The hazard function of sales: An analysis of UK supermarket food prices

Lan, H., T.A. Lloyd and C.W. Morgan*

University of Nottingham, UK

Contributed Paper prepared for presentation at the 87th Annual Conference of the Agricultural Economics Society, University of Warwick, United Kingdom

8 - 10 April 2013

*The authors are research student, Associate Professor and Professor respectively at the School of Economics, Sir Clive Granger Building, University Park, University of Nottingham, NG7 2 RD, UK. Email tim.lloyd@nottingham.ac.uk for correspondence.

Abstract

In this paper we examine the empirical pattern of sales behaviour among the UK's seven largest retail chains using a scanner dataset of weekly food prices on over 500 products over a 2.5 year period. Motivating the analysis is the question 'are products more likely to go on sale that longer they remain unpromoted?'. Theory is not unanimous and empirical and recent empirical studies also offer conflicting evidence. To address the question we estimate the hazard rate of a sale - probability that a product goes on sale in the tth week since the last sale - over the market as a whole and then separately across different national retailers. We pay particular attention to the effects of sales in like-for-like products in rival retailers on the hazard of a sale. We also find that accounting for multiple sales has a pivotal role in determining the slope of the hazard function, which actually reverses sign when proper account is taken of this seemingly innocuous technicality. Correcting for this we find that food products are more likely to be discounted the longer they remain without a sale. This result helps square the circle between price setting and modern theories of sales behaviour. Furthermore, we find that the positive time-dependent pattern varies across product format and brand status. With sales in rivals, branded products in a representative retailer are more likely to be discounted if it has been on sale previously in the rival retailers, however the hazard of a sale in private labels is unrelated to its rival sales. In the individual retailer level, the hazard results show that while most supermarkets exhibit some form of a 'hi-lo' pricing there is one retail chain does not (showing no time-dependence) preferring an every day low pricing strategy (EDLP).

Key words: sales, the hazard function, multiple sales

JEL classification: L16; L66; E30.

1 Introduction

Sales (price promotions) are an important element of pricing and none more so than in UK food retailing where almost 40% of consumer retail food expenditure on food and drink is on products that are on promotion (Nielsen 2011). Data from the UK also show that 40% of the annual variation in retail prices are accounted for by sales (Lloyd et al., 2011). In the US market, the figure is between 20% and 50% (Hosken and Reiffen, 2004); the median frequency of price changes including sales is roughly twice the corresponding frequency excluding sales (Nakamura and Steinsson, 2008).

Sales are used as a strategic tool for attracting customers from rivals as part of the profit maximising process and as an essential tool of new product introduction (examples are seen in Bass, 1980 and DeGraba, 2006). In the theory of sales, models may be characterised in to three categories: static, dynamic and state-dependent. Correspondingly, they give different predictions of sales behaviour by retailers: sales are randomly distributed in to price competition models, time-dependent in dynamic models and related to specific characteristics such as perishability and brand (private-label products) in state dependent models. In this paper, our principal focus is to evaluate whether sales in UK food retailing are time dependent and the form of such time dependence. Current empirical evidence from the US does not offer a consistent story in favour of one or the other (see Pesendorfer, 2002 and Berck et al., 2008).

Microeconomic behaviour has macroeconomic implications. While not our focus it is interesting to note that amidst the current recession in the UK, British supermarkets have been keen to stress the important role of price promotions in keeping UK food (and in turn general) inflation down. While dynamic theories of sales predict that sales will be more likely the longer the non-sale regular price remains, a large and growing body of literature in the macroeconomic literature suggests that the longer a price remains the *less* likely it is to change (Nakamura and Steinsson, 2008). While this finding relates to price changes *per se* and not simply those that are terminated by sales, one might expect similar responses given the importance of sales to price changes (see Nakamura and Zerom, 2010; Guimaraes and Sheedy, 2011). The macroeconomic literature also suggests a staggered pattern of price changes in multi-product sellers whereby . . . . (seeLach and Tsiddon, 1996, Loy and Weiss, 2002). Retail promotion pricing may follow a similarly staggered pattern due to strategic substitutes of sales (Guimaraes and Sheedy, 2011, Nakamura and Steinsson, 2012).

We also seek to investigate whether the relationship between the duration of regular (i.e. non-sale) price spells are invariant to (state-dependent) factors such as retail chain, product perishability and brand status (national brands and private label products). In an economy with low inflation the timing of price changes is dominated by idiosyncratic shocks, implying that the potential role of heterogeneity complicates the frequency of price changes (Nakamura and Steinsson, 2012). For example, evidence shows that branded products are more likely to cause later sales of other brands compared to than private labels (Berck et al., 2008). Results from analysis at such a micro level of aggregation offers insights in to the strategic use of sales in the UK market, one in which a handful of national retail chains dominate (CR4=76%, Kantar WorldPanel, 2013). The analysis also sheds light on the pricing strategies such as Everyday Low Price (EDLP), Promotional Hi-Lo Pricing and Hybrid EDLP/Hi-Lo summarised in current marketing literature (see Ellickson and Misra 2008).

To evaluate these aspects of sales behaviour we conduct an empirical analysis based on extensive sample of food prices acquired from the AC Nielsen *Scantrack* database. The sample contains information on over 500 barcode-specific products in up to seven national retailers at 137 weekly intervals during a 2.5 year (2001-4) period, giving weekly time series on over 1700 unique product codes (UPC)[1] and 231,000 price observations in all. Using these data we calculate how long regular price spells last by estimating the hazard function of sales, a principal attraction of which is that it explicitly recognises that price promotions are not evenly distributed across products: some products are promoted frequently while others rarely. It transpires that the key result of this paper, concerning the likelihood of sale, turns on this issue (which we call

[1]Each UPC represents the product's barcode identity and the supermarket chain that is was purchased in

multiple sales). More specifically, we find that the probability of a sale is, in the main, time dependent. Whether the probability rises or falls with the duration of non-sale prices depends on whether the effect of multiple sales are taken in to account. Controlling for multiple sales we find that the longer a product remains without a sale, the greater the likelihood of it being promoted, a result that is consistent with the dynamic theory of sales of Pesendorfer (2002). Accounting for multiple sales helps square the circle between price setting and modern theories of sales behaviour. Furthermore, we find that positive time-dependence varies across various classifications of the dataset such as product format, brand status and most markedly by retail chain. While most supermarkets exhibit some form of a 'hi-lo' pricing there is one retail chain that does not. Showing no time-dependence its hazard function reveals an every day low pricing (EDLP) strategy. We also find that for branded products the hazard rate of sales is affected by sales of the like-product in rival retailers: Branded products are more likely to be on sale given they have been on sales in the rivals within the previous month. No such relationship is detected for private label products.

The paper is organized as follows. Section 2 reviews the theories of sales both theoretically and empirically. Section 3 describes the scanner price data and the regular price spells. Section 4 illustrates the methodology. Section 5 reports the empirical results and discussion. The final section concludes.

2 Literature Review

Theories of why retail firms may offer sales are found in Microeconomics literature. Considering market structure, firm behavior or consumer behavior, three main types of microeconomic models of sales are built up around a few key assumptions. In addition, retailers are considered to offer price promotions in order to attract enough customers to clear inventory (Pashigian and Bowen, 1991) or as an essential part of introducing a new product (Bass, 1980). Besides that, 'multiple sales' is a very important feature in the sales theories as sales are generally self-fulfilling and being strategic substitutes. Frequent sales on specific products tend to reduce consumers' sensitivity to price changes as they come to expect more sales in the future for these products (Richards, 2006). The timing of sales may also be staggered because a firm's incentive to have a sale is decreasing in the number of other firms having sales (Guimaraes and Sheedy, 2011). In the following, we mainly discuss the microeconomic models of sales and then to the link to multiple sales and the staggering of sales.

The early literature of the theory of sales is dominated by the static models of Varian (1980). By assuming that sales are a result of pure price competition amongst retailers, both models predict that sales are an event that are randomly distributed. Sobel (1984) develops the earlier contributions introducing a dynamic element in to the competitive process, and as a result predict that sales are strongly time-dependent and in particular more likely the longer the regular (non-sale) price persists. In a framework characterised by symmetric firms, there exists a duration of regular prices in which retailers charge high prices after a sale, only later cutting prices to sell to a large group of customers with a low reservation price. Since the demand of customers with low reservation prices accumulates steadily over time, the longer the regular price lasts the greater is the incentive to trigger a sale. This is time-dependent pricing model, sales behaviour is due to the inter-temporal price discrimination by supermarkets.

It had been long confused why static and dynamic time-dependent models offer different predictions of sales behaviour. Hosken and Reiffen (2007) augment the time-dependent framework with a state-dependent dimension of product perishability to explain the puzzle. According to them, state-dependence such as product perishability affects sales sales pattern and thus this type of models is categorised as dynamic state-dependent model of sales. For products that can be stored, inter-temporal optimisation of customers affects the retailer's sales behaviour, making sales being time-dependent. However, sales are random in the perishable products. Perishable products cannot stay over night so both retailer's pricing and consumer's buying decision should be made within the period and independent of time, which fits the static model assumption. Lal (1990) suggests that supermarket pricing behaviour may varies by brand status, positing that branded products are more likely to be promoted than private labels. It is mainly because price promotions are a long-run strategy pursued by national brands to defend their market shares from possible completion from private labels.

A number of paper show empirical analysis of sales related to supermarket pricing using micro-level data. In the US, Pesendorfer (2002) used a scanner dataset of daily prices and quantities for Ketchup products in Springfield Missouri between 1986 and 1988 and confirmed that the timing of ketchup sales is well explained by the number of time periods since the last sale, thereby emphasising that sales behaviour is positively dependent on the duration of the regular price. Berck et al. (2008) examine the behaviour of sales in the US retailing market using comprehensive dataset of retail prices covering a broad range of goods over a long period of time. They find significant heterogeneity of perishability and brand status in the sales pattern and supermarket pricing in the disaggregate level. For the timing of sales, they reject the sales random hypothesis and their results are the exact opposite of Pesendorfer's (2002) prediction that the probability of a sale increases with the time since the last sale of various brands. Lloyd et al. (2011) investigate the empirical sales patterns using scanner price observations from UK food retailing market, highlighting the importance of the heterogeneity of product durability, brand statues but most particular retailers in the estimation of sales patterns.

Both microeconomic theoretical models and corresponding empirical evidence are useful to understand sales behaviour. However, few can explain or make a link to another important feature of sales: why retailers typically offer many promotions for one product over time. In empirical macroeconomics, where a key issue is that of sticky prices a number of studies have estimated the hazard function of price changes and provide some useful insights. While an important contributor to price changes the occurrence of sales is not the primary focus of these but they are relevant they find that multiple price changes affect the hazard function of price changes significantly. Nakamura and Steinsson (2008) address that based on pooled data from many heterogeneous products, even if the hazard functions of all the goods are flat or upward sloping, heterogeneity in the level of the hazard function of different products can cause the estimated hazard function to be downward sloping. By examining the hazard function of price changes in food group, prices are typically found less likely to change the longer they persist when the 'multiple price changes' or heterogeneity bias in their words does not account for. Examples include Klenow and Kryvtsov (2008), Klenow and Malin (2010) for the US, Bunn and Ellis (2009) for the UK and a clutch of studies under the auspices of the Eurosystem Inflation Persistence Network (Alvarez et al., 2005 for Spain; Baumgartner et al., 2005 for Austria; Veronese et al., 2004 for Italy). Unlike them, Nakamura and Steinsson (2008) and Ikeda and Nishioka (2007) propose parametric methods that acknowledge uneven distribution of price changes, in that some product groups experience more frequent price changes than others. The former still find the negative time-dependence in the hazard function of price changes while the latter have exact opposite result: positive time-dependence in the hazard function. Besides that, Fougere et al. (2007) offer some evidence to the issue using an alternative methodology. They estimate the hazard patterns of price changes in the retailer-product level based on their huge price duration dataset in the French retailing market in order to eliminate the heterogeneity bias and find that patterns of price changes are either upward sloping or flat in such disaggregate level. Those results are very important to the multiple sales and standard sales theories. Following the literature, we will test static, time-dependent and state-dependent models for sales pattern empirically and particular consider the estimating of multiple sales.

