drive prices further away from a relationship with costs. Because what ultimately matters to a supermarket chain is not how much it charges for an individual product, but rather the overall margin it earns on the range of products it sells. If marketing pressures drive the chain to lower a massive number of prices by a single penny, whilst simultaneously raising a smaller number of prices by a larger amount, the impact on price flexibility as described by macroeconomists is very unclear. Thus our findings do not fit with models of time-dependent pricing and do not fit well with established models of state-dependent pricing either. Nevertheless, our suggestion of marketing pressures as an underlying force remains rather speculative.

Cardiff Business School

Loughborough University

University of Warwick

⁴⁰Indeed, some of the products in our sample nearest to being “fresh” are amongst those with the fewest price changes, for example milk and cream.

APPENDICES

1. The relationship between Nakamura-Steinsson (NS) regular prices and Kehoe Midrigan (KM) regular prices and calculation of KM regular prices

As stated in the text, there are two versions of NS regular prices. One (algorithm B) is more straightforward and unambiguous than the other, because it simply replaces any short-lived “sale” price, by the price from which the price falls and to which it returns. Algorithm A seeks also to replace sale prices in cases where the price on return is different from the price before the sale, and has a method for determining whether the price on return is a new regular price. Both are comprehensively described in Nakamura and Steinsson (2010). It is clear from their description that this should only remove prices that are below the regular price: “Sale filter B removes price patterns in which the price returns to the original price within a set number of months without going above the original price. Sales filter A is designed to also remove price patterns in which a sale is followed by a change in the regular price, i.e. asymmetric V's. For example, for the 2 month case, we require that the price return to the original regular price in the first two months after the price decline occurs. If the product remains at a low price or is not available when the price collector returns in the first two months, then the original price decline is not defined as a sale.”⁴¹

The idea of the NS “A” algorithm is to cut out sequences of low prices between higher values that are not themselves identical. However, we found a significant number of sequences in posted prices where there was some ambiguity involved in determining what might be the stable price, if it was not the price before the set of lower prices. The following real sequence of prices over a 13 week period illustrates the point: {3.26, 3.26, 3.26, 3.28, 3.24, 3.24, 3.26, 3.23, 3.23, 3.69}. The difficulty relates to the presence of the 3.28 in the sequence, which would not be eliminated by the algorithm. Is it a sale price or a new regular price?

KM regular prices are defined differently from either of these, although they are nearer to NSA regular prices than NSB regular prices.⁴² The key difference, apparent both from the algorithm by which they are calculated and from the graphed examples in the

⁴¹ This is their definition based on monthly prices. Its timing is modified for weekly prices (Nakamura, 2008)

⁴² Kehoe and Midrigan (2010 footnote 3) discuss the relationship between the definitions.

paper, is that their approach also excludes short-lived increases in price. In that sense, they are smoothed to a greater extent than either NSB or NSA prices.

In our approach to KM regular prices, we initially took the window length, minimum appearance of the regular price in the posted price series, etc to be the same as theirs. However, we had some difficulty in generating the prices given that the mode (a key element in their procedure) was not always defined. For example, the thirteen week (again real) sequence {1.49, 1.48, 1.47, 1.88, 1.87, 1.86, 1.85, 1.84, 1.81, 1.78, 1.76, 1.75, 1.74} presents the problem that the mode is simply undefined because no price is repeated! To deal with this we manually replaced missing values for the mode with the most nearly previous regular price (this was required on more than 200 occasions). We also checked by inspection the last stage of the process, which is whether there were artificial changes in the regular price when the posted price did not change. We found that commonly, there was a lag (occasionally a lead) in the regular price. We also noted that, in the closing weeks when the underlying mode is calculated using less data, the calculation produced a slightly volatile series. Having accounted for this, we did not observe aberrant movements in our calculations. However, we also found that there were a number of six week periods that appeared to constitute “sale” pricing. Hence we adopted a thirteen week running mode as the underlying framework, rather than the eleven week mode KM use.

Two example graphs are shown in figures A1 and A2 below, each for a single product from one of the three supermarkets- one is for a product where there are a large number of small price changes, the other for a product where larger price changes are more frequent. In the first, we observe that the KM11 algorithm removes most temporary falls (e.g. weeks 123-125) and temporary rises (weeks 61 and 62) but that some longer-lived price fluctuations are not eliminated, including a six week cut (weeks 43-48 inclusive) and an erratic period starting in week 187 where the posted price rose briefly, fell to a lower price than before for five weeks, then rose to a new higher price. A similar event starts in week 298, some of which is eradicated using our 13 week window that was not eliminated when an 11 week window was used. In the case of the second product, the price moves very often and in many ways, but significant amounts of this movement are eliminated given our definition of KM13 regular prices.

Figures A1 and A2 about here

The first product illustrates a further feature. Although the regular price never moves when the posted price does not, there is an example of a “singleton” around period 212 where the regular price moves downward for just one period before moving up (to a different level than before).⁴³ Although somewhat ad hoc, it seems reasonable also to eliminate such observations when calculating KM regular prices. There were 51 such cases in total across the set of products that we eliminated manually by replacing the value with the previous regular price. Having done this, it gave us our final series, as used in the text, of KM13 (smoothed) regular prices. Even after these various changes, at least 60% of the products (80% in Asda and Tesco) in our sample have median price durations of less than 20 weeks. No direct comparisons with NSB regular prices are possible, because of the rather different methods by which they are developed.

For comparison with figures in the text and previous analyses, we append figure A3 constructed using (unsmoothed) KM11 regular prices below.

Figure A3 about here

2. Additional material referred to in the text

Figure A4, A5, A6; Tables A1 and A2 here

⁴³The posted price stays down for several periods, which coupled with the fact that the “before” and “after” prices differ, can yield a singleton low mode, because this is, for that period, the most commonly observed price in the 13 week sequence centred on it.

REFERENCES

- Alvarez, Fernando E and Francesco Lippi (2013), "Price setting with menu cost for multi-product firms", EIEF Working Paper 02/13 (forthcoming, *Econometrica*).
- Alvarez F, Le Bihan H, Lippi F. (2013). "Small and large price changes and the propagation of monetary shocks." CEPR DP 9770.
- Baharad, E. and Eden, B (2004). Price Rigidity and Price Dispersion: Evidence from micro data, *Review of Economic Dynamics*, 7, 613-641.
- Barro, Robert J (1972). "A Theory of Monopolistic Price Adjustment." *Review of Economic Studies* 39 (1):17–26.
- Bils, Mark and Klenow, Peter (2004), "Some evidence on the importance of sticky prices", *Journal of Political Economy*, 112, 947-985.
- Bliss, Christopher, "A Theory of Retail Pricing," *Journal of Industrial Economics*, 36(4) (1988), 375-91.
- Bulmer, M. G., *Principles of Statistics* (Dover, 1979).
- Bunn, Philip and Colin Ellis, "Examining the behaviour of individual UK consumer prices," *Economic Journal* 122 (558) (2012), F35- F55.
- Cavallo, Alberto, "Scraped data and sticky prices", Mimeo, MIT Sloan, April 2012.
- Chakraborty, Ratula, Paul Dobson, Jonathan S. Seaton and Michael Waterson, "Pricing in inflationary times: the penny drops", Warwick Economics Research Paper no 975, 2011.
- Competition Commission (CC), *The Supply of groceries in the UK Market Investigation - Final Report*, April 2008 (London, Competition Commission.) (http://www.competition-commission.org.uk/rep_pub/reports/2008/538grocery.htm).
- Dixon, Huw (2009), A unified framework for understanding and comparing dynamic wage and price setting models, Working Paper 257, Banque de France.
- Dixon, Huw and Kara, Engin (2010), Can We Explain Inflation Persistence in a Way that Is Consistent with the Microevidence on Nominal Rigidity? *Journal of Money, Credit and Banking*, 42, 151-170.
- Dixon, Huw and Le Bihan, Herve (2012), "Generalised Taylor and Generalised Calvo Price and Wage Setting: Micro-evidence with Macro Implications," *Economic Journal*, vol. 122 (560), pages 532-554,
- Dixon, Huw and Tian, Kun (2013), "What We Can Learn about the Behavior of Firms from the average Monthly Frequency of Price-Changes: An Application to the UK CPI Data", CESifo Working Paper No. 4226.

Dixon, Huw, Franklin, Jeremy and Millard, Steven (2013), Sectoral shocks and monetary policy in the United Kingdom, mimeo Bank of England.

Dhyne, Emmanuel, Luis J. Alvarez, Herve Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffman, Nicole Jonket, Patrick Lunnemann, Fabio Rumler and Jouko Vilmunen, 2005. "Price setting in the euro area: Some stylized facts from Individual Consumer Price Data," Working Paper Research 74, National Bank of Belgium.

Eichenbaum, Martin, Nir Jaimovich and Sergio Rebelo, "Data appendix" to Eichenbaum et al., 2011. Mimeo, Kellogg Graduate School, Northwestern University, 2009.

Eichenbaum, Martin, Nir Jaimovich and Sergio Rebelo, 2011. "Reference prices, costs, and nominal rigidities", *American Economic Review* 101 (February 2011), 234-262.

Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo and Josephine Smith, 2012. "How frequent are small price changes?" Mimeo, Kellogg Graduate School, Northwestern University, February 2012.

Ellis, Colin, "Do supermarket prices change from week to week?", Bank of England Working Paper no 378, November 2009.

Gagnon, Etienne, "Price Setting during Low and High Inflation: Evidence from Mexico," *Quarterly Journal of Economics* 124 (August 2009), 1221 – 1263.

Griffith, Rachel, Ephraim Leibtag, Andrew Leicester and Aviv Nevo, (2009). "Consumer Shopping Behavior: How Much Do Consumers Save?" *Journal of Economic Perspectives*, 23(2), 99-120.

Guimaraes, Bernardo and Sheedy, Kevin D (2011). Monetary Policy and Sales, *American Economic Review*, 111, 844-76.

Kehoe, Patrick J. and Virgili Midrigan, "Prices are sticky after all", Federal Reserve Bank of Minneapolis, Research Department Staff Report 413, September 2010.

Klenow, Peter J. and Oleksiy Kryvtsov, "State-Dependent or Time-Dependent Pricing: Does It Matter For Recent U.S. Inflation?" *Quarterly Journal of Economics* 123 (August 2008): 863-904.

Klenow, Peter J. and Benjamin A. Malin, "Microeconomic evidence on price setting", in B.M. Friedman and M. Woodford, eds, *Handbook of Monetary Economics* 3A (2010) Amsterdam: Elsevier.

Lach, Saul and Daniel Tsiddon, "Staggering and synchronization in price-setting: Evidence from multi-product firms" *American Economic Review* 86(5) (1996), 1175-1196.