The staggering of pricing strategy are treated as an another important dimension in sales behaviour. Retailers may have incentive to stagger the timing of sales (Nakamura and Steinsson, 2012). Do British retailers react to the sales of rivals, which may strongly related to the probability of sales in the food retailing market. Taylor (1979) shows that the staggering of prices between retailers causes aggregate price level inertia. Similarly, the timing of sales staggered or synchronised may increase potential sales in the entire retail chain. Guimaraes and Sheedy (2011) point out in their model that sales are being strategic substitutes so retailers may have incentive to avoid holding sales simultaneously and thus stagger the timing of sales. Empirically, Lach and Tsiddon (1996) find that the timing of the price changes is staggered across stores selling the same product and synchronised within the same store for different products sold using a store-level dataset of food prices in Israel. Similarly, Loy and Weiss (2002) find that the perfect staggering of price changes are present significantly across different sizes of retailers in German grocery food market. Similar evidence of price changes are also found in plenty of papers representing by Fisher and Konieczny (2000), Cavallo (2009) and Klenow and Malin (2010). Although the results are price changes *per se*, sales, mainly dominating price changes to some extent, may drive the staggering pattern of price change by sales staggering pattern. Based on a scanner dataset of weekly price observations for refrigerated and frozen orange juice in the US market, Berck et al. (2008) examine the staggering or synchronisation pattern of sales in rival retailers using Granger Causality test; they find that a sale of a major national brand is more likely to cause later sales of other brands than is a sale of a private label products in the US orange juice retailing market. Although based on a small market and traditional time series technique applied, their analysis is a beginning for further research based on bigger market and more panel data techniques.

3 The Data Description

3.1 The scanner dataset

The current research uses a high-frequency disaggregate-level scanner dataset purchased from Nielsen *Scantrack* and represents a comprehensive panel of 231,069 weekly UK supermarket prices on 507 food products in 15 categories of food spending over a 2.5-year (2001-4) period[2]. The products in the sample are principally processed foods and beverages, from a range of various formats (tinned, ambient, frozen, chilled and fresh) but exclude loose fruit and vegetables, ready-meals and uncooked meats, so while not fully representative of UK consumer spending on food, the sample includes products from a broad range of categories.

Three dimensions of the database are particularly interesting: products are evaluated at a highly specific – barcode - level; in up to seven of the UK's major food retailers (Tesco, Sainsbury, ASDA, Safeway, Somerfield, Kwik Save and Waitrose) and; include both branded and private-label products. Being based on EPOS data, prices are based on 100% of transactions in each supermarket and are thus national averages by retailer. Specifically, they are average revenue prices in that they represent the total value of transactions divided by the quantity purchased of the product in each supermarket over a week.[3] As such, they reflect promotional discounting (or sales) whether this is terms of a pure discount on price ('10% off') or quantity ('buy one, get-one free'). Each time series of prices is identified at the product-supermarket level by a unique product code (UPC) and there are 1,704 of these in the database. Since the presence of a sale is not explicitly recorded by Nielsen, Lloyd et al. (2011) apply a algorithm to indicate the presence of based on the sales profile of each UPC. They define a sale as a period of x% price decline of no longer than 12 weeks long, where x refers to sales thresholds of 10%, 25% and 35%. For the sake of brevity we report the results using the 10% sales indicator in this paper as the

[2]The 15 categories are orange juice, instant coffee, breakfast cereal, teabags, yoghurt, wrapped bread, tinned tuna, tinned tomatoes, tinned soup, corned beef, fish fingers, frozen peas, frozen chips, Jam and frozen pizza.

[3] For a detailed description and examination of the dataset see Lloyd et al. (2011).

results for deeper sales remain qualitatively unchanged using 25% and 35% sales indicators.[4]

Figure 1 illustrates the prices of just one of these UPCs (Del Monte Orange Juice Tetra 1L 3Pack in Safeway). Prices appear to switch between two states: non-sale (regular) prices and sale (discount) prices: shaded areas are referred to as the non-sale (regular) price period. Specifically, the price series appears to be punctuated by 6 sales over the sample and in the same way that sales are not necessarily of the same depth, the regular price appears to vary over the sample period too. Since sales occur at irregular intervals the duration of each regular price (i.e. the time between each sale) also varies: two of them are short, lasting only one week, others last between 10 and 20 weeks and one has a relatively long duration of some 63 weeks. In this paper, it is this duration of the regular price, and the factors that might account for it, that we seek to analyse.

[4] A full set of all results are available from the authors on request.

Figure 1: Price observations from a typical UPC

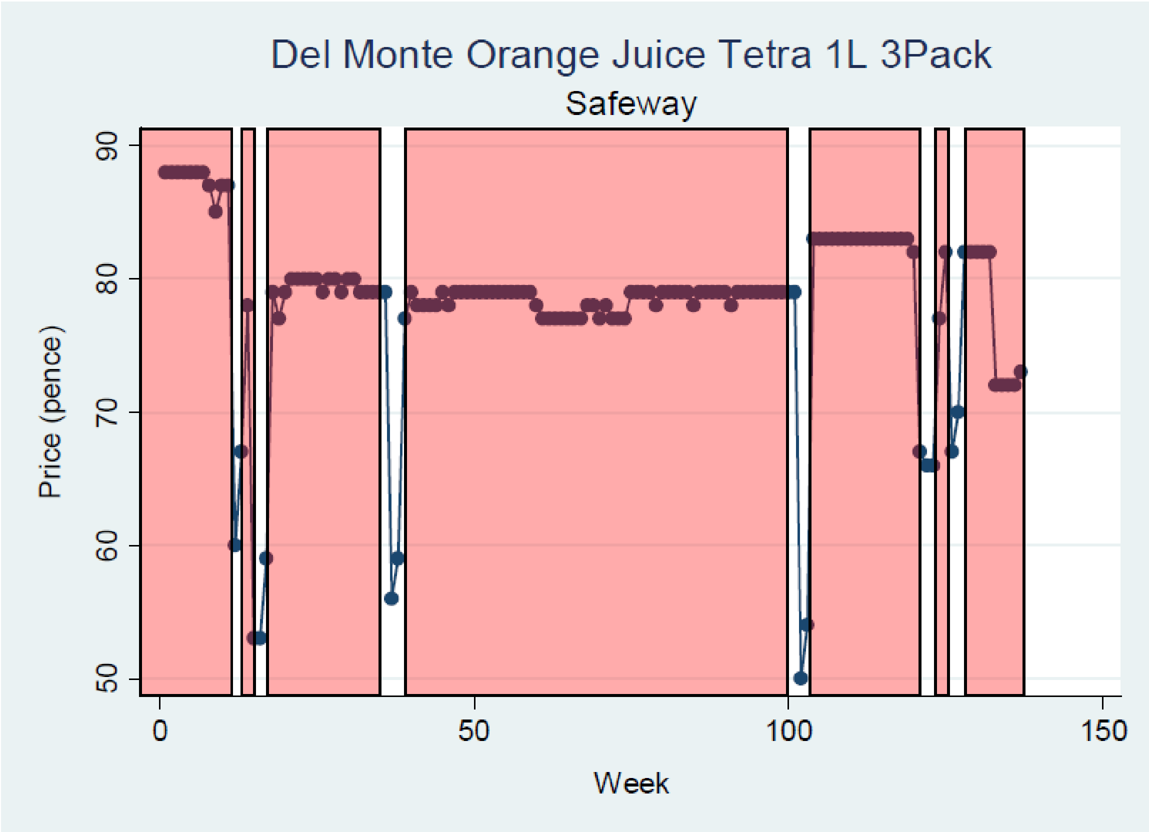

3.2 Censoring

Defining sales as a temporary decline in prices of at least 10%, there are 6,007 regular price spells in the sample. The histogram of these regular spells is given in Figure 2 which shows two important features of the duration data. First, for UPCs that have experienced a sale, the frequency of regular price spells declines with duration. Short-lived regular price spells are most common with around 18% of being less than three weeks in duration; around half are less that 15 weeks and 77% are less than a year. The Figure also indicates a significant proportion (almost 10%) of all spells are either 120 or 137 weeks in duration and these represent UPCs that did not experience a sale during the sample period.[5] As a result, the mean and median durations differ

[5]Around 90% of the UPCs are available for the full 137 week window, the remainder being 120 weeks long.

considerable being 35 and 15 weeks respectively.

Regular price spells from UPCs that were never on-sale during the sample period are commonly called *double-censored spells*; *left-censored spells* being regular price spells that started before the sampling frame began and were terminated within the sample and *right-censored spells* represent spells that began at a point within the sample but continue beyond the end of the sample frame. In contrast, *complete spells* are those that begin and end within the sample[6]. As reported in Table 1 55% of the spells are complete, 18% left-censored, 17% right-censored and 10% double-censored. The percentage of the left-censored spells is roughly equivalent to the right-censored spells indicating that they are evenly distributed in the data.

Figure 2: Histogram of regular price spells

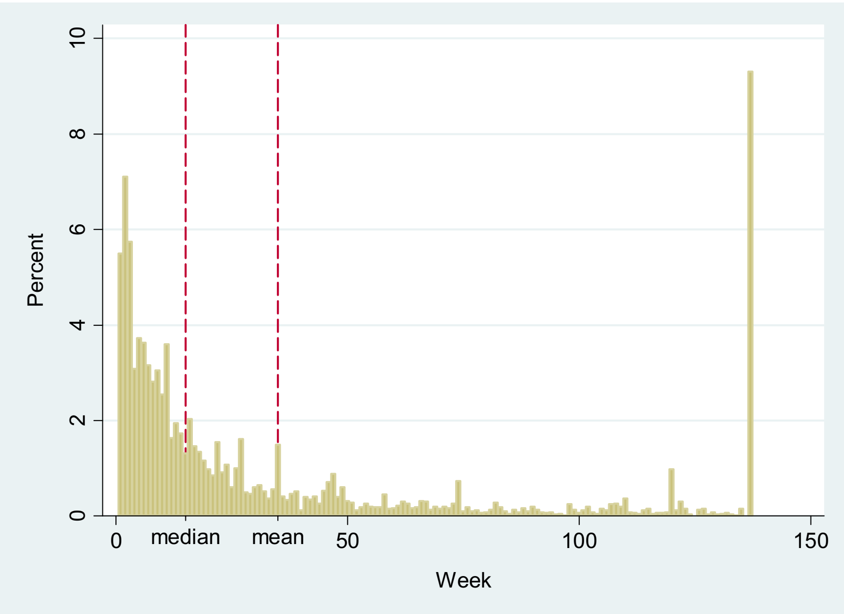

[6]One example of censored spells is shown in Appendix A.

Table 1: Distribution of complete and censored spells

	No. of spells	Percentage	Mean duration (Weeks)
Complete spells	3290	55%	16
left-censored spells	1094	18%	39
Right-censored spells	1013	17%	35
Double-censored spells	610	10%	136
Total	6007	100%	35

Since our interest is in the 'time to sale' and the factors that affect this duration, not all the regular price spells are relevant for the analysis. Specifically, double censored spells may be discarded since this represent the regular price spells of products that have never been on sale. Left-censored spells are also removed because the time at which the spell starts is not known (Nakamura and Steinsson, 2008; Fougere et al., 2007; Ikeda and Nishioka, 2007), so that the formal analysis is conducted on those that are either complete or right censored.[7] Table 2 reports the average duration of regular price spells for various classifications of the (censored) data. With the mean duration at 20 weeks twice that of the median, the skewed nature of the duration distribution is apparent. While this leptokurtosis is common to the classifications of the data in the table, it is also apparent that there is variation in the duration of regular prices within these classifications: private labels lasting longer than brands; frozen lasting long than chilled and most notable of all, regular prices in one retail chain (ASDA) lasting over 8 times as long as another (Safeway).

[7]As a robustness check, we have also estimated models using all the regular price spells (i.e. ignoring the censoring issue) and find that results are qualitatively similar, albeit that the hazard function is shifted down owing to the inclusion of the (longer) double censored spells. Results are available upon request.