- Levy, Daniel, Mark E. Bergen, Shantanu Dutta, and Robert Venable, “The magnitude of menu costs: Direct evidence from large US supermarket chains”, *Quarterly Journal of Economics* 112(3) (1997), 791-825.
- Manez, J. “Issues in UK food retail pricing”. Unpublished PhD thesis, University of Warwick, Coventry UK, 1999.
- Mankiw, N. Gregory, “Small menu costs and large business cycles: a macroeconomic model of monopoly” *Quarterly Journal of Economics* 100(2) (1985), 529-537.
- Midrigan, Virgilu, “Menu costs, multiproduct firms and aggregate fluctuations” *Econometrica* 79 (2011), 1139-1180.
- Nakamura, Emi, “Pass-through in retail and wholesale”, *American Economic Review: Papers and Proceedings*, 98(2) (2008), 430-437.
- Nakamura, Emi and Jon Steinsson, “Five facts about prices”, *Quarterly Journal of Economics* 123(4) (2008), 1415-1464.
- Nakamura, E. and Steinsson, J. (2010), “More facts about prices”, mimeo, Columbia University (<http://www.columbia.edu/~en2198/papers/fivefactssupplement.pdf>).
- Seaton, Jonathan and Waterson, Michael (2013), Identifying and characterising price leadership in British supermarkets, *International Journal of Industrial Organization*, 31, 392-403.
- Varian, Hal (1980), A Theory of Sales, *American Economic Review*, 70, 651-9.
- Vavra, Joseph (2013), Inflation dynamics and time-varying volatility: new evidence and an S,s Interpretation, forthcoming, *Quarterly Journal of Economics*.
- Verdict Research (2008) *UK Retail Futures 2012: Food and Grocery*. (London, Verdict).

TABLES

Table 1: The two distributions for the price-spells of Paxo Stuffing.

weeks	1	2	5	6	14	82	106	144
Duration	0.1111	0.1111	0.2222	0.1111	0.1111	0.1111	0.1111	0.1111
CS.	0.0027	0.0055	0.0274	0.0164	0.0384	0.2247	0.2904	0.3945

Table 2: Basic descriptive statistics of our data sample- 366 weeks and 370 products

	min	median	mean	max
£	price	price	price	price
10/11/2003	0.15	0.99	1.85	25.99
8/11/2010	0.17	1.24	2.23	29.99

Table 3: Mean durations of price spells under different metrics (weeks).

	2003-2010	Moderation	Crisis
Estimated (freq)	9.8	18.6	5.8
Measured	11.1	17.0	9.6
Measured median	3.5	7.5	2.5
Cross-section	37.5	48.4	21.4
Estimated NSB (freq)	13.4	26.3	7.8
Estimated KM13 (freq)	25.0	32.4	18.4
Estimated EJR (freq)	29.8	42.3	20.8

Note: The EJR is based on a quarterly window.

Table 4: Average number of products experiencing price changes in the ranges shown, by supermarket and year, NSB prices.

	2004	2005	2006	2007	2008	2009	2010
Asda							
0	136	71	82	25	4	27	29
1 – 4	223	284	258	254	85	87	170
5 – 8	11	15	28	89	79	88	149
9 - 12	0	0	2	2	81	76	20
More	0	0	0	0	121	92	2
Sainsbury							
0	89	108	93	37	13	28	30
1 – 4	275	250	260	291	232	194	247
5 – 8	6	12	17	41	85	80	81
9 - 12	0	0	0	1	31	54	11
More	0	0	0	0	9	14	1
Tesco							
0	118	66	74	23	2	30	17
1 to 4	237	269	252	267	67	94	170
5 to 8	15	34	40	72	92	102	139
9 to 12	0	1	3	7	80	67	41
More	0	0	1	1	129	77	3

Table 5: Average percentage NSB regular price change across chain and year-

NSB Regular prices

Percentage magnitude

Rises	2004	2005	2006	2007	2008	2009	2010
Asda	13.3	15.0	14.7	18.4	24.0	22.8	21.4
Sains	11.9	11.5	12.3	14.5	17.3	11.2	9.1
Tesco	13.6	12.9	13.2	15.0	22.4	19.0	17.3
Average	13.2	13.2	13.4	15.9	21.2	17.5	15.4
Falls							
Asda	-8.5	-8.5	-8.7	-8.6	-8.7	-8.8	-9.0
Sains	-7.8	-7.8	-7.8	-7.7	-7.7	-7.8	-7.7
Tesco	-7.0	-7.0	-7.0	-7.0	-7.0	-7.0	-6.9
Average	-7.7	-7.7	-7.8	-7.7	-7.8	-7.8	-7.8

Table 6: Average percentage posted price change across our products by year and chain**Percentage magnitude**

Rises	2004	2005	2006	2007	2008	2009	2010
Asda	13.0	14.3	18.0	21.3	23.9	24.4	21.3
Sains	17.6	13.8	16.8	21.3	23.8	20.1	20.5
Tesco	14.0	17.3	16.2	19.7	22.3	22.4	21.7
Average	14.6	15.2	17.0	20.8	23.3	22.3	21.1
Falls							
Asda	-8.2	-8.2	-8.2	-8.0	-8.3	-8.3	-8.2
Sains	-8.7	-8.7	-8.7	-8.8	-8.8	-8.9	-8.9
Tesco	-9.7	-9.7	-9.6	-9.6	-9.7	-9.7	-9.7
Average	-8.8	-8.8	-8.8	-8.8	-8.9	-8.9	-8.9

Note: these mean values ignore cases where no change has occurred in that year, so that a percentage cannot be calculated.

Table 7: Proportion of posted price falls that are 1p in size

	2004	2005	2006	2007	2008	2009	2010
Asda	0.16	0.19	0.22	0.12	0.24	0.36	0.17
Sains	0.11	0.12	0.27	0.17	0.19	0.27	0.16
Tesco	0.18	0.22	0.25	0.14	0.27	0.32	0.20
Total	0.15	0.19	0.25	0.14	0.24	0.33	0.18

Table 8: Showing correlations between two common official CPI indices and constructed NSB regular price indices for our three chains

Correlations		CPI2	Asda	Sains	Tesco
2003-2010	CPI1	0.998	0.986	0.977	0.987
	CPI2		0.984	0.974	0.985

Note: CPI1 is the index for the food, drink and tobacco group of products, and CPI2 is a narrower index covering processed food and non-alcoholic drink only

Table 9: Time-series analysis of monthly frequency data (posted prices).

D.V.	OLS			IV		
	FR	Up	Down	FR	Up	Down
IO(m)	-1.25 (0.92)	-0.95** (0.47)	-0.66 (0.95)	0.46 (2.14)	-0.88 (1.08)	1.20 (2.21)
IO(a)	-0.52 (0.52)	-0.02 (0.27)	-0.55 (0.54)	-0.48 (0.96)	-0.06 (0.48)	-0.52 (0.99)
FB(m)	-0.10 (0.17)	-0.05 (0.09)	-0.06 (0.17)	-0.38 (0.24)	-0.19 (0.12)	-0.29 (0.25)
FB(a)	0.14** (0.06)	0.04 (0.03)	0.13** (0.06)	-0.13 (0.11)	0.04 (0.06)	0.12 (0.12)
CRISIS	0.18*** (0.04)	0.04** (0.02)	0.18*** (0.04)	0.18** (0.07)	0.03 (0.04)	0.19** (0.07)
Trend	0.004*** (0.0004)	0.0002*** (0.00002)	0.002*** (0.0004)	0.004*** (0.0007)	0.002*** (0.0004)	0.003*** (0.0008)
\bar{R}^2	0.85	0.75	0.77	0.83	0.73	0.75

There are 82 calendar months, from Jan 2004 to October 2010 inclusive. The exogenous variables: IO is industrial output; FB is the IMF food and beverage world price index expressed into sterling. Both FB and IO are expressed as current monthly growth (m) and annual growth over last 12 months (a). Crisis is a dummy that takes the value 1 from January 2008 onwards. All variables are stationary (FB trend-stationary). The constant and monthly dummies were included but are not reported. The IV estimates are 2SLS with the instrument set being lags 1-4 of the four exogenous variables. ** indicates significance at the 5% level, *** at the 1% level. The dependent variables are Fr (proportion of prices experiencing at least one price change in the calendar month); Up (proportion of prices experiencing at least one price rise in the calendar month); Down (proportion of prices experiencing at least one price rise in the calendar month). Posted prices were used.

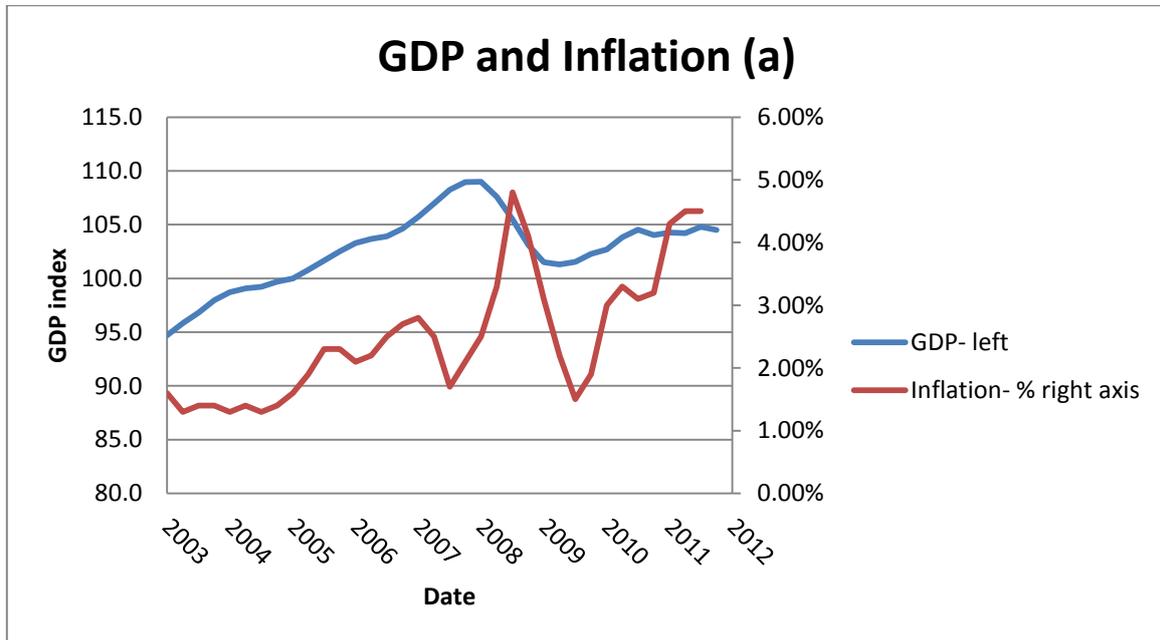
Table A1: Number of weeks in which there are KM13 regular price changes in the ranges shown

KM13 regular prices	No changes	1 to 9	10 to 19	20 to 29	30 to 39	40 to 49	50 to 99	100 or more changes
Rises	1	66	129	101	34	23	11	0
Falls	2	44	129	88	49	30	21	2

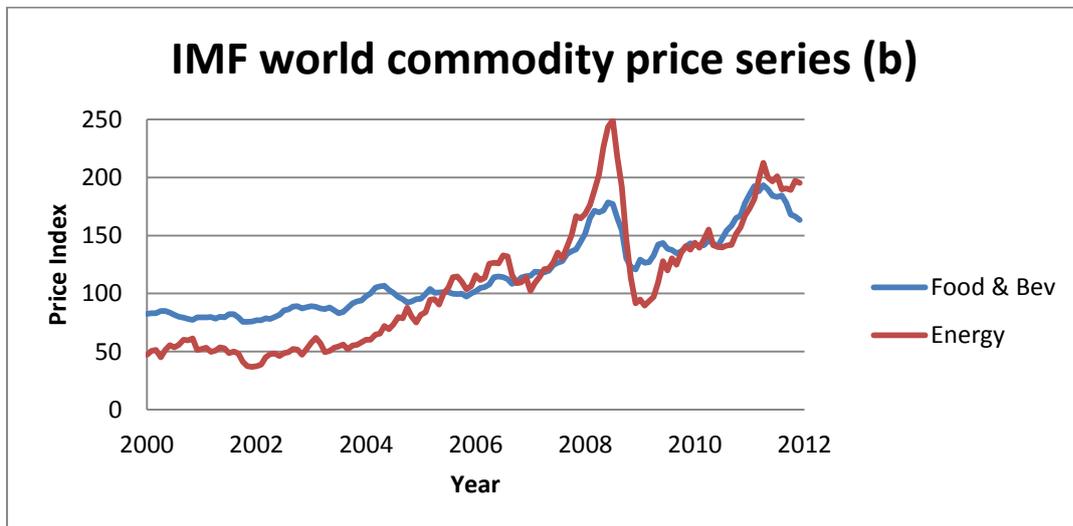
Table A2: Price cuts of 2p in magnitude

	2004	2005	2006	2007	2008	2009	2010
Asda	32	47	34	68	496	240	78
Sains	37	27	17	74	121	154	79
Tesco	40	58	45	99	409	217	86

FIGURES



Source: ONS. Inflation is CPI index



Source: IMF

Figure 1: Key macroeconomic factors underlying our framework

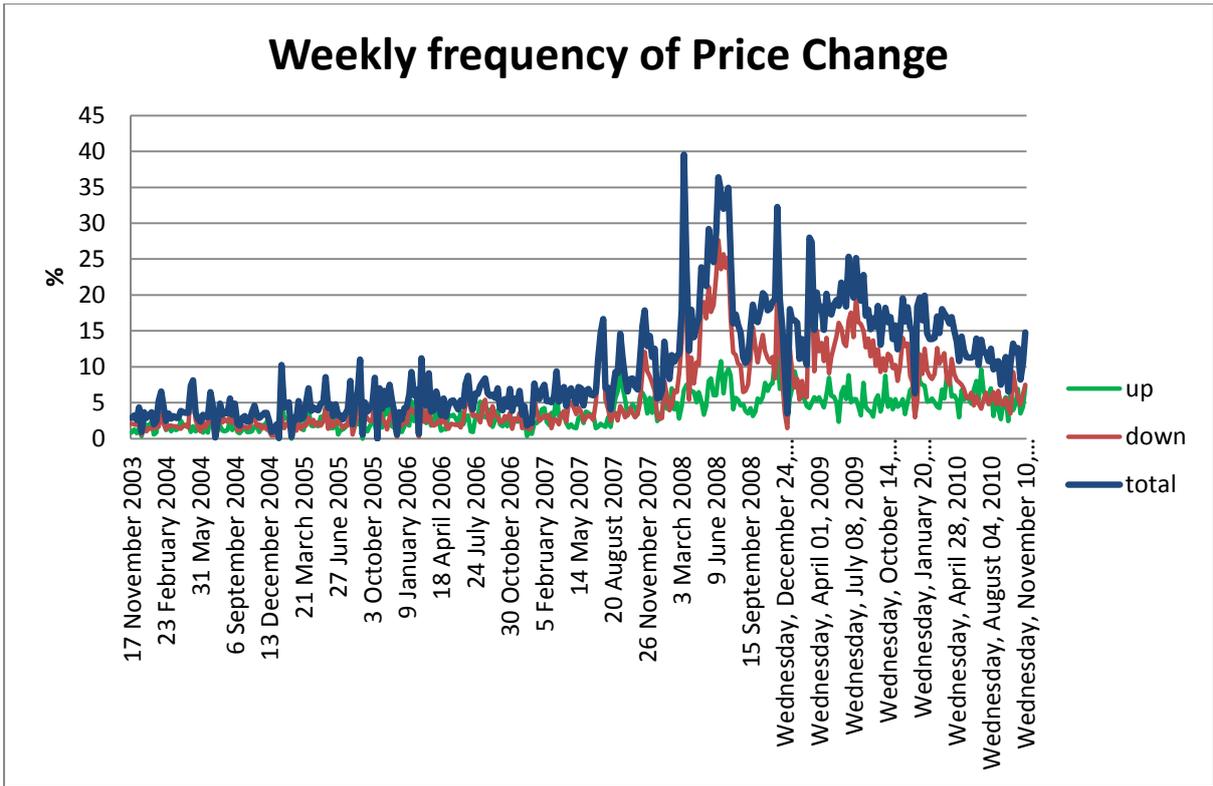


Figure 2: The weekly frequency of posted price changes.

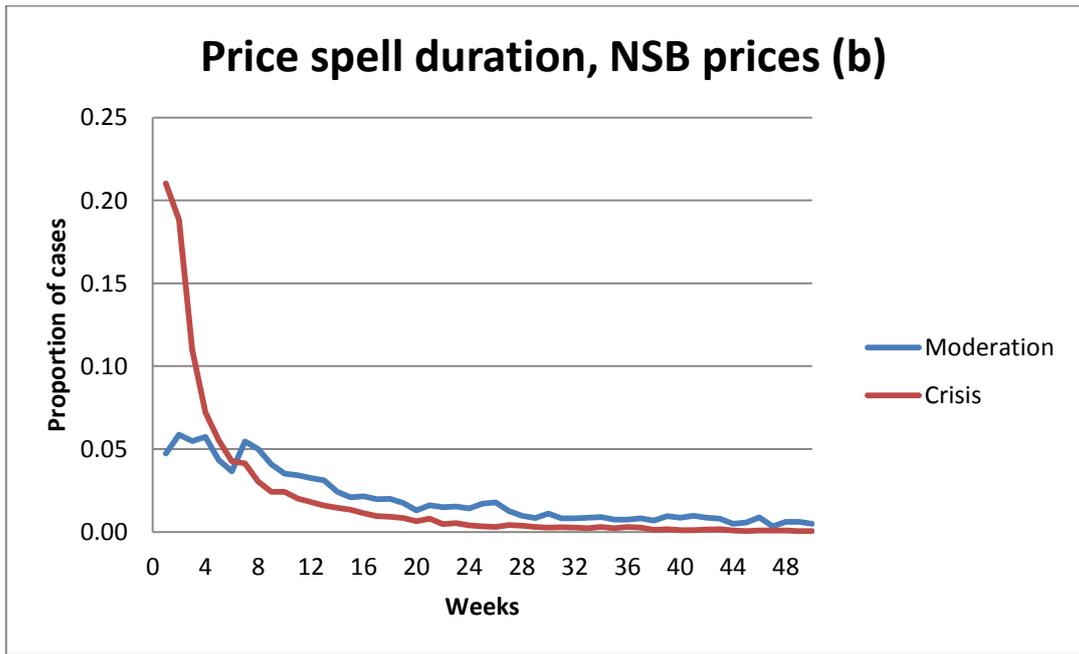
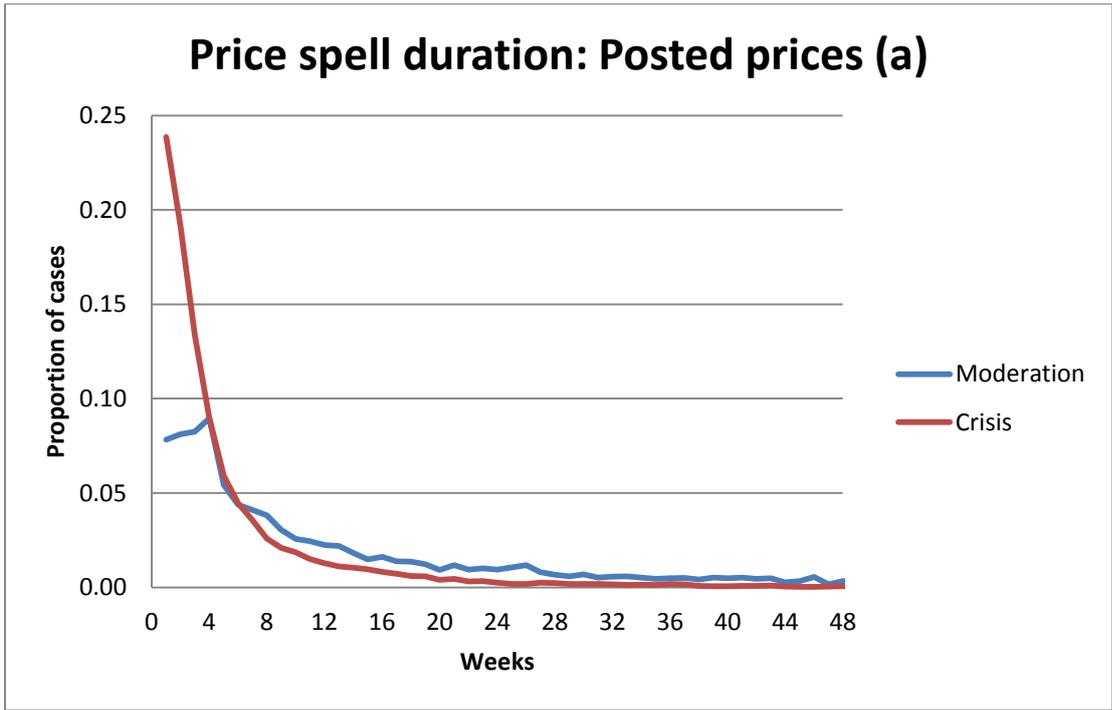


Figure 3: The distribution of price-spells compared in the moderation and the crisis, using posted and alternatively NSB prices.