Table 2: The duration of regular price spells

	Censored Data	
	Mean Duration (weeks)	Median Duration (weeks)
TOTAL	**20**	**10**
TESCO	24	15
SAINSBURY	29	17
ASDA	54	41
SAFEWAY	14	5
SOMERFIELD	17	9
KWIK SAVE	17	10
WAITROSE	26	16
TINNED	18	9
AMBIENT	20	11
FROZEN	21	13
CHILLED	21	9
FRESH	24	12
BRAND	20	10
PRIVATE LABEL	26	14

Note: 'Censored Data' denotes that the dataset contains complete and right censored spells only.

3.3 Heterogeneity in retailers, formats and brands

As the disaggregate statistics in Table 2 suggest there is considerable heterogeneity in regular price duration - and thus the use of sales - by retail chain. To give a flavour of this heterogeneity consider the histograms of regular price duration by retail chain, product format and brand status shown in Figures 3, 4 and 5 respectively. Differences by retail chain are most apparent: short-lived regular price durations (frequent sales) being most conspicuous in Safeway in contrast to the almost uniform distribution of ASDA. Other retailers fit in between with Tesco, Sainsbury and Waitrose (the more mainstream retailers) resembling ASDA whereas Somerfield and Kwik Save (both soft discounters) more like Safeway. In contrast to the distributions by retailer, there is little difference across product format (Figure 4), although brands appear to have shorter regular price duration (more frequent sales) than private labels (Figure 5). The picture that emerges is that while there may be important differences between brands and private label product, as far as UK food retailing is concerned, the use of sales appears principally determined by the retailers' fascia.

3.4 Multiple sales

As events go, sales happen frequently. In fact, using the 10% sales threshold, while only 8% of prices are actually sale prices, over two-thirds of UPCs experience at least one sale. Hence while sales are something of an exception to the normal rule of pricing, they are broadly applied and it is this that gives rise to the widely held perception that price promotions are commonplace in UK retailing. In a recent survey of food shoppers 81% of food shopper agree that they 'constantly look for the best prices and promotions' (Nielsen Homescan Survey of Great Britain, 2011). However, as the data in Figure 6 makes clear, their use and thus the distribution of regular price spells by UPC, is far from uniform: While the modal number of regular price spells is four (a feature that around 11% of the UPCs have in common) there are some UPCs that exhibit sales very frequently and thus many regular price spells, the largest number being 29. Frequent sales on specific products tend to reduce consumers' sensitivity to price changes as they come to expect more sales in the future for these products (Richards, 2006). Thus, for some (i.e. frequently promoted) products at least, sales are self-fulfilling in that the presence of a sale makes it even more likely that further additional promotional activity will be forthcoming. For this reason, treating multiple sales as independent events unrelated to the UPC in question is unlikely to be appropriate when estimating the time to sale.

Figure 3: Histogram of regular price spell duration by retailer

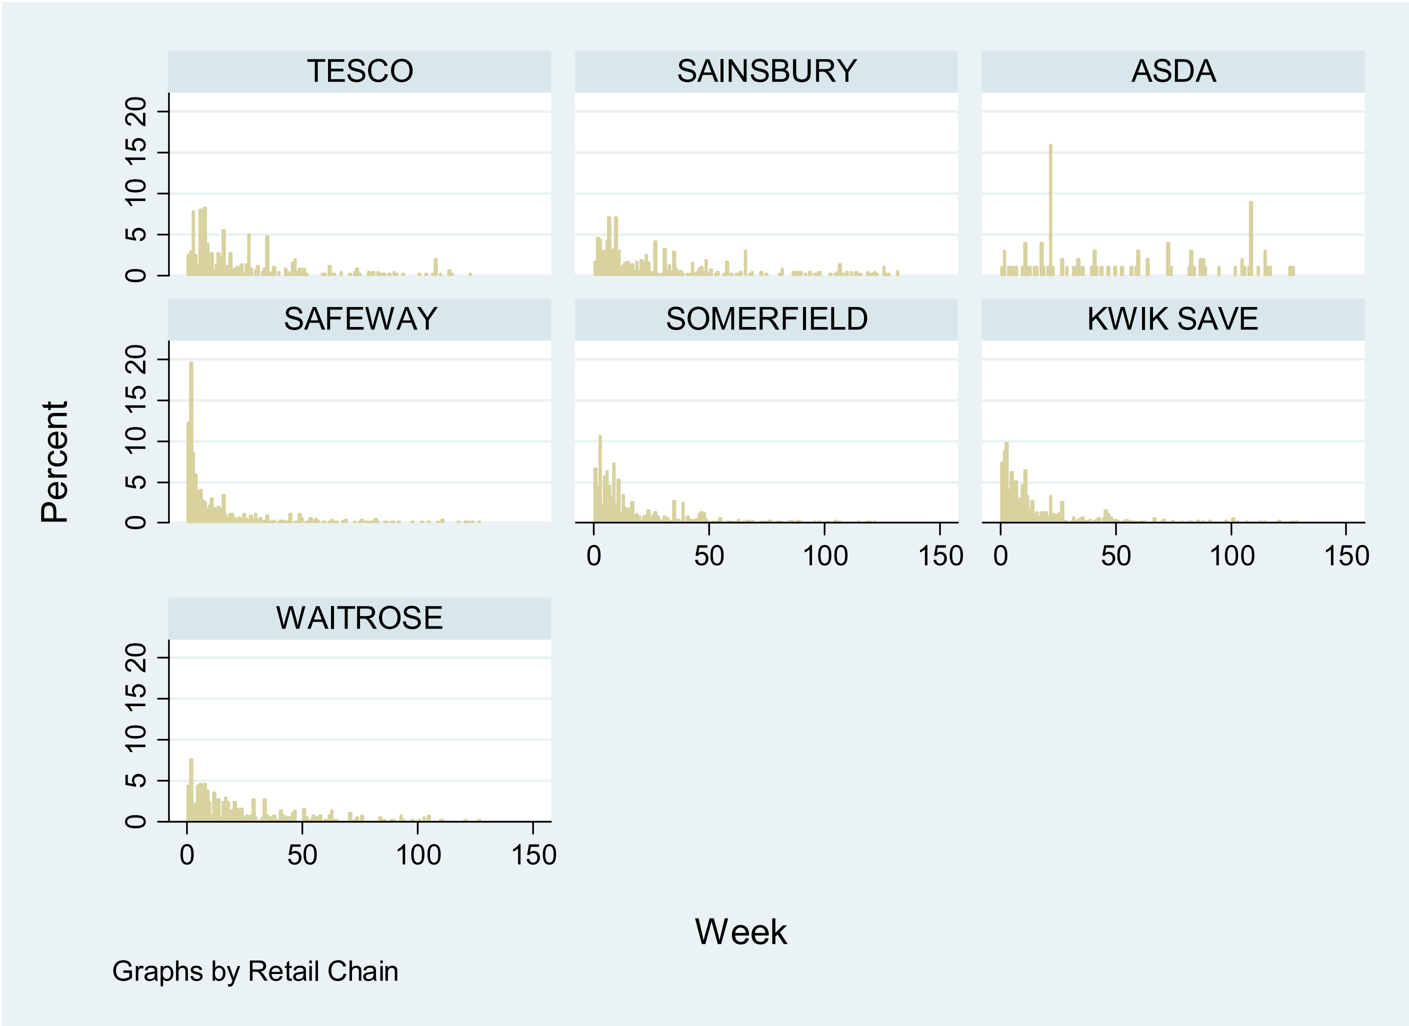

Figure 4: Histogram of regular price spells across format

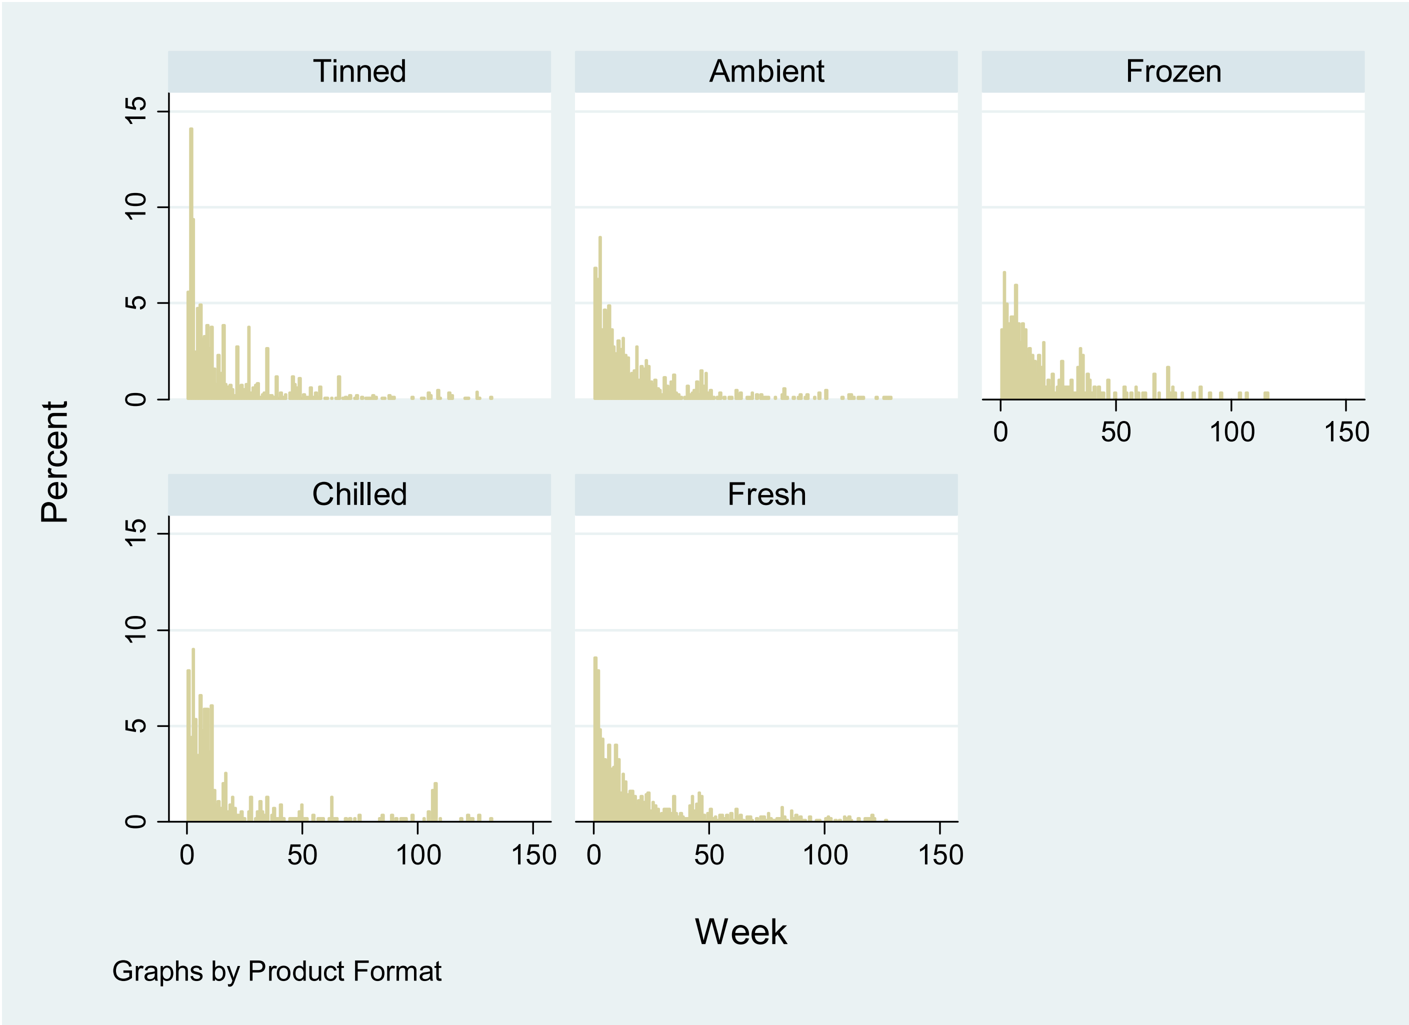

Figure 5: Histogram of regular price spells across brands

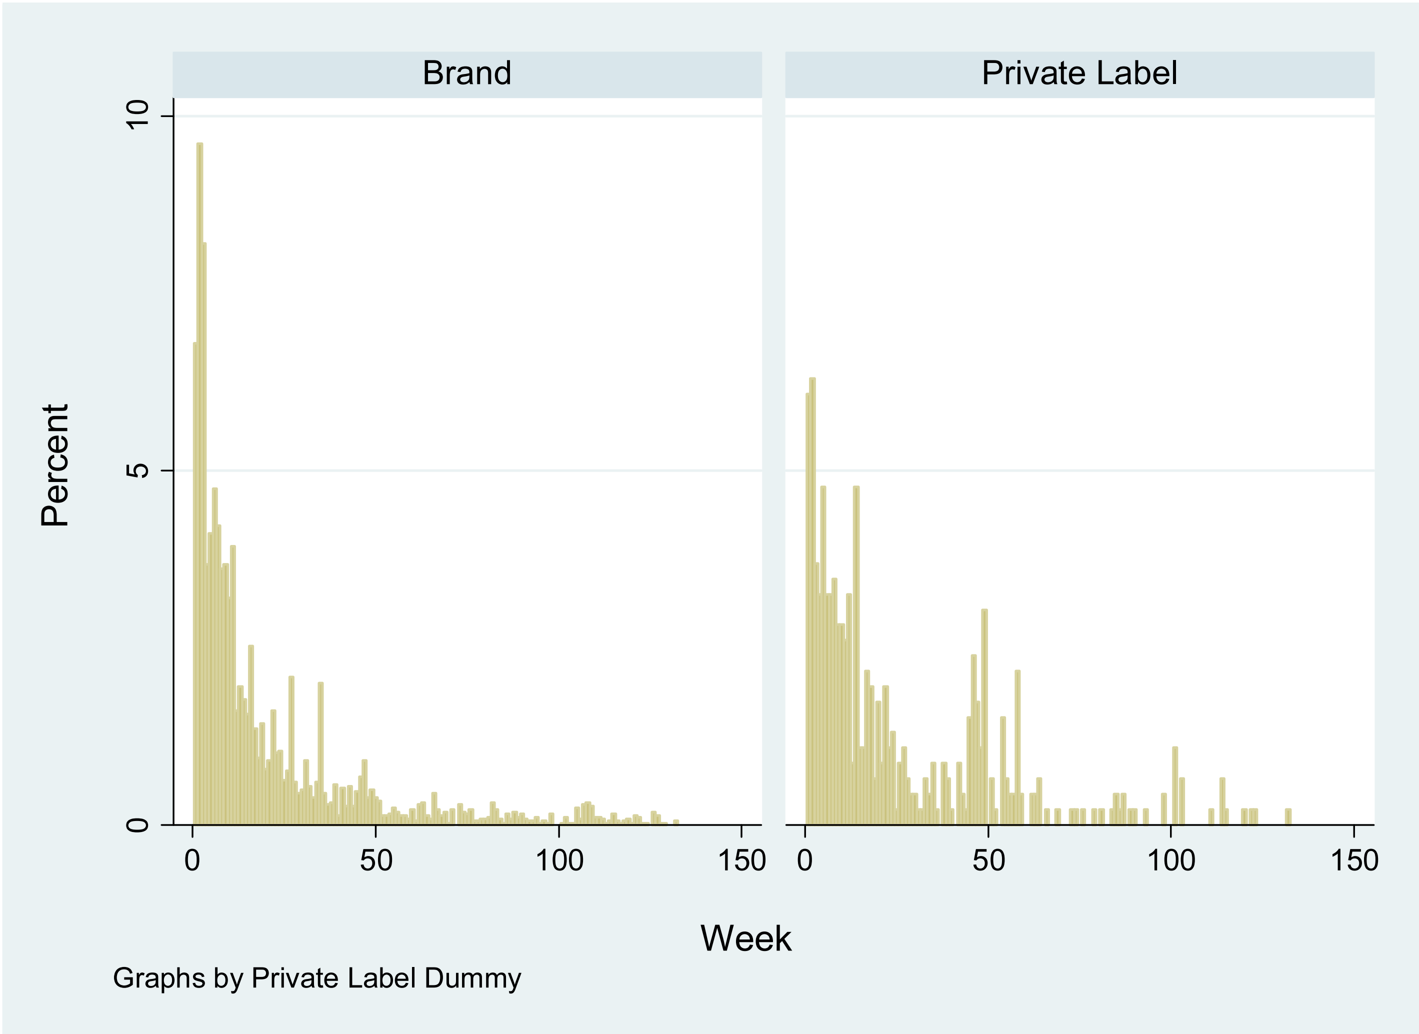

Figure 6: Distribution of regular price spells

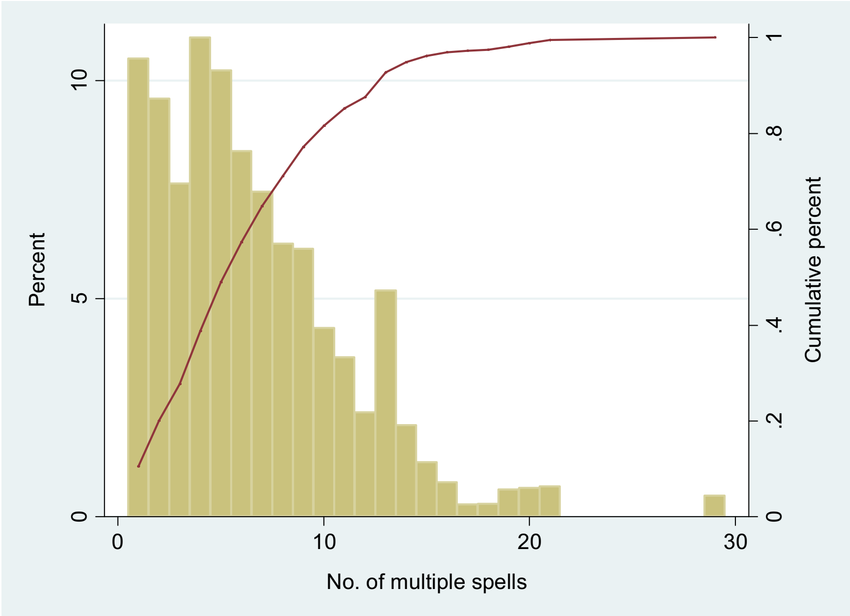

4 The Methodology

Having sketched the key features of the regular price spells we now undertake a formal statistical investigation using Duration Analysis. Originating in biomedical science where the duration of post-operative patient life is a popular application of the technique (e.g. Kalbfleisch and Prentice, 2002) duration analysis has been applied in economics to investigate a number of topics where the time to the occurrence of an event is a quantity of interest. Applications include the duration of spells of unemployment (e.g. Kiefer, 1988, Meyer, 1990 and Jenkins, 2005) and following the advent of large, highly detailed datatsets developed for the calculation of price inflation, the duration of prices (e.g. Dias et al., 2004, Alvarez et al., 2005, Fougere et al., 2007, Ikeda and Nishioka, 2007, Nakamura and Steinsson, 2008 and Bunn and Ellis, 2009). In our dataset of prices each spell of regular (i.e. non-sale) prices is terminated by the onset of a sale, and it is the duration of such spells of regular prices that is the focus of the analysis here. This is in contrast to the works cited above which examine the duration between price changes *per se,* irrespective of the cause. To be clear, we are neither interested in determining the duration of sales themselves (which are typically four weeks long) nor the duration of prices terminated by factors other than sales (such as changes in costs, inter-retailer competition or rounding errors) but the time between sales, what we refer to as the duration of a regular price spell.

4.1 Duration analysis

At the heart of duration analysis is the *hazard function* which, in the current application, models the rate at which sales occur as the duration of the regular spell increases. Rather than estimate the length of time to a sale directly, the hazard function models the rate at which sales occur as a probability. While the former (time-metric) formulation may be legitimately estimated using in some special cases, the 'time to an event' (time to a sale in the current application) is unlikely to be anything close to normally distributed in empirical settings (see Cleves et al., 2008).[8] In contrast, the hazard function actually incorporates the distribution into its functional form, so that providing a sufficiently flexible form is chosen, the problem is overcome. Moreover, the hazard function offers a convenient way to interpret the process that generates the termination of regular price spells by estimating the impact of factors that extend or reduce the duration as a probability that a sale will occur, given that the regular price has persisted for a certain amount of time.

Let T be a non-negative random variable measuring the duration of a regular price spell of length t, with density function $f(t) = \lim_{\Delta \to 0} P(t < T < t + \Delta)/\Delta$ that defines the probability of a sale terminating regular price spells of length T . Accumulating those probabilities over all spell durations gives the cumulative density function $F(t) = P(T < t)$, which defines the probability that regular price spells that have lasted up to t periods will have terminated by a sale. It then follows that the probability that regular prices persist longer that t, what is called the *survival*

[8]For example, where the probability of a sale is constant (say 10%) over time as in the Varian model, time to a sale is exponentially distributed.

function is $S(t) = 1 - F(t)$, i.e. the cumulative probability that the regular price spells will last beyond t .

Given the preceding statements we can define the hazard rate of regular price spells $h(t)$ as the probability that the regular price spell of length t is terminated by a sale, given that it has lasted (survived) t periods since the previous sale. In other words the hazard (rate) can be interpreted as the probability that a product goes on sale in the t^{th} week since the last sale. Thus, it denotes both the *failure rate* of regular price spells and the *occurrence rate* of sales for spells of length t. The continuous hazard function is written as:

$$h(t) = \lim_{\Delta \to 0} \frac{P(t < T < t + \Delta | T > t)}{\Delta} = \frac{f(t)}{1 - F(t)} = \frac{f(t)}{S(t)} \tag{1}$$

Unlike the density function describing the duration of regular prices $f(t)$, the hazard function defines the probability of a sale conditional on the regular price spell having lasted for t periods. In so doing, the hazard function recognises that while the (unconditional) probability of a sale may be say 5%, the probability of sale occurring say immediately after the previous one, or an arbitrarily long time after the previous one, may be considerably different from 5%. It is the conditional nature of the hazard that distinguishes it from the density function and gives it special appeal since it is this that is typically of interest. Notice that if the spell is right-censored, the hazard function is simply the survival rate of regular price spells at that duration, in which case $h(t) = P(T > t) = S(t)$.

A popular parameterisation of $h(t)$ is the proportional hazard model,

$$h_i(t) = h_0(t) \exp(\mathbf{x}_i \beta) \tag{2}$$

where here $h_0(t)$ represents a baseline hazard common to all i regular price spells as the duration increases; $\mathbf{x}_i$ is a vector of covariates used to control for the heterogeneity in the duration of the i^{th} spell and β is a vector of coefficients to be estimated with a sample of regular price spell durations. In this specification the exp() function ensures that the hazard is non-negative and thus the covariates have a proportional effect, shifting the baseline hazard up or down according to characteristics contained in $\mathbf{x}_i$. The attraction of the proportional hazard model lies in its flexibility (many functional forms can be accommodated), tractability (it is easily estimated by maximum likelihood in modern software) and because it allows state-dependent influences ($\mathbf{x}_i$) to augment a time-dependent baseline hazard, the latter being a proxy for a potentially large number of omitted variables that may be correlated with the passing of time (time itself not having any causal effect). Equation (2) is commonly estimated using semi-parametric methods developed by Cox (1972) and Cox (1975) whereby the parameters on the covariates of the model are estimated by (partial) maximum likelihood and then the baseline hazard function is recovered non-parametrically using Nelson-Aalen estimators (see Cleves et al., 2008). In the empirical results we report smoothed baseline hazard functions that are calculated using the Epanechnikov Kernel (see Stata 11.0 for details) and coefficients in the form of the hazard ratio (*i.e.* $\exp(\hat{\beta})$ rather than the estimated coefficient $\hat{\beta}$) which indicates the multiplicative effect on the hazard rate of a change in the covariate. Hence, if $\hat{\beta}$=0 the hazard ratio is equal to one ($\exp(\hat{\beta}) = 1$) indicating that a unit change in covariate has no effect on the hazard rate. Similarly, if the hazard ratio is greater than one ($\exp(\hat{\beta}) > 1$), the covariate increases the hazard by $\exp(\hat{\beta})$% *ceteris paribus*; when $\exp(\hat{\beta}) < 1$ and the hazard is reduced by $\exp(\hat{\beta})$% *ceteris paribus.*

In the current application we let the hazard vary by retailer, product format and brand status by including three sets of dummy variables in $\mathbf{x}_i\beta$ in (2). We also investigate the existence of multiple sales and sales synchronisation on the hazard of sales where the later is examined by evaluating whether the probability that a UPC goes on sale is affected by previous sales of the product in rival retailers. We explain how this is achieved in Section 5.