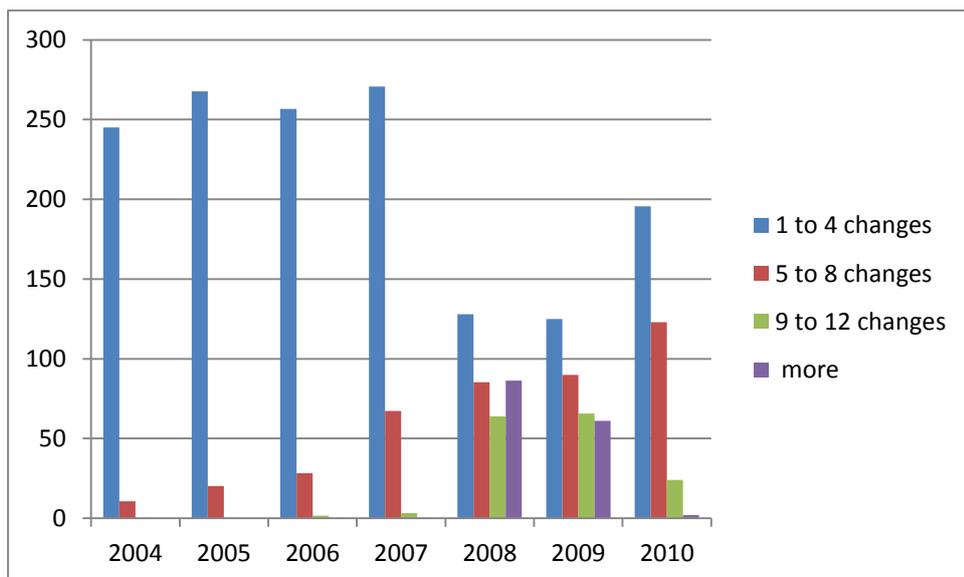


Figure 4: Count of numbers of products (vertical axis) experiencing the number of price changes indicated per year, NSB prices.

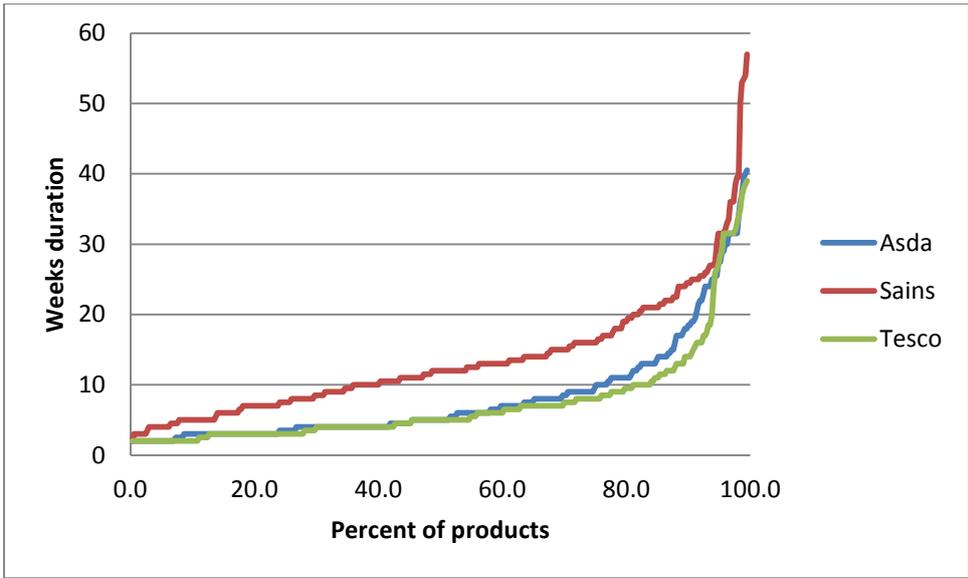


Figure 5: Distribution of median duration of NSB regular prices across our sample of products

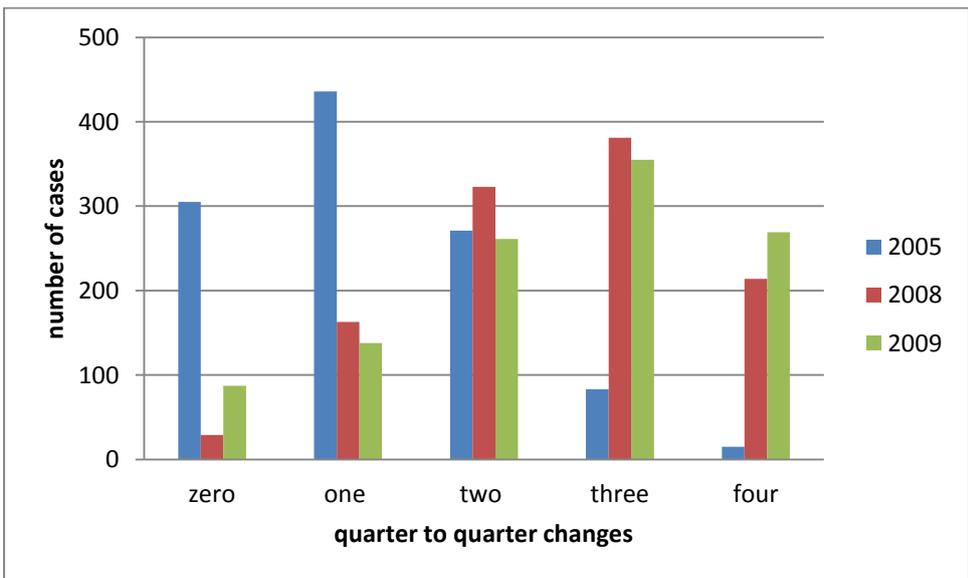


Figure 6: Number of quarterly price changes in EJR reference prices per year across 370 products and three firms

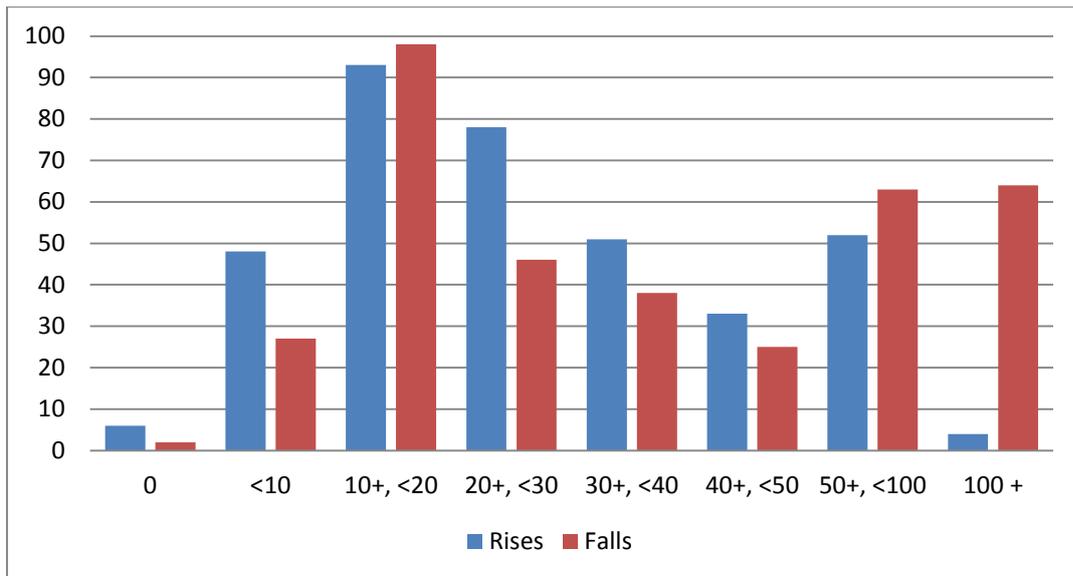


Figure 7: Number of weeks (vertical axis) in which there are NSB regular price changes (horizontal axis) in the ranges shown

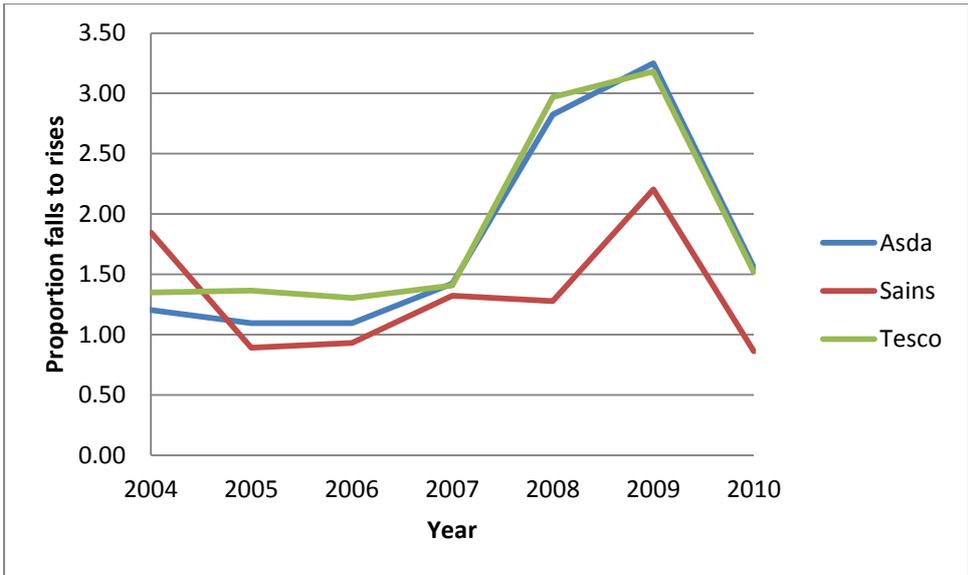


Figure 8: The ratio of NSB regular price falls to price rises across our sample and time

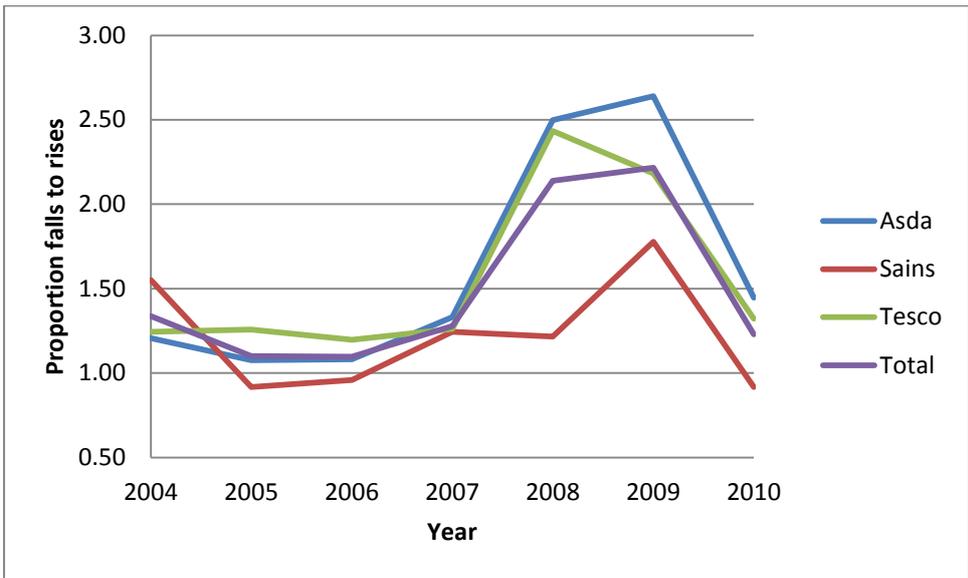


Figure 9: The ratio of posted price falls to price rises across our sample and time

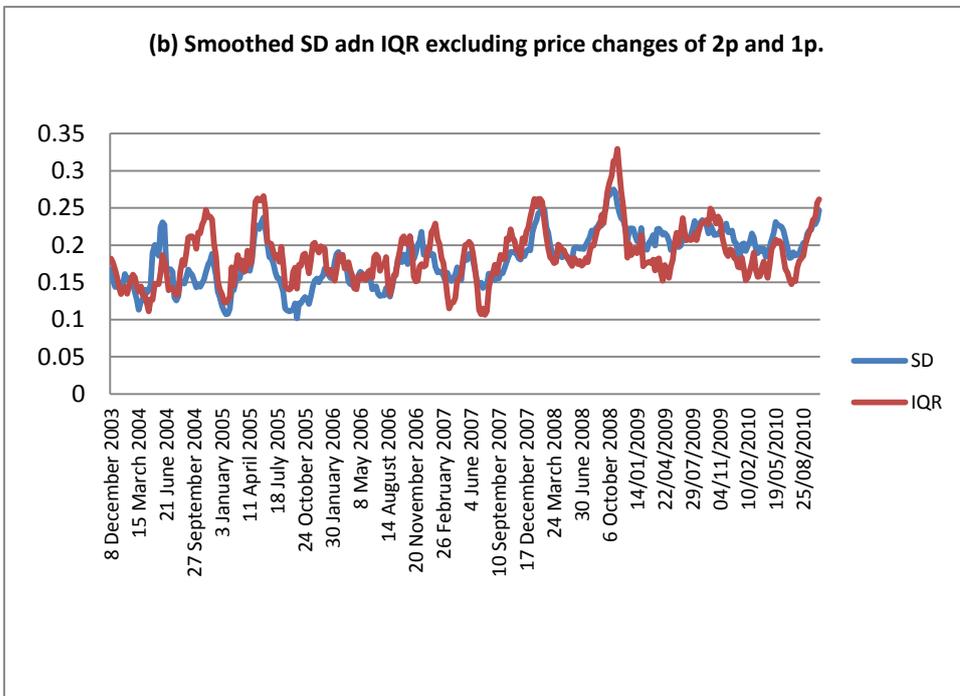
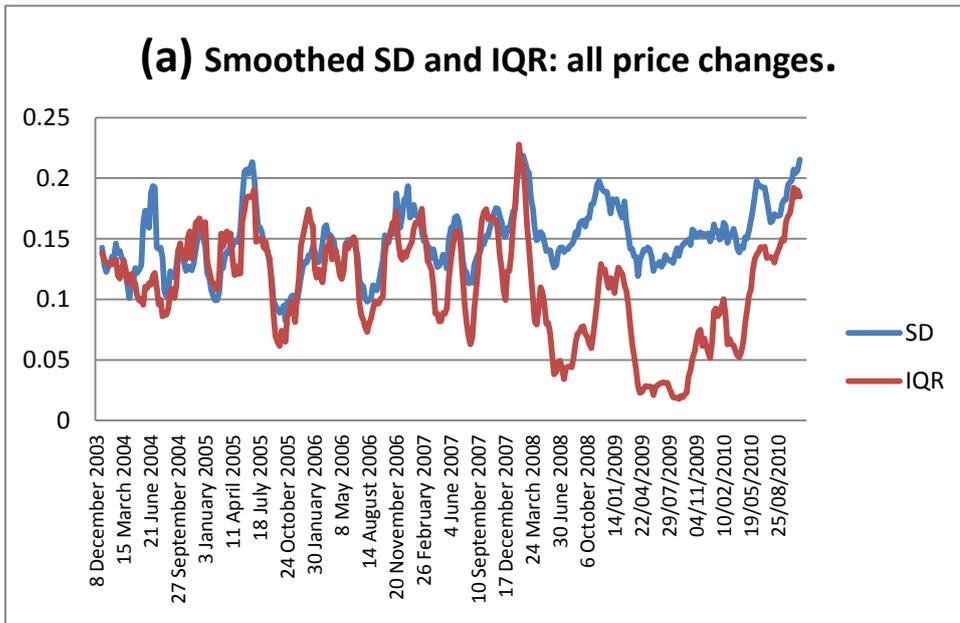


Figure 10: The Dispersion of NSB price changes with and without small changes.

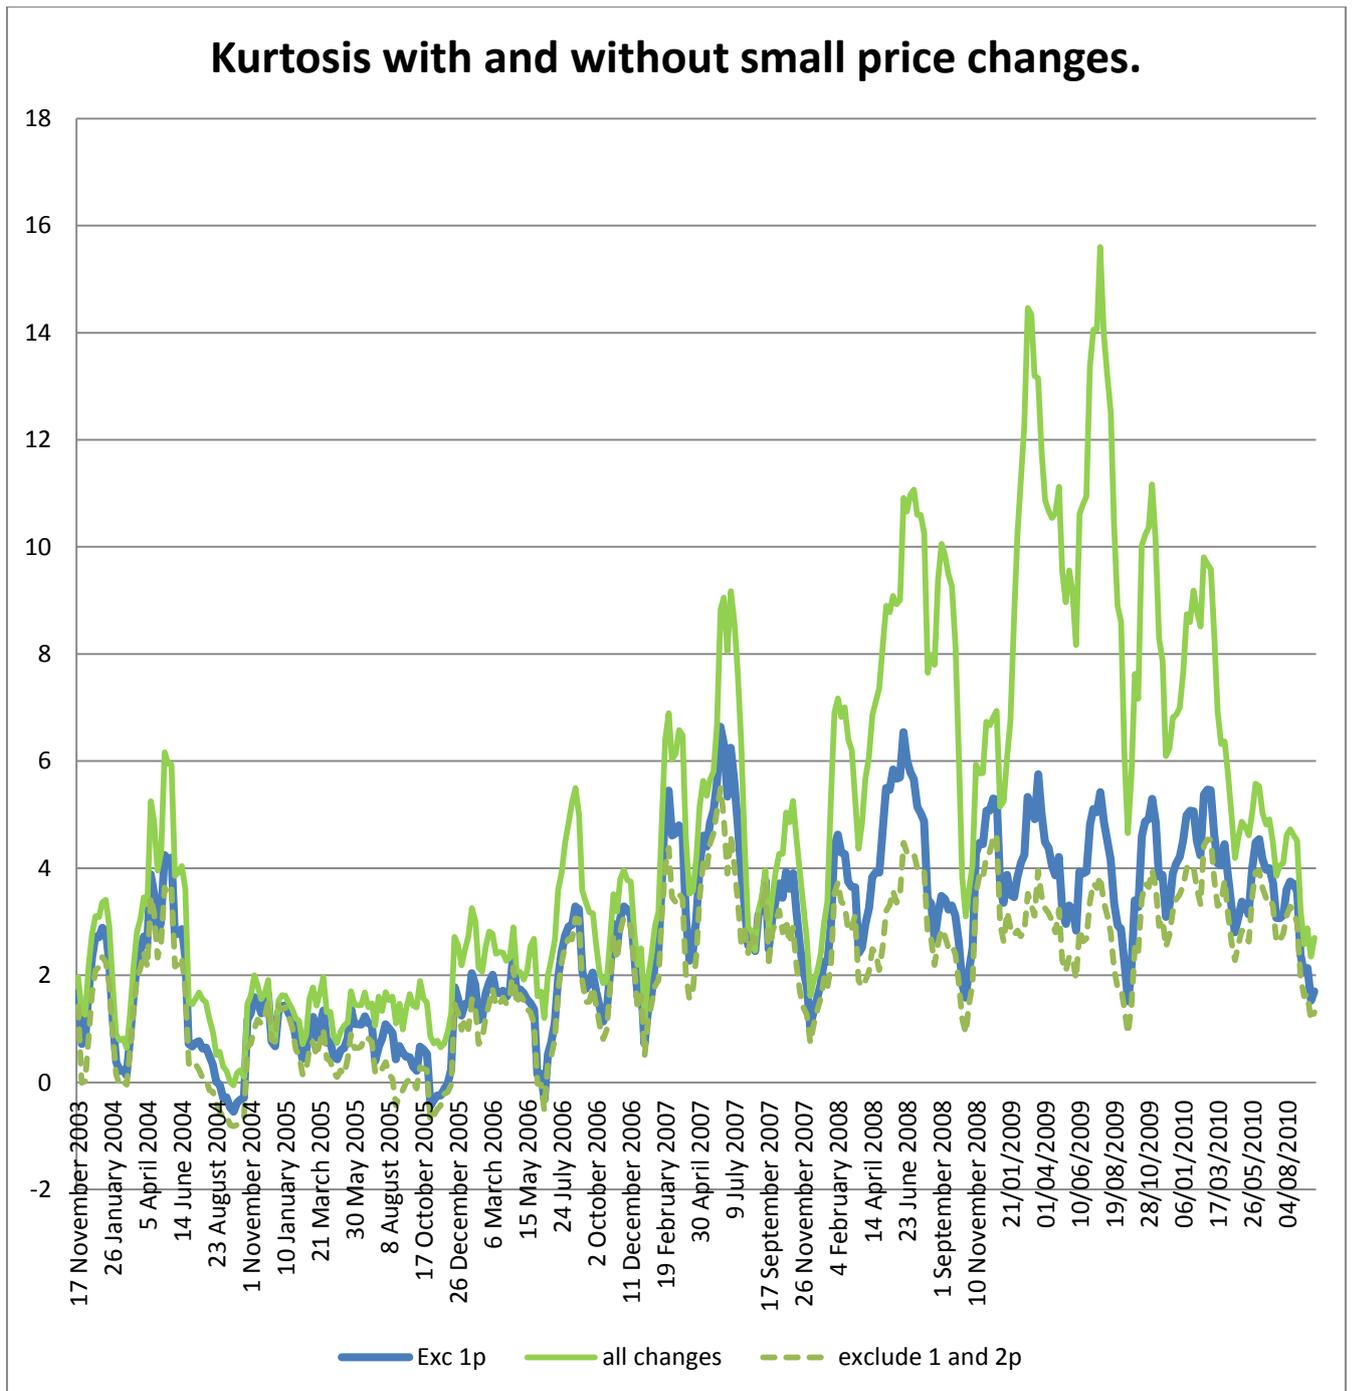


Figure 11: The smoothed Kurtosis of Price growth with and without small price changes.