4.2 Accounting for Multiple Sales (Spells)

Unlike many (time-to-death) biomedical applications of duration analysis, a sale for a typical UPC is not a once-and-for-all occurrence. Indeed, the frequent use of sales mean that most UPCs experience multiple spells of regular prices. Moreover, the fact that marketing activity is unevenly distributed across UPCs (some being promoted a lot while others a little) represents an additional source of heterogeneity in duration data. While it is common in the literature to ignore this heterogeneity, a number of approaches have been proposed to tackle the issue head on. Nakamura and Steinsson (2008) treat multiple-spell effects as unobserved heterogeneity and use a random effects model to estimate the hazard function.[9] Fougere et al. (2007) exploit the vast dataset at their disposal and pick a sample of spells at random while Ikeda and Nishioka (2007) use a finite mixture model in which spells are clustered by product groups to provide weights for each UPC in the likelihood function. While appropriate for their setting, they are not well-suited to situations in which control variables are scarce, data is limited and tractability an advantage. In this paper we propose a simple alternative to these methods that exploits information in the *sales profile* of each UPC and thereby allows the hazard rate to vary with sales intensity. To fix the ideas that underpin the approach consider the following. In duration analysis, multiple spells are described as having *delayed entry*, in that while the current spell began at some point t_0, the UPC became at risk of being promoted prior to this date, such that $t_0 > t = 0$. The hazard h^d of a (complete) delayed entry spell is[10]

$$h^d(t_0 + t) = \frac{P(T = t_0 + t \mid T > t_0 + t)}{P(T > t_0)} = \frac{f(t_0 + t)/S(t_0 + t)}{S(t_0)} = \frac{h(t_0 + t)}{S(t_0)} \tag{3}$$

where t_0 denotes the starting point of the spell, t refers to the duration of the spell and $S(t_0)$ is the probability of having survived to t_0. (3) shows that $h^d(t_0 + t)$ is determined by the hazard rate $h(t_0 + t)$ and the survival function $S(t_0)$. By incorporating $S(t_0)$, the hazard rate of a multiple spell takes in to account one important facet that distinguishes it from a single spell, namely that the spell started some time after the onset of risk ($t = 0$). So while (3) says nothing about the order of the spell in the sequence or the number spells the UPC has experienced, it does recognise that the spell occurred with delayed entry, which is the case for any multiple spell.

[9] Similarly, Klenow and Kryvtsov (2008) allow for fixed effects by calculating the unconditional hazard rate for goods in different deciles of the distribution of regular price changes. It is a simple non-parametric approach though.

[10] In the case of a right-censored delayed entry spell, the hazard is written as: $h^{d'}(t_0 + t) = \frac{P(T > t_0 + t)}{P(T > t_0)} = \frac{S(t_0 + t)}{S(t_0)}$

If the spell is the first in a sequence of multiple spells, $t_0 = 0$ and $S(0) = P(T > 0) = 1$ so that the hazard for delayed entry collapses to that of a single spell, $h^d(t) = h(t)$. For all other spells in the sequence of multiple spells $t_0 \neq 0$ and the survival probability is less than unity, i.e. $S(t_0) < 1$. Moreover, as the number of spells increases so does t_0 and the chances of not being on sale $S(t_0)$ diminish. So, despite only accounting for the effect of delayed entry, the hazard function applicable to multiple spells differs from the single spell case; the implication being that treating multiple spells as a series of single (independent) spells will introduce a bias in the hazard function.

Before we introduce a simple method to allow for other aspects of multiple spells, consider a hypothetical UPC such as that depicted in Figure 9, which is on sale (denoted by the dotted lines) three times over sample period creating four regular price spells including one left-censored, one right-censored and two complete spells. Spell d_0 refers to the left-censored spell with duration $t = 1$ weeks; d_1 and d_2 are complete spells with $t = 2$ weeks each; and d_3 denotes the right-censored spell with $t = 4$ weeks. At the bottom of Figure 9 the hypothetical data is expressed in duration time. As is standard in duration analysis, the left-censored spell d_0 is discarded, as indeed are the sale periods themselves.[11] Being the first complete spell, $t_0 = 0$ for d_1; whereas $t_0 = 2$ for the second spell d_2 and $t_0 = 4$ for and the third spell d_3.

[11]Whereas left censored spells are discarded for lack of information on when they started, a UPC can only be at risk of a sale if currently not on sale, so all sale periods are excluded from duration time.

Figure 7: Effects of multiple spells: an example

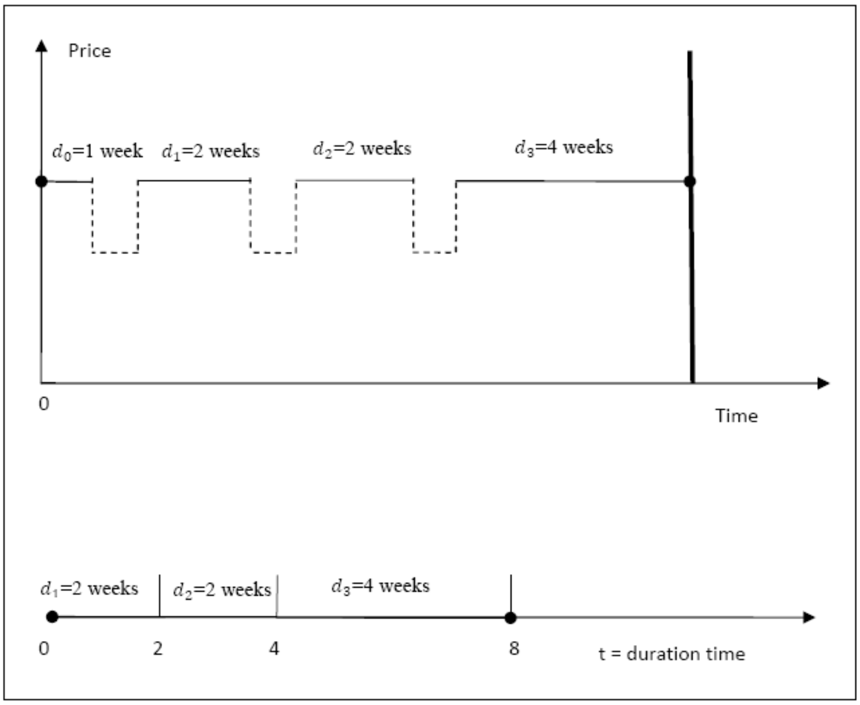

According to (3), the hazard rate of spell d_1 is

$$h^d(2) = \frac{P(T = 2)}{P(T > 2)} = \frac{f(2)}{S(2)}$$

since being the first in the sequence of multiple spells $t_0 = 0$ and so is merely the hazard of an independent spell of two week duration. In contrast, d_2 and d_3 are delayed entry spells with associated hazards given by

$$h^d(4) = \frac{P(T = 4 \mid T > 4)}{P(T > 2)} = \frac{h(4)}{S(2)}$$

and

$$h^d(8) = \frac{P(T > 8)}{P(T > 4)} = \frac{S(8)}{S(4)}$$

respectively. Note that being right-censored the hazard rate for d_3 is simply the ratio of survival functions. Table 3 summarises the hazard rate of these three spells when the effect of delayed entry is taken in to account and when it is ignored (as is the case when multiple spells are treated as a series of single spells). The effect is illustrated with a numerical example in which sales are assumed to occur at a constant rate of 5% per period, so that $f(2) = f(4) = 0.05$ with corresponding survival rates $S(2) = 0.9025$, $S(4) = 0.8145$, $S(8) = 0.6634$. Notice that the hazard rate for spells d_1 and d_3 with and without delayed entry are the same: for d_1 this is because it is the first spell in the sequence while for d_3 it reflects that the hazard for right censored spells it is merely the duration time that counts, and since these are the same for right censored spells, so is the hazard. For all intervening regular price spells such as d_2 the hazard rate with delayed entry is higher reflecting that the UPC has actually been 'at risk' from a sale for longer than the duration of the current regular price spell. Where such spells are numerous, as is the case when products are promoted frequently, the greater is the bias on the hazard caused by treating multiple spells as a sequence of independent single spells.

Table 3: Effects of multiple spells: Delayed entry

No.	Censoring	No delayed entry				Delayed entry			
		t_0	t	Hazard		t_0	t	Hazard	
				Analytic	Numerical			Analytic	Numerical
d_1	complete	0	2	$\frac{f(2)}{S(2)}$	0.0505	0	2	$\frac{f(2)}{S(2)}$	0.0505
d_2	complete	0	2	$\frac{f(2)}{S(2)}$	0.0505	2	2	$\frac{h(4)}{S(2)}$	0.0680
d_3	right-censored	0	4	$S(4)$	0.8145	4	4	$\frac{S(8)}{S(4)}$	0.8145

As noted previously, the delayed entry effect is only one, albeit important, aspect of sales heterogeneity, and ideally we wish to capture other information contained in a UPC's sales profile, such as the total number of sales it experiences ($Number_i$) and the order of each episode in the sequence of sales ($Order_{it}$), both of which may potentially affect the hazard rate of sales. Empirically, highly promoted products are something of a rarity and their regular price spells short-lived, so we expect $Number_i$ to increase and $Order_{it}$ to decrease the hazard rate of sales. To appreciate why consider Figure 8 which displays the sales profile of two hypothetical UPCs in which regular price spells of UPC1 and UPC2 are denoted as d_{1i} and d_{2i} respectively. Since UPC1 is promoted more frequently than UPC2 the hazard of spell d_{11} would be expected to be greater than d_{21}. If highly promoted products are rare, so are spells such as d_{12} which would be expected to have a lower hazard than more numerous spells of equivalent length (such as d_{11}). The value of the sales profile is thus apparent in estimation of the hazard function: without taking account of it, spells such as d_{11}, d_{12} and d_{21}would be predicted to have identical hazard rates.

Figure 8: Covariates: Example of total number of spells

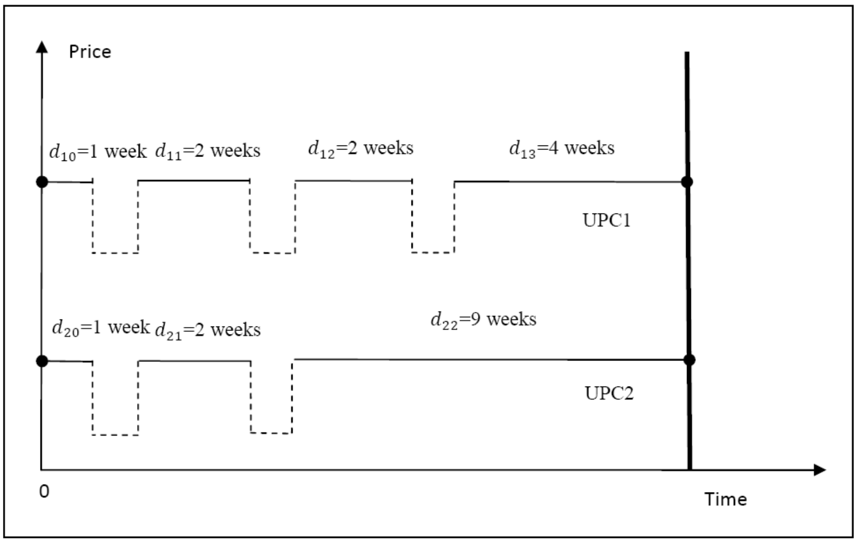

In setting-out our solution to the heterogeneity problem, the foregoing discussion underlines the importance of sales heterogeneity when multiple spell data is used to estimate the hazard function of sales. Given that most spells of regular prices relate to products that have experienced a number of sales during the sample period, they are numerous in the data and because sales behaviour tends to be UPC-specific, in that some UPCs are promoted frequently while others are not, ignoring the issue it likely to give rise to a potentially serious source of bias in the hazard function.