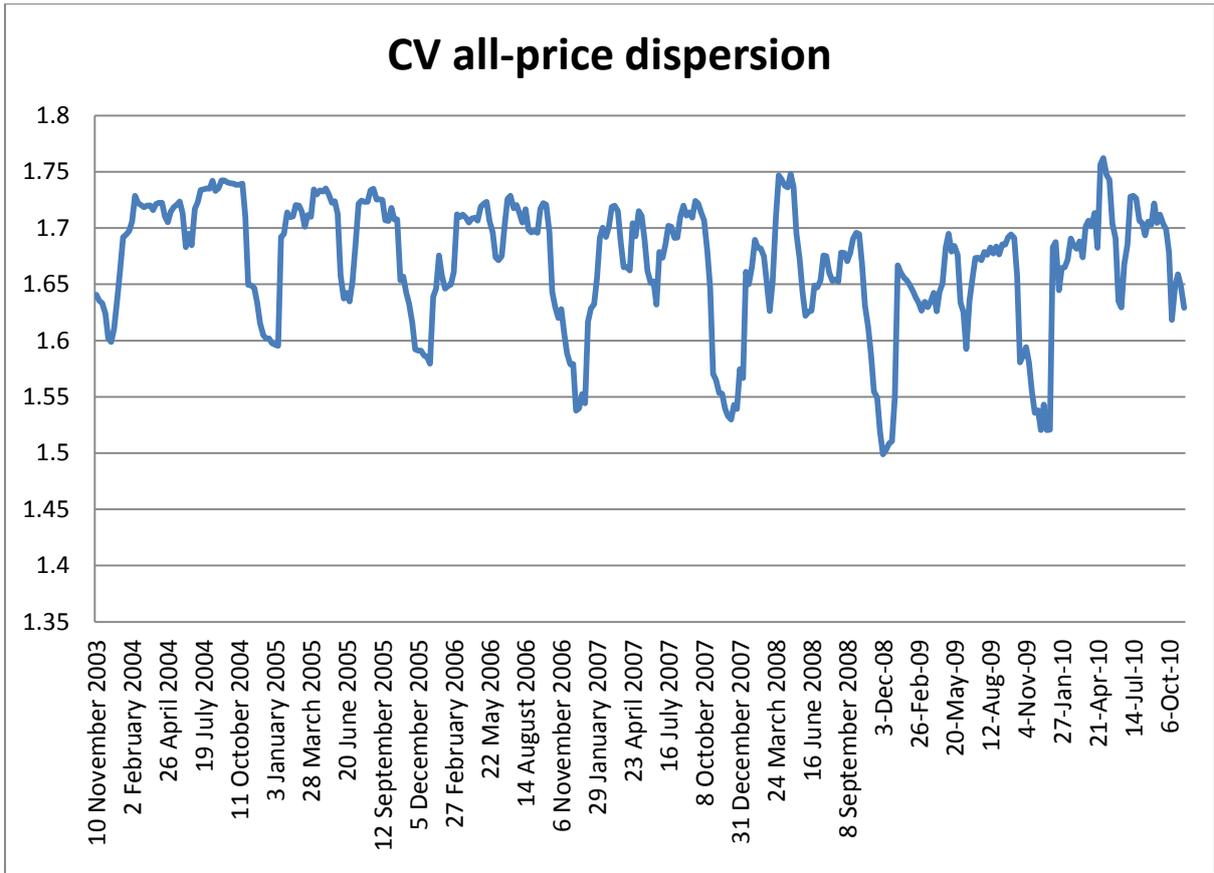


FIGURE 12: Posted price dispersion across all prices as measured by the coefficient of variation.

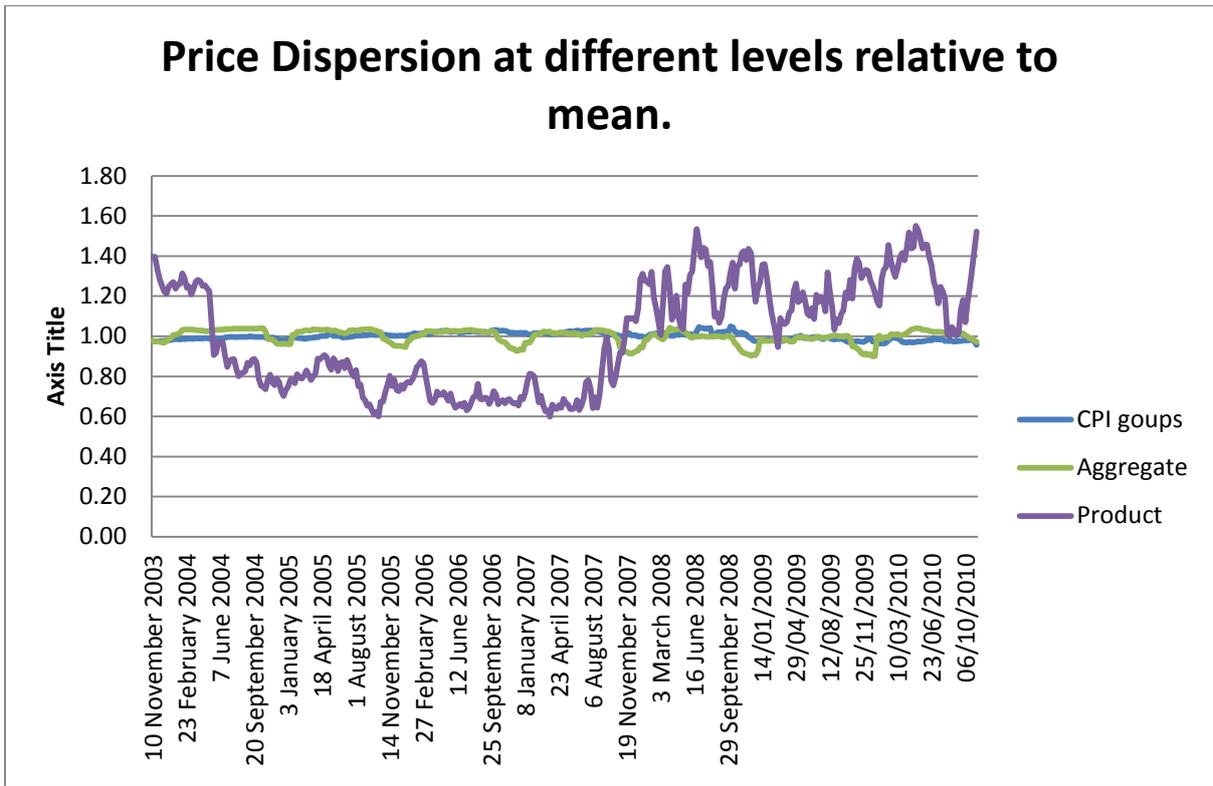


Figure 13: The normalised CV for NSB prices at different levels of aggregation.

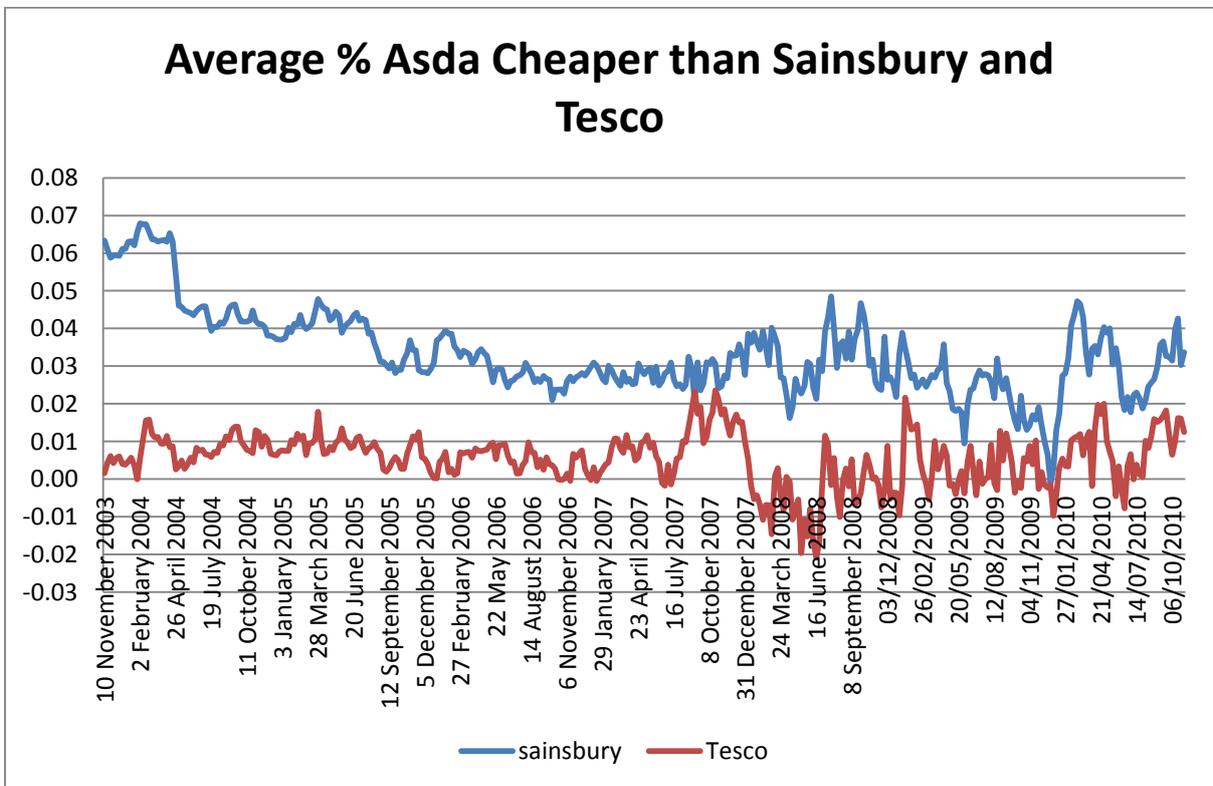
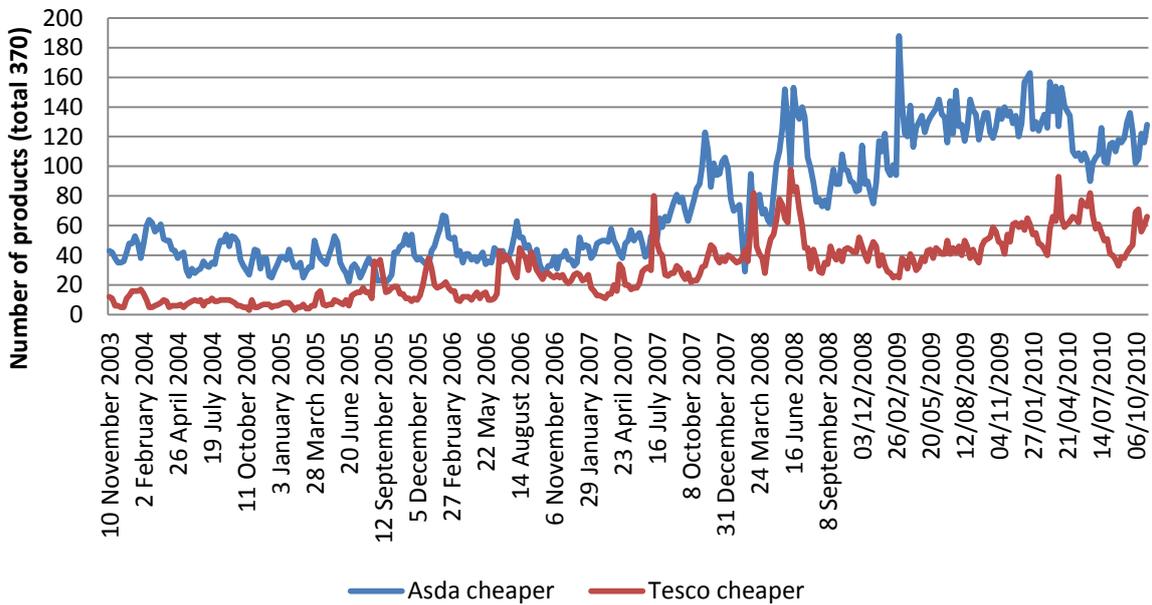


Figure 14: The Average % by which Asda is cheaper than Sainsbury and Tesco.

(a) Number of products cheaper at ASDA and TESCO.



(b) Number of products with same price at Asda and Tesco.

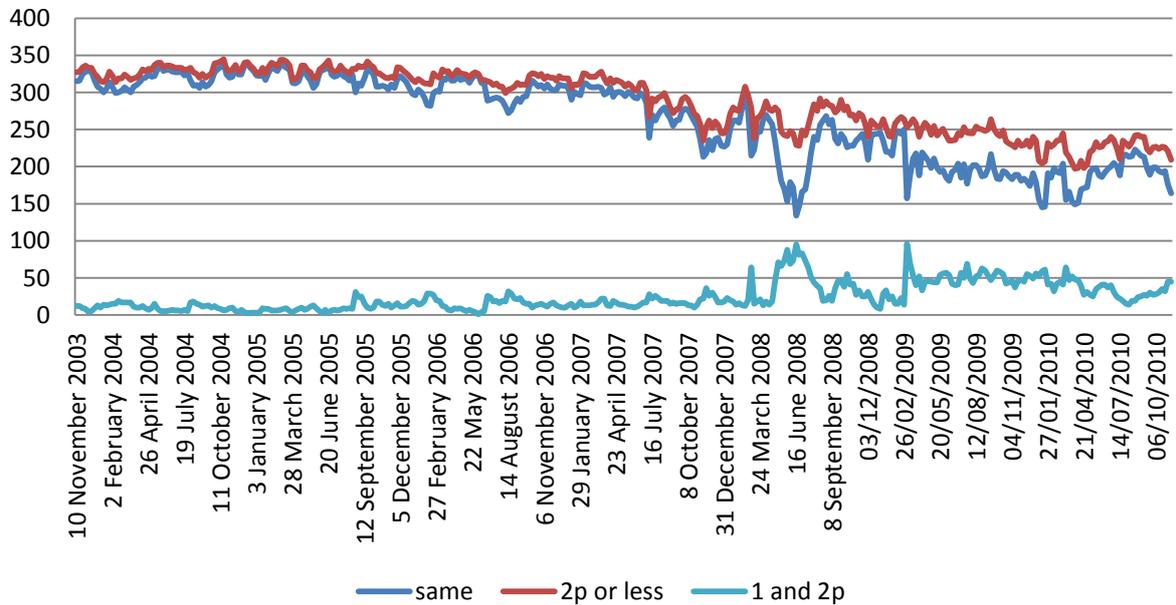


Figure 15: Price Competition between Asda and Tesco.

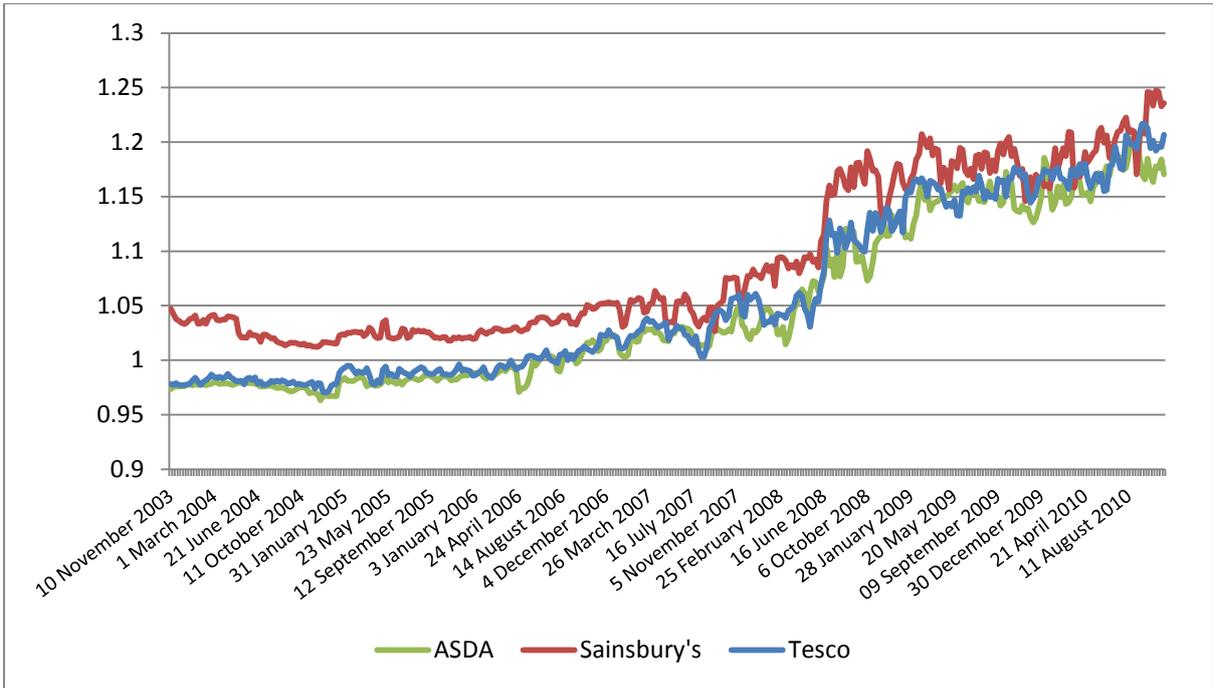


Figure 16: Price indices calculated using CPI weights from our sample of NSB regular prices