4.3 Rival sales

Sales sit at the heart of any marketing strategy designed to bolster market share. As a strategic device, sales reflect the competitive setting at both retail and upstream levels such that their timing, intensity and duration are the result of a potentially complex interaction between each retail chain and their suppliers. As such sales are unlikely to be random events but synchronised in some way. Whether this is the result of price-matching among rival retailers or to regulate manufacturing capacity, sales may be expected to be co-ordinated, all the more so given that nearly 40% of all goods sold by UK retailers are done so under promotion (Nielsen Homescan GB Survey 2011). Here we test whether the hazard of a sale is affected by recent sales of the like-product in rival retailers.[12] To do so we create a binary [1,0] dummy variable s_{ijt} which switches on to denote a sale price for product $i = 1, \cdots, I$, in retailer $j = 1, \cdots, J$ at time $t = 1, \cdots, T$ and is zero otherwise. Using the ordering of s_{ijt} a $K \times 1$ vector of dummy variables $\mathbf{r_{it}}$, $[\mathbf{r_{it}} = r1_{it}, r2_{it}, \cdots rK_{it}]$ is created to indicate whether product i is on sale in a rival retailer k $(k = 1, \cdots, K)$ at time t. Being matched with retailers in s_{ijt} each variable in $\mathbf{r_{it}}$ refers to the rivals of supermarket j, such that $rk_{it} = 0$ for $k = j$ (as k cannot be a rival for itself). Note also that $rk_{it} = 0$ when the product is either not stocked by rival retailer k or stocked but not 'on sale'. Observations where $rk_{it} = 1$ signal that product i is stocked in retailer j and k and on sale in retailer k $(k \neq j)$. For example, when considering sales in supermarket 1 (i.e.s_{i1t}) , $r1_{it} = 0$ since it cannot be a rival for itself. Hence $r1_{it} = 1$ when considering a product i in at time t that is on sale in supermarket 1 and stocked by supermarket $j = 2, \cdots, J$. Summing over all rivals in the market for the four previous weeks gives $Rival_{it}$ which is a [1,0] dummy variable indicating that product i has been on sale in the four weeks prior to period t in at least one rival, zero otherwise.[13] If recent sales of like-products in rival retailers make sales more (less) likely in period t, the hazard ratio of $Rival_{it}$ will be greater (less) than unity.

[12] The results reported in the following section relate to sales in the previous four week period. Models estimated for sales in the previous week and previous fortnight produce results that are qualitatively similar and are available upon request.

[13] In principle, it is possibile to examine the effect of sales in rival k on supermarket j however the small number of occurences make such bilateral tests of limited value empirically.

Berck et al. (2008) find that branded products are more synchronised than private labels and thus sales on branded products are triggered more than private labels by them being on sale previously . To examine such effect, we interact the rival sale dummy with the private label dummy. The new dummy, $(rival \times label)_{it}$, accounts for the impact on the hazard of a sale on private labels caused by their rival sales previously.

5 Empirical Results and Discussion

In this section we report findings from the estimation of the proportional hazards model introduced previously, namely

$$h_i(t) = h_0(t)\exp(\mathbf{x}_i\beta)$$

which is applied to data on the duration of regular price spells.[14] We begin by estimating a (single spell) model with retailer (Tesco, Sainsbury, ASDA, Safeway, Somerfield, Kwik Save, Waitrose), format (tinned, ambient, frozen chilled fresh) and brand status (private label=1) dummies, so that $\mathbf{x_i} = [\mathbf{retailer_i}, \mathbf{format_i}, label_i]$. In this model (Model 1) spells are treated as independent in that they take no account of the multiple spell problem. To assess the magnitude of heterogeneity bias in the single spell model we compare its baseline hazard $h_0(t)$ with that from (Model 2) a delayed entry (multiple spell) model in which $\mathbf{x_i} = [\mathbf{retailer_i}, \mathbf{format_i}, label_i, number_i, order_{it}]$. To assess whether recent sales in rival supermarkets impact differently on the likelihood of sales in branded and private labelled product we further augment the multiple spell model with rival retailer variable ($rival_{it}$) which we interact with the private label dummy (Model 3) giving $\mathbf{x_i} = [\mathbf{retailer_i}, \mathbf{format_i}, label_i, number_i, order_{it}, rival_{it}, (rival \times label)_{it}]$. Drilling down through the data we then finally estimate individual hazard functions for each retail chain in order to gain some insight in to the extent to which the UK supermarket sector segments by sales behaviour (Model 4). Since our focus is on the factors that affect the occurrence of sales, the duration data comprise complete and right-censored spells only; left-censored and double-censored spells being discarded, as discussed above.

5.1 Hazard Functions

To assess the role of multiple spells on the slope of the baseline hazard of sales and hence the

[14] The modelling reported in this section has been informed by an extensive non-parametric analysis by retailer, format and brand status. This and all other results are available from the authors upon request. We are however unable to release the data to third parties under the contract of the data acquistion agreement but all data has been made available to the referees in anonymised form.

heterogeneity bias of single spell hazard model consider Figures 9 and 10 which plot the baseline hazard rates in the single and multiple spell models (Models 1 and 2 respectively). [15] As is clear, taking account of multiple spells appears to matter. While both baseline hazards are time dependent, accounting for multiple spells reverses the slope of the baseline function for all but the longest lasting spells of regular prices. Two peaks approximately 50 weeks apart are also visible in the mulitple spell model possibly suggesting an annual cycle in sales behaviour. Note that after the second of the peaks, the hazard declines sharply suggesting that the likelihood of a product going on sale vanishes rapidly if it has not been on sale in the previous two years.

Figure 9: Baseline hazard function of sales in the single spell model (Model 1)

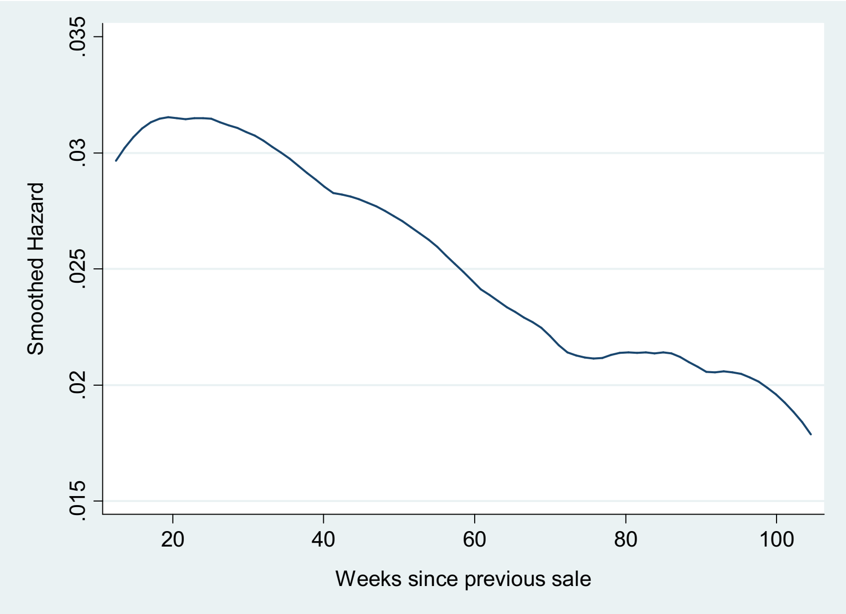

[15] A similarly upward sloping baseline hazard is obtained when sales profile covariates are excluded, suggesting it is the effect of delayed entry that is decisive.

Figure 10: Baseline hazard function of sales in the multiple spell model (Model 2)

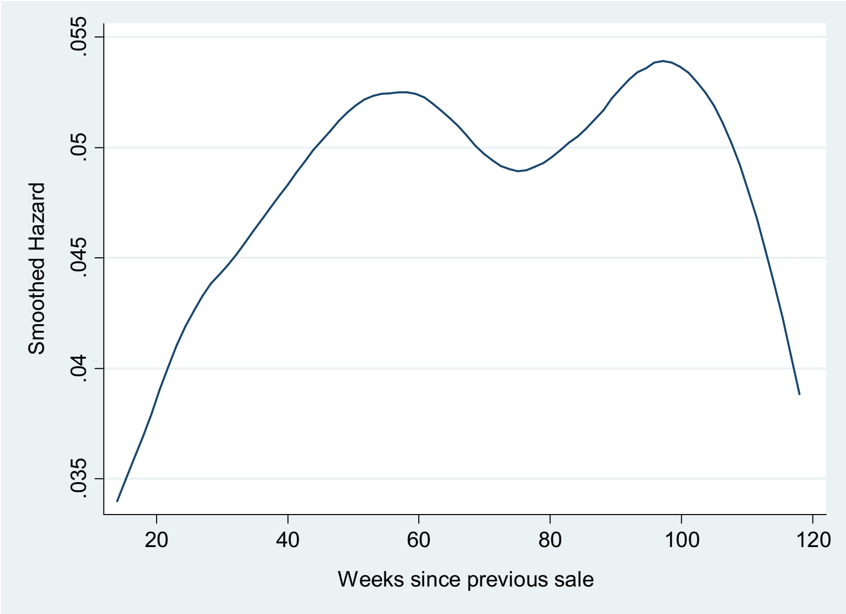

Before we consider the estimated coefficients from the multiple spell models, it is worthwhile considering why the treatment of multiple spells imparts such a dramatic impact on the hazard rate of sales. The reason is simple: short-lived regular price spells tend to cluster in frequently promoted products rather than being evenly distributed over UPCs. Since the single spell model does not control for the characteristics of a UPC's sales profile, as the duration time increases its baseline hazard reflects spells from infrequently promoted products, which by definition are unlikely to be promoted, thus the baseline falls as the duration of the regular price lengthens. In contrast, by controlling for each UPCs sales profile the multiple spell model has a baseline that more accurately reflects the underlying tendency in the data as a whole rather than merely the spells of infrequently promoted products. By doing so, the hazard rate of sales exhibits a positive slope, indicating that in general the longer the product is on the supermarket shelf, the more likely it is being on sale (positive time-dependence). Notice how in the weeks following a sale the baseline hazards from both models are similar: they only begin to differ as the heterogeneity in sales behaviour starts to take effect.

Results from the multiple spell hazard functions (Model 2 and 3) are reported in Table 4. Estimates of retailer, format and sales profile covariates are very similar in both regressions although we will focus on those of Model 2 initially. To ease interpretation, Tesco, the market leader is used as the base retailer, with tinned products and branded products representing the base categories for format and brand status, so that the hazard ratios reported in the table are relative to these comparators. Hazard ratios by retailer suggest that among all the national chains, only Safeway is more likely to have a sale than Tesco. Specifically, with a hazard ratio of 1.141, Safeway is estimated to be 14.1% more likely than the market leader to offer discounts, while other similar soft discounters (Somerfield and Kwik Save) are indistinguishable from Tesco. Both of the other mainstream retailers (Sainsbury and ASDA) have a significantly lower hazard rate of sales than Tesco; ASDA actually being 67.9% less likely to have a sale than Tesco, other things remaining equal. This corroborates inference from casual inspection of the data and underlines that not only does ASDA have more products that are never promoted, but for products that are, sales are used less than in other retailers. Note also that Waitrose, the only luxury food retailer of the seven in the sample, has a hazard rate of sales that is estimated to be marginally higher than the market leader, suggesting that 'hi-low' sales strategies are not solely adopted by retailers with a reputation for discounting.