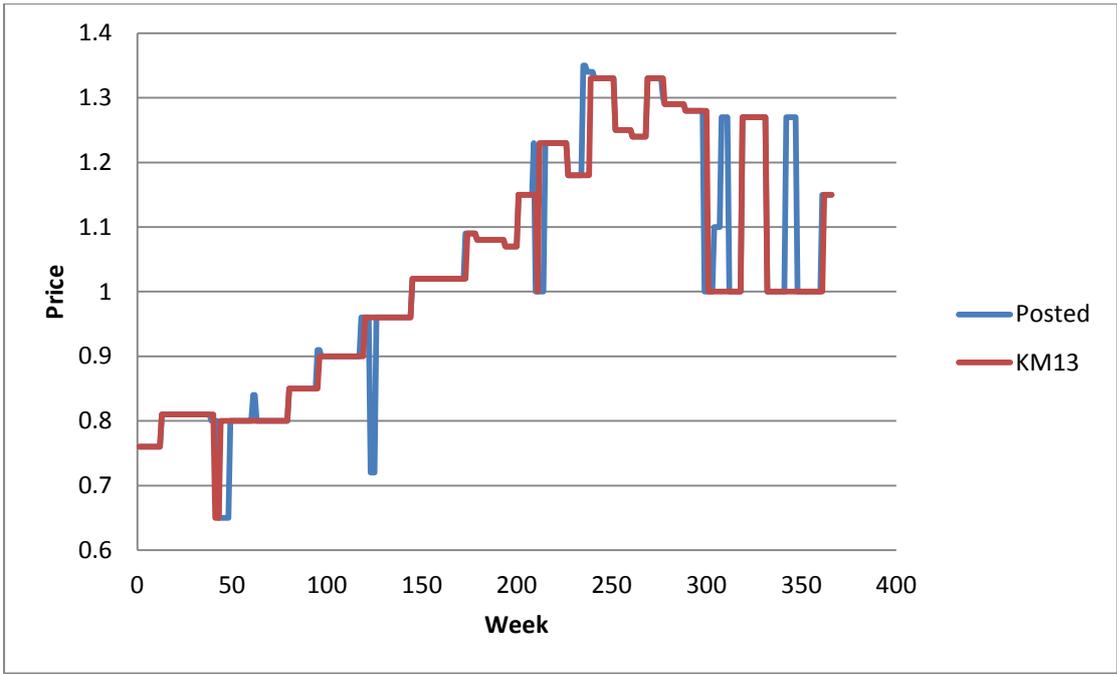


Figure A1: Posted and KM 13 regular prices relating to a bread product

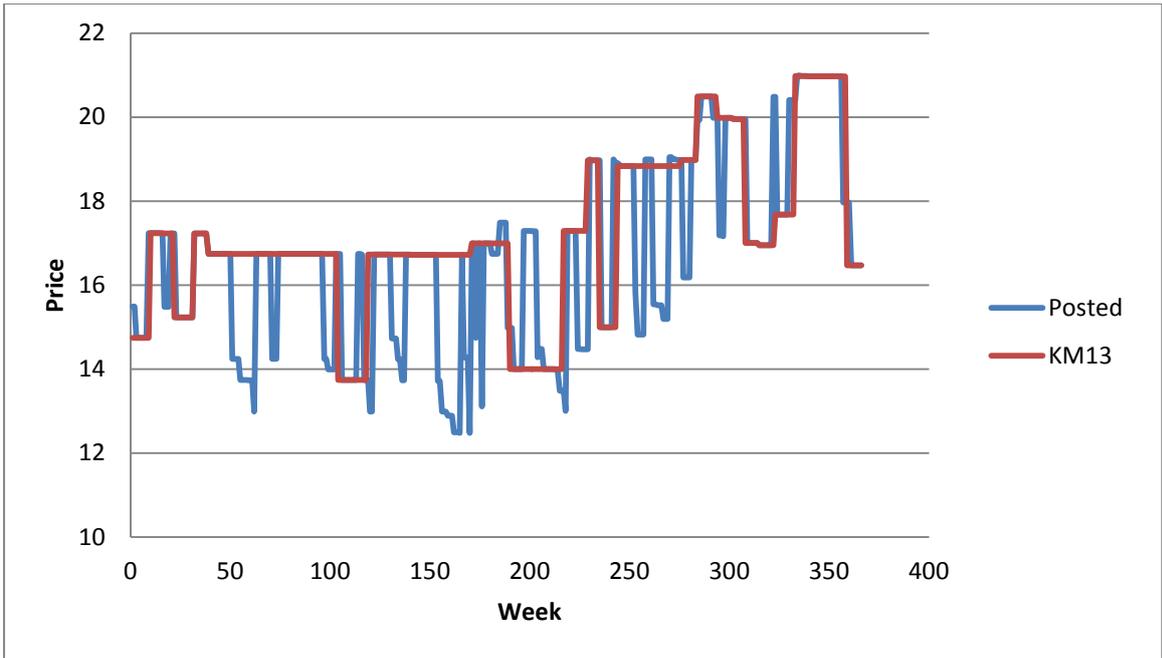


Figure A2: Posted and KM13 regular prices relating to an alcoholic drink

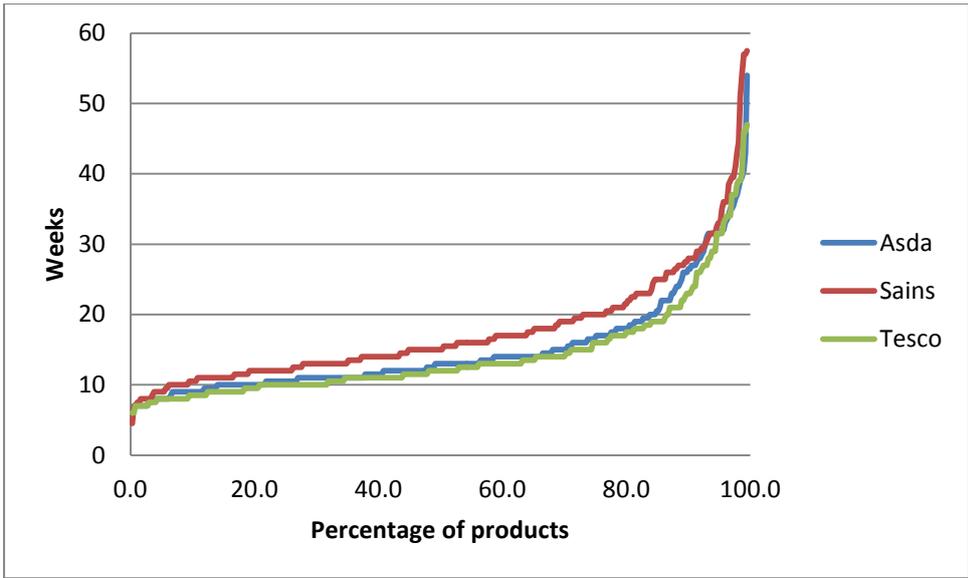


Figure A3: Median duration of price curve using KM11 regular prices.

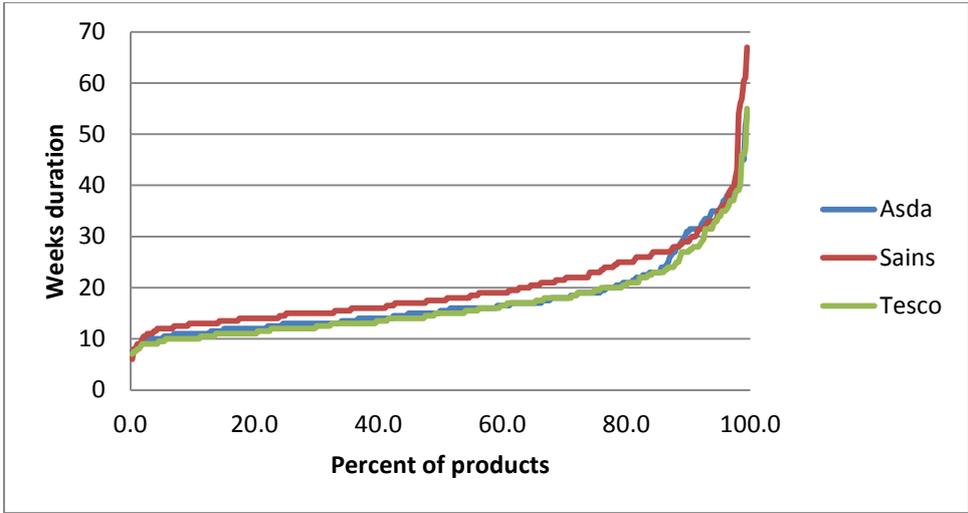


Figure A4: Distribution of median duration of KM13 regular prices across our sample of products

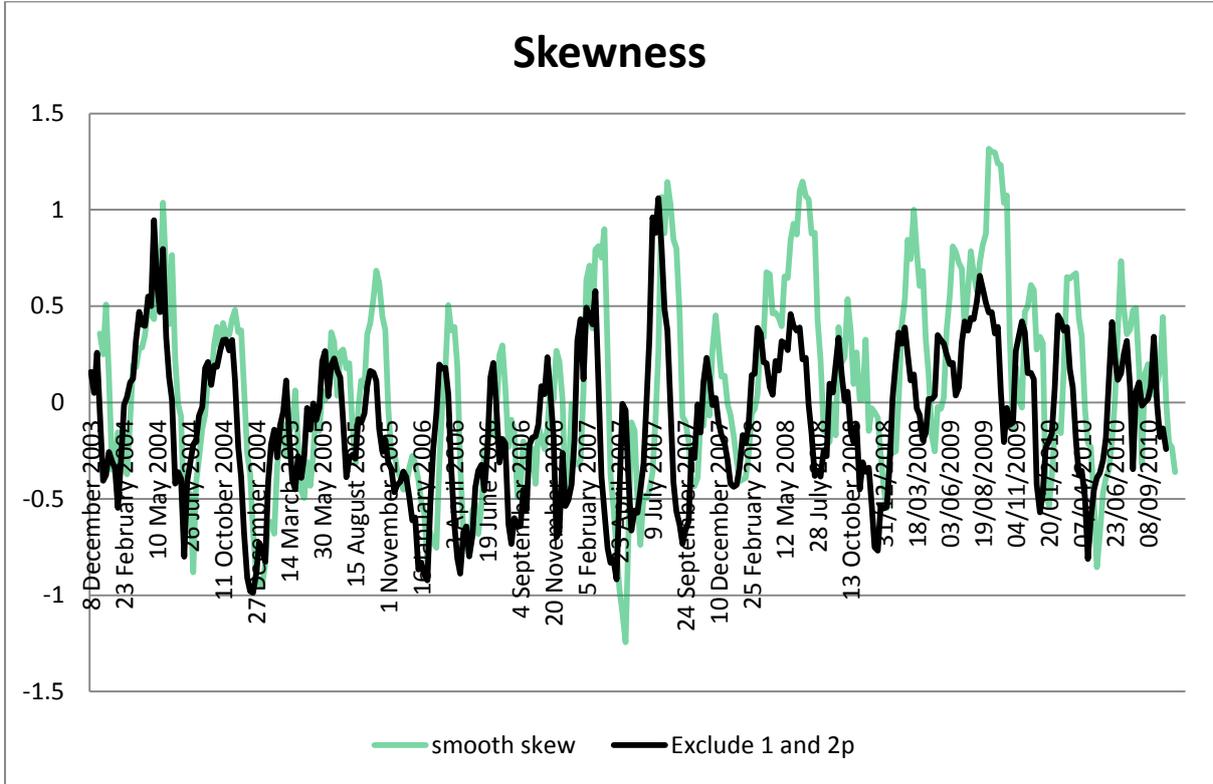


Figure A5: The smoothed Skewness of growth in NSB prices including and excluding small price changes.

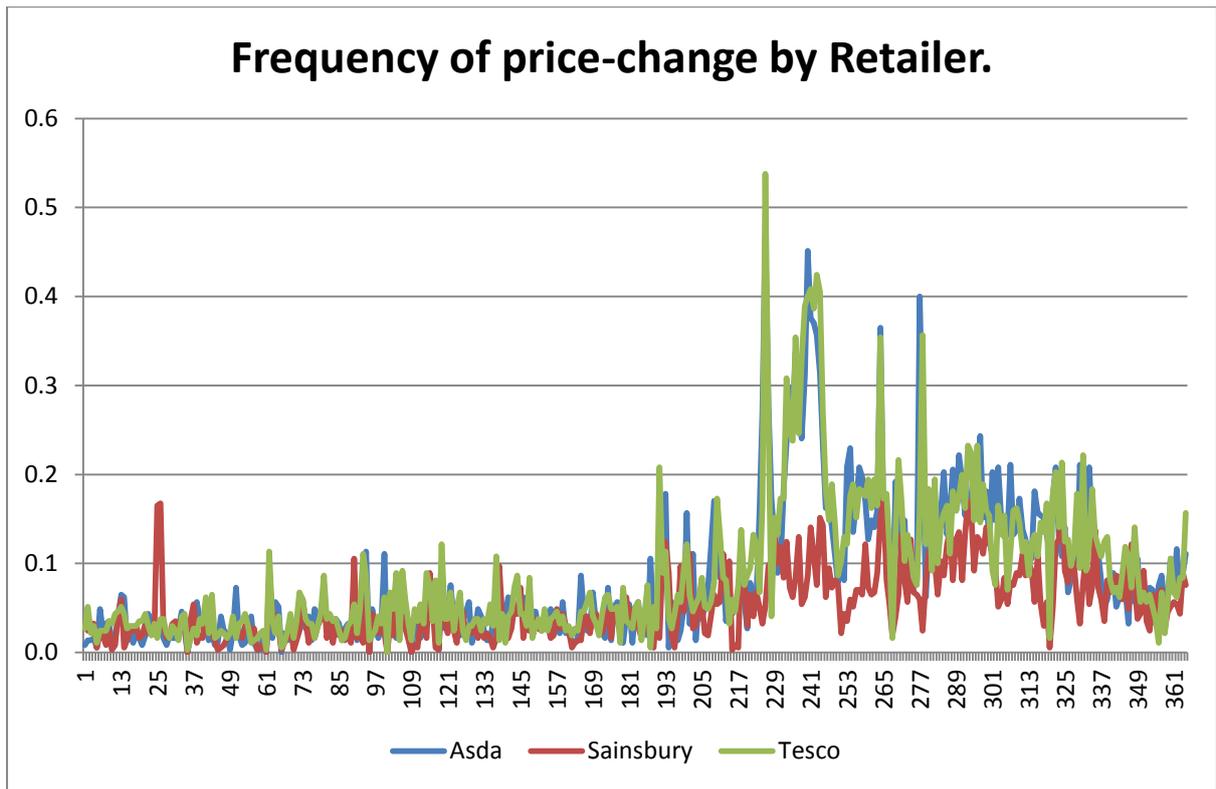


Figure A6: Frequency of price-change by retailer.