Results by format also show some differences, albeit considerably less than that by retailer. Results are relative to tinned products and suggest that all formats except chilled have hazard rates lower than tinned products, fresh products being the lowest at 31.1%. While this is at odds with Hosken and Reiffen (2007) who propose a state-dependent model in which the likelihood of sales increases with perishability, note that since our sample does not include fresh fruit, vegetables and meat (indeed, the fresh format comprises bread exclusively) the result is likely to be driven by factors other than pure perishability. Of perhaps more importance is that tinned (and to some extent, frozen) products are more likely to be discounted than the other formats, a result that is consistent with arguments that supermarket promotions are skewed towards storable processed foods and drinks (Dobson and Gerstner, 2010). The results also confirm the

Table 4: Multiple spell proportional hazard models of sales

	Model 2		Model 3	
	Hazard Ratio	Robust Standard Error	Hazard Ratio	Robust Standard Error
$\mathbf{retailer_i}$				
SAINSBURY	0.871**	(0.054)	0.863**	(0.052)
ASDA	0.301***	(0.057)	0.321***	(0.061)
SAFEWAY	1.141	(0.107)	1.170*	(0.111)
SOMERFIELD	0.988	(0.060)	0.985	(0.058)
KWIK SAVE	1.078	(0.086)	1.035	(0.079)
WAITROSE	1.032	(0.067)	1.058	(0.068)
$\mathbf{format_i}$				
AMBIENT	0.833***	(0.051)	0.815***	(0.053)
FROZEN	0.906	(0.055)	0.965	(0.062)
CHILLED	0.763***	(0.042)	0.836***	(0.045)
FRESH	0.689***	(0.053)	0.781***	(0.059)
$label_i$	0.852***	(0.053)	1.001	(0.077)
$number_i$	1.200***	(0.021)	1.196***	(0.021)
$order_{it}$	0.971*	(0.017)	0.965*	(0.018)
$rival_{it}$	-	-	1.629***	(0.064)
$(rival \times label)_{it}$	-	-	0.912	(0.123)
N	4303		4303	
UPC	1,703		1,703	

Note: ***, ** and * denote that the null hypothesis (of unity in the hazard ratio) is rejected at the 1% , 5% and 10% level. Results are relative to base group which is tinned branded products stocked by Tesco.

theoretical prediction of Lal (1990), who posited that sales are less likely to occur in private labels than brands. The estimate in Table 4 suggests private labels are 15% less likely to be promoted compared to national brands.

Turning to the effect of sales profile covariates, results suggest that while the number of previous spells is positively related to the probability of future sales and highly (1% level) significant and effect of spell order is slightly negative and weakly (10% level) significant. As such they confirm that sales are indeed not randomly assigned but tend to be concentrated in certain UPCs (reflecting factors such as brand competition) and that as that he number of spells rises the chance that the current regular price spell will be terminated by a sale falls (reflecting that the frequently promoted products are more unusual).

As set out above, Model 3 assesses the effect of recent sales of like-products in rivals via the addition of a dummy variable $rival_{it}$, which is also interacted with brand status $(rival \times label)_{it}$ to test whether private labels exhibit the same response as brands. From Table 4 the hazard ratio on $rival_{it}$, which is statistically different from unity at the 1% level, indicates that branded products which have been on sale in the previous month in at least one rival retailer are 62.9% more likely to be promoted brands that have not. The staggering of sales implied by this result does not carry over to private labels however given that the the coefficient on $(rival \times label)_{it}$ is not statistically different from unity at conventional levels. Note also that unlike Model 2 the coefficient on $label_i$ in Model 3, is no longer different from unity suggesting that different in the propensity for sales by brand status arise solely because of the staggering of sales of branded products: sales are equally as likely in private label products (irrespective of recent sales in like-products) as they are for brands that have not recently experienced a sale in rival retailers. The results imply that while sales are staggered across UK retailers, this only true for branded (not private label) products, a result that is consistent with Berck et al. (2008) for the US.

5.2 Hazard models by retailer

One of the key findings from the previous section is that differences in the hazard of sales is most keenly affected by retailer. As Berck et al. (2008) addressed based on their empirical results, retailers rather than manufacturers determine sales. As a result we drill down in to the data and estimate a separate hazard function for each retail chain using the same covariates as in Model 3. Baseline multiple spell hazard functions for each retailer are graphed in Figure 11. While there is a general tendency of positive time dependence in the baseline hazard, as a group their differences are more striking than their similarities. Interestingly, the market appears to segment: Sainsbury mimics the baseline hazard of Tesco, as might be expected from the top two retail chains; Kwik Save and Somerfield also follow this pattern albeit at a higher level consistent with their soft discounter reputations; Waitrose, the luxury retailer, lies in between these pairs. The remaining retailers, Safeway and ASDA represent polar cases: Safeway's hazard lies at a high level, reflecting a manifestly 'hi-lo' sales strategy whereas the flat baseline hazard of ASDA is consistent with an EDLP pricing strategy.

Figure 11: Multiple spell proportional hazard models of sales by retailer

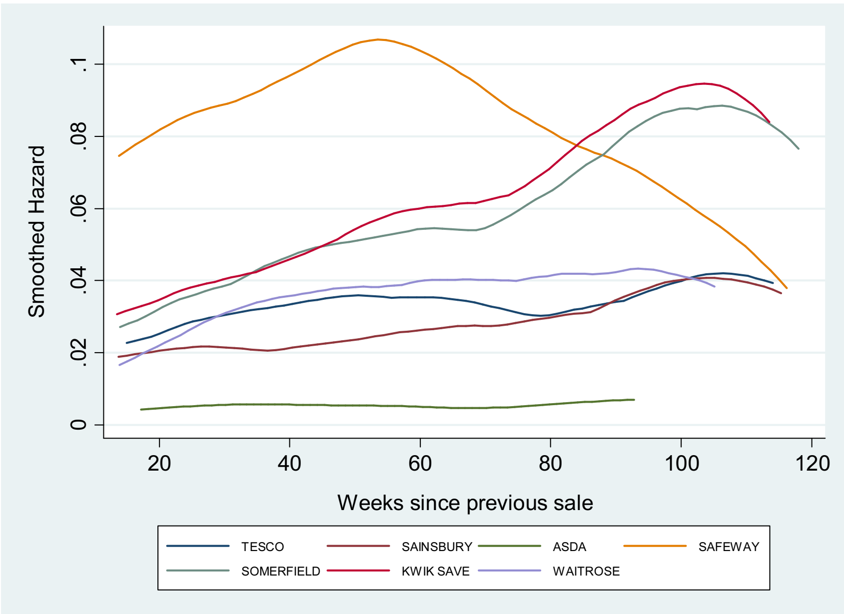

Table 5 reports estimates from the retailer hazard functions. There are two salient points to make. First, with reference to the effect of multiple spells, coefficient estimates on the two covariates for sales profile are consistently signed and statistically significant in almost all retailers, confirming the importance and rarity of frequently promoted products, a result that echoes that found in the analysis on the market as a whole.[16] Second, and in contrast to the findings on multiple spells, there is little unanimity by product format with respect to either the sign or statistical significance: some supermarkets (e.g. ASDA) show little difference across the formats, others (e.g. Safeway) a lot. There is now some support (from Sainsbury and Waitrose at least) for Hosken and Reiffen's prosed link between perishability and the frequency of sales, although contrary results in (Safeway and Somerfield).

Echoing the results from the aggregate, the private label dummy is generally not statistically significant in the retailer models, although even here there are idiosyncrasies: Safeway's intensive use of sales applying far more to brands than private labels whereas in Waitrose it is its own label products that are more promoted, possibly reflecting the high quality of its private label range. The staggering of sales for branded products is a feature that is common to most retailers, the coefficient on $rival_{it}$ being statistically greater than one at the 1% level in all except ASDA, which does not appear to respond to its rivals's sales behaviour, in-line with its EDLP strategy. As with the aggregate results the coefficient on $(rival \times label)_{it}$ is not significant statistically in most retailers.[17] At the retailer level too, branded products are more likely to be on sale if they have recently been on sale in a rival. The pattern is not shown in private labels.

In sum, the disaggregated analysis by retailer presents a richer picture reflecting the idiosyncratic nature of sales strategy in UK food retailing. Importantly for our focus, even the positive time dependence of the hazard function of sales does not apply to all retail chains. However, we do find one result that is common to all, namely that multiple spells matter. Put slightly differently, this means that all retail chains have some products that are more frequently promoted than others. While this is so obvious that it hardly needs to be stated, its effect on the principal instrument of duration analysis - the hazard function of sales - seems to have been overlooked. But it is not just a story of heterogeneity and there are some common tendencies in particular,

[16] ASDA's coefficients appear somewhat exaggerated, which most likely simply reflects its low baseline hazard.

[17] Note that the hazard ratio of $(rival \times label)_{it}$ is only based on four observations for Waitrose and none in ASDA. Data is more numerous in the other retailers: Specifically there are 7 non-zero observations for Tesco, 22 for Sainsbury, 41 for Safeway, 23 in Somerfield, and 21 in Kwik Save,

Table 5: Multiple spell hazard function of sales across supermarkets

	Model 4						
	TESCO	SAINSBURY	ASDA	SAFEWAY	SOMERFIELD	KWIK SAVE	WAITROSE
$\mathbf{format_i}$							
AMBIENT	1.339***	0.768**	0.728	0.503***	1.032	0.920	1.460**
	(0.150)	(0.093)	(0.401)	(0.049)	(0.100)	(0.122)	(0.241)
FROZEN	1.176	0.972	1.190	0.807*	1.315*	1.089	1.760***
	(0.205)	(0.142)	(0.686)	(0.094)	(0.215)	(0.141)	(0.316)
CHILLED	0.876	0.902	0.955	0.600***	1.485***	1.164	1.114
	(0.153)	(0.111)	(0.675)	(0.046)	(0.138)	(0.177)	(0.301)
FRESH	1.362**	1.356***	1.114	0.635***	0.943	0.941	1.847***
	(0.202)	(0.117)	(0.330)	(0.053)	(0.097)	(0.095)	(0.319)
$label_i$	0.987	1.151	2.069	1.115	0.448***	1.099	1.371*
	(0.165)	(0.184)	(1.496)	(0.147)	(0.121)	(0.174)	(0.225)
$number_i$	1.383***	1.630***	3.650***	1.178***	1.194***	1.252***	1.534***
	(0.040)	(0.066)	(0.657)	(0.018)	(0.038)	(0.025)	(0.089)
$order_{it}$	0.870***	0.737***	0.337***	1.041	0.883**	0.879***	0.765***
	(0.046)	(0.041)	(0.091)	(0.027)	(0.043)	(0.029)	(0.060)
$rival_{it}$	1.933***	1.585***	1.003	1.641***	1.324***	1.539***	1.977***
	(0.238)	(0.175)	(0.487)	(0.118)	(0.092)	(0.137)	(0.277)
$(rival \times label)_{it}$	0.814	0.682	-	1.091	1.563	0.518	0.456***
	(0.296)	(0.194)	-	(0.218)	(0.538)	(0.210)	(0.112)
N	527	600	101	1151	820	739	365

Note: Robust standard errors are reported in parentheses. ***, ** and * denote that the null hypothesis (of unity in the hazard ratio) is rejected at the 1% , 5% and 10% level respectively.

that most supermarkets promote their private label products less than brands, and most are more likely to place a product on sale as time passes.

6 Conclusion

In this paper, we examine the timing of sales and supermarket pricing in UK food retailing using a comprehensive dataset of weekly prices obtained via barcode scanners of purchases of over 500 products in seven national retailers during a recent two and half year period. Using semi-parametric approaches we estimate the hazard function of sales and find that, in the main, the pattern of sales is time-dependent. Importantly however, the shape of the hazard critically depends on the whether sales are treated as events that are independent of the product in question or not. Sales are not distributed uniformly over products so the distinction is empirically relevant. Relaxing the independence assumption, reverses the slope of the hazard function, so that when account is taken of sales profile of each product, the (conditional) hazard increases with time, meaning that products are more likely to be put on sales, the longer the regular price remains. This finding is consistent with the theoretical prediction of time-dependent models of sales proposed by Sobel (1984) and Pesendorfer (2002). While this result relates specifically to the probability of sales (rather than price changes *per se*) it is also similar to that of Ikeda and Nishioka (2007) who find that the prices of food products have an increasing hazard function. Thus it appears that the increasing hazard of price changes in general is in part attributable to the pressure on inflation causing increases in the regular price (Cavallo, 2009) and partly due to an increasing probability of sales.

The hazard of sales is also analysed by accounting for rival sales in the previous periods. The staggering pattern of sales are found in the UK food retailing chain, improving the estimates of the hazard function of sales in one aspect. It may interpret partially the high hazard of sales in the market. The branded products in a representative retailer are more likely to be discounted if it has been on sale previously in the rival retailers, however the hazard of a sale in private labels is unrelated to its rival sales. It implies that the staggering pattern of sales are only found in branded products rather than private labels, consistent with Berck et al. (2008).

Our analysis also addresses the issue of heterogeneity in terms of retailer, product format (perishability) and brand status (private label). It turns out that despite their national presence, the sales strategy of UK food retailers is, to a large extent, idiosyncratic. While all use sales selectively (in that in every retail chain, some products are more frequently promoted than others) and most promote brands more than private labels, there is little consistency by format. Most importantly, there are key differences in the shape of the conditional hazard function of sales across retailers, reflecting a wide range of sales strategies being used in UK food retailing. Retailers appear to occupy niches in the market whereby variants of 'hi-lo' and 'EDLP' sales strategies appear to co-exist. The fact that there is only one EDLP retailer, may simply reflect the ease with which the EDLP claim is to refute or verify. The results do however have implications for theory in that they offer little support for the representative firm in models of sales behaviour. While the Varian (1980) model offers a useful benchmark, we find for the UK that it's random sales pattern with symmetric retailers to be some way from the actuality that we observe.

References

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

Bass, F. M. (1980). The relationship between diffusion rates, experience curves, and demand elasticities for consumer durable technological innovations. *The Journal of Business 53*(3), S51–S67.

Baumgartner, J., E. Glatzer, F. Rumler, and A. Stiglbauer (2005). How frequently do consumer prices change in austria? evidence from micro cpi data. European Central Bank Working Paper Series, No.523.

Berck, P., J. Brown, J. M. Perloff, and S. B. Villas-Boas (2008). Sales: Tests of theories on causality and timing. *International Journal of Industrial Organization 26*(6), 1257 – 1273.

Bunn, P. and C. Ellis (2009). Price-setting behaviour in the united kingdom: a microdata approach. *Quarterly Bulletin 2009 Q1, Bank of England 49*(1), 28–35.

Cavallo, A. (2009). Scraped data and sticky prices: Frequency, hazards, and synchronization. Working paper in Harvard University.

Cleves, M., W. W. Gould, R. G. Gutierrez, and Y. Marchenko (2008). *An Introduction to Survival Analysis Using Stata* (second edition ed.). Stata Press, Texas.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological) 34*(2), pp. 187–220.

Cox, D. R. (1975). Partial likelihood. *Biometrika 62*(2), 269–276.

DeGraba, P. (2006). The loss leader is a turkey: Targeted discounts from multi-product competitors. *International Journal of Industrial Organization 24*(3), 613 – 628.

Dias, M., D. Dias, and P. D. Neves (2004). Stylised features of price setting behaviour in portugal: 1992-2001. European Central Bank working paper series No. 332.

Dobson, P. and E. Gerstner (2010). For a few cents more: Why supersize unhealthy food? *Marketing Science 29*(4), 770–778.

Ellickson, P. and S. Misra (2008). Supermarket pricing strategies. *Marketing Science 27*(5), 811–828.

Fisher, T. and J. Konieczny (2000). Synchronization of price changes by multiproduct firms: evidence from canadian newspaper prices. *Economics Letters 68*(3), 271–277.

Fougere, D., H. Le Bihan, and P. Sevestre (2007). Heterogeneity in consumer price stickiness. *Journal of Business and Economic Statistics 25*(3), 247–264.

Guimaraes, B. and K. D. Sheedy (2011). Sales and monetary policy. *American Economic Review 101*(2), 844–876.

Hosken, D. and D. Reiffen (2004). Patterns of retail price variation. *The RAND Journal of Economics 35*(1), pp. 128–146.

Hosken, D. and D. Reiffen (2007). Pricing behaviour of multi-product retailers. *The B.E. Journal of Theoretical Economics 7*(1), 1–43.

Ikeda, D. and S. Nishioka (2007). Price setting behavior and hazard functions: Evidence from japanese cpi micro data. Bank of Japan Working Paper Series No.07-E-19.

Jenkins, S. P. (2005). Survival analysis. University of Essex, Institute for Social and Economic Research, http://www.iser.essex.ac.uk/files/teaching/stephenj/ec968/pdfs/ec968lnotesv6.pdf [Accessed: March 23, 2010].

Kalbfleisch, J. D. and R. L. Prentice (2002). *The Statistical Analysis of Failure Time Data* (2nd ed ed.). John Wiley and Sons, Inc., Hoboken, New Jersey.

Kiefer, N. M. (1988). Economic duration data and hazard functions. *Journal of Economic Literature 26*(2), pp. 646–679.

Klenow, P. J. and O. Kryvtsov (2008). State-dependent or time-dependent pricing: Does it matter for recent u.s. inflation? *The Quarterly Journal of Economics 123*(3), 863–904.

Klenow, P. J. and B. A. Malin (2010). Microeconomic evidence on price-setting klenow. NBER Working Paper No. 15826.

Lach, S. and D. Tsiddon (1996). Staggering and synchronization in price-setting: Evidence from multiproduct firms. *The American Economic Review*.

Lal, R. (1990). Price promotions: Limiting competitive encroachment. *Marketing Science 9*(3), 247–262.

Lloyd, T., C. W. Morgan, S. McCorriston, and E. Zgovu (2011). Do sales matter? evidence from uk food retailing. The University of Nottingham Discussion Paper in Economics, No. 11/01.

Loy, J. and C. Weiss (2002). Staggering and synchronisation of prices in a low-inflation environment: Evidence from german food stores. *Agribusiness 18*(4), 437–457.

Meyer, B. D. (1990). Unemployment insurance and unemployment spells. *Econometrica 58*(4), 757–782.

Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics 123*(4), 1415–1464.

Nakamura, E. and J. Steinsson (2012). Price rigidity: Microeconomic evidence and macroeconomic implications. *Working paper in Columbia University*.

Nakamura, E. and D. Zerom (2010). Accounting for incomplete pass-through. *Review of of Economic Studies 77*(3), 1192–1230.

Pashigian, B. P. and B. Bowen (1991). Why are products sold on sale? explanations of pricing regularities. *The Quarterly Journal of Economics 106*(4), pp. 1015–1038.

Pesendorfer, M. (2002). Retail sales: A study of pricing behavior in supermarkets. *The Journal of Business 75*(1), pp. 33–66.

Richards, T. J. (2006). Sales by multi-product retailers. *Managerial and Decision Economics 27*, 261–277.

Sobel, J. (1984). The timing of sales. *The Review of Economic Studies 51*(3), 353–368.

Taylor, J. (1979). Staggered wage setting in a macro model. *The American Economic Review*, 108–113.

Varian, H. R. (1980). A model of sales. *The American Economic Review 70*(4), pp. 651–659.

Veronese, G., S. Fabiani, A. Gattulli, and R. Sabbatini (2004). Consumer price behaviour in italy: Evidence from micro cpi data. European Central Bank Working Paper Series, No.449.

Appendix A: Censored spells in scanner prices

There are examples of complete, left-censored and right-censored regular price spells. They are created from the retail price series in the scanner dataset. In the figure, the top price series is the Gerber Libbys Organic Orange juice (Tetra 1L 4Pack) in Waitrose and the bottom one is the Gerber Libbys Organic Orange Juice (Tetra 1L Single) in Tesco. A left-censored, complete and right-censored spells created from the top price series respectively; a double-censored spell is created from the bottom price series.

Figure 12: Examples of types of censored spells

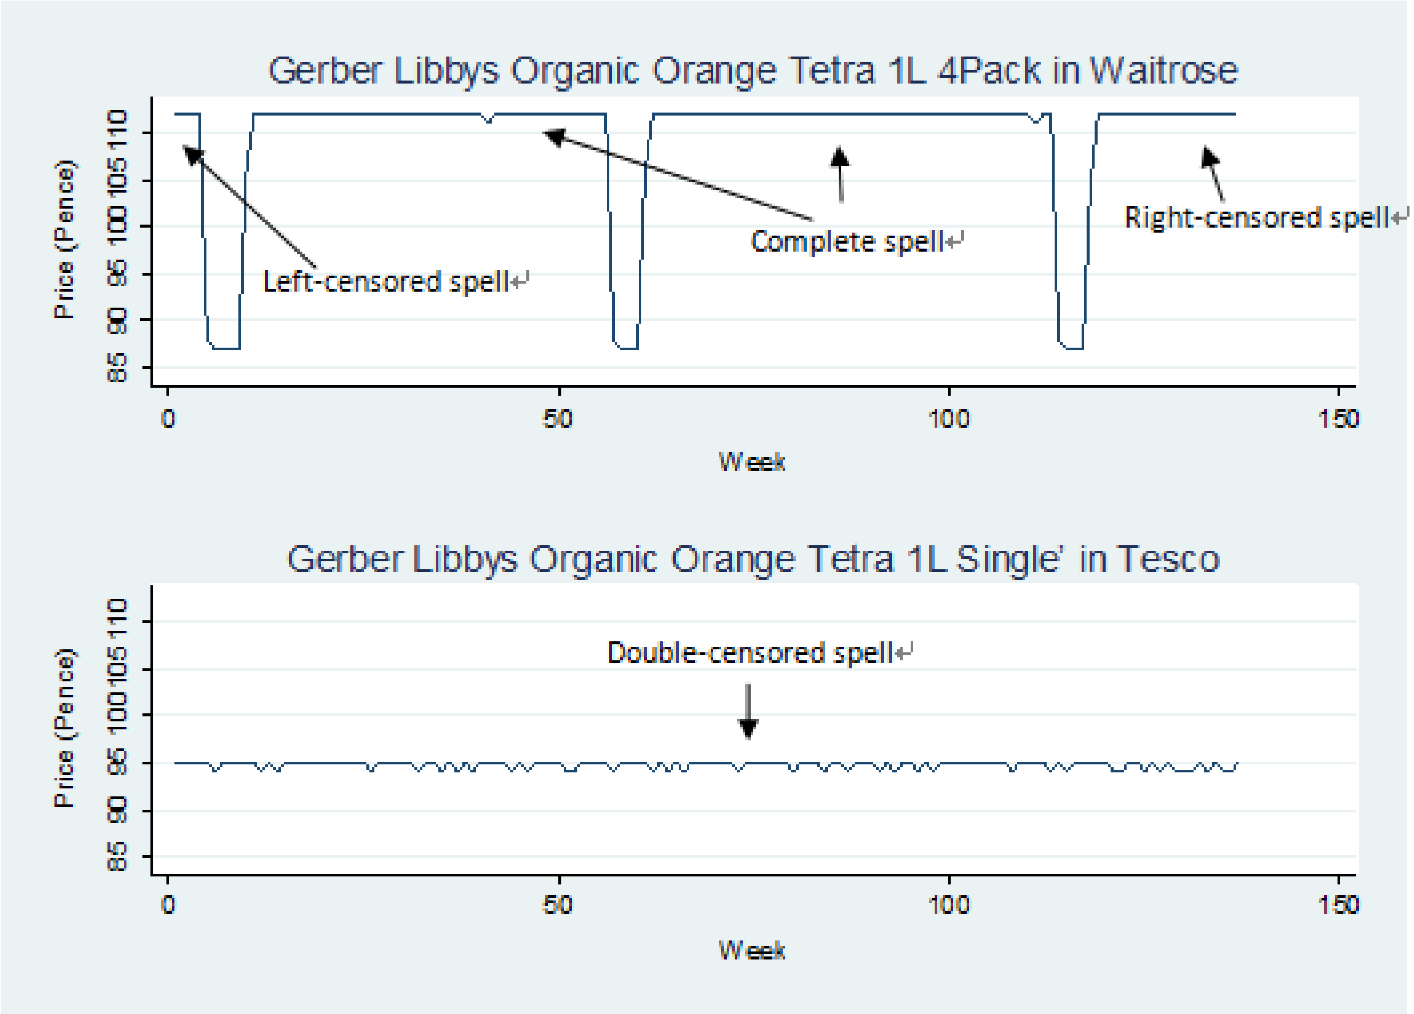

Note: the top price series illustrates the left, right-censored and complete spells; the bottom one demonstrate a double-censored spell.