Price promotions and supermarket pricing: A duration analysis of UK supermarket prices

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

University of Nottingham, UK

Draft August 23 2012

*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.

1 Introduction

Price promotions or 'sales' are an important element of pricing, and none more so than in food retailing. Data from the UK 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 theories 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, 2000 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, 2008; Guimaraes and Sheedy, 2008).

We also seek to investigate whether the relationship between the duration of regular price

spells are invariant to (state-dependent) factors such as retail chain, product perishability and brand status (national brands and private label products). 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.

To evaluate these aspects of sales behaviour we conduct an empirical analysis based on extensive sample of food prices acquired from the AC Nielsen *Scantrak* 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 over 1700 unique product codes (UPC)[1] and 231,000 price observations in all. From this dataset we calculate the durations of each regular price spell that is applied to a statistical technique commonly known as duration analysis. In doing so we recognise that price promotions are not evenly distributed across products: some products are promoted frequently while others rarely. It transpires that the key result, concerning the likelihood of sale, turns on whether this issue (which we call multiple sales) is addressed. We propose a new approach to accommodate multiple sales.

In the results, we find 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. When they are not, the relationship is negative, in stark contrast to the positive relationship that obtains when multiple sales are explicitly recognised. 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 (2000). Accounting for multiple sales helps square the circle between price setting and modern theories of sales behaviour. Furthermore, we find that the positive time-dependent pattern varies across heterogeneity such as product format, brands and especially supermarkets. 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).

The paper is organized as follows. Section 2 reviews the theories of sales both theoretically

---

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

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

The early literature of the theory of sales is dominated by the static models of Shilony (1977) and most notably 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) and Pesendorfer (2000) develop 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 m 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.

Hosken and Reiffen (2007) augment the time-dependent framework with a state-dependent dimension of product perishability. For products that can be stored, sales are characterised as being time-dependent while they are random in the perishable products. 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.

A number of paper show empirical analysis of sales related to supermarket pricing using micro-level data. In the US, Pesendorfer (2000) 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

4

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 in the sales pattern and supermarket pricing in the disaggregate level. 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 and brands in the estimation of sales patterns.

In empirical macroeconomics, where a key issue is that of sticky prices a number of studies have estimated the hazard function of price changes. While an important contributor to price changes the occurrence of sales is not the primary focus of these but they are relevant because they typically find that prices are less likely to change the longer they persist. Examples include Nakamura and Steinsson (2008), 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). Fougère et al (2007) analyses price duration in the French retailing market and find that patterns of price changes are either downward sloping or flat. In the macroeconomics literature of price duration two papers are particularly relevant owing to the fact that unlike the others, they propose methods that acknowledge uneven distribution of price changes, in that some product groups experience more frequent price changes than others, an issue that is important to the duration of sales. Ikeda and Nishioka (2007) analysis relates to Japan and Nakamura and Steinsson (2008) is for the US. Their results of price patterns are mixed.

3   The Data Description

3.1   The scanner dataset

The current research uses a high-frequency disaggregate-level scanner dataset purchased from Nielsen *Scantrak* 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 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,

---

[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).

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

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.

Figure 1: Price observations from a typical UPC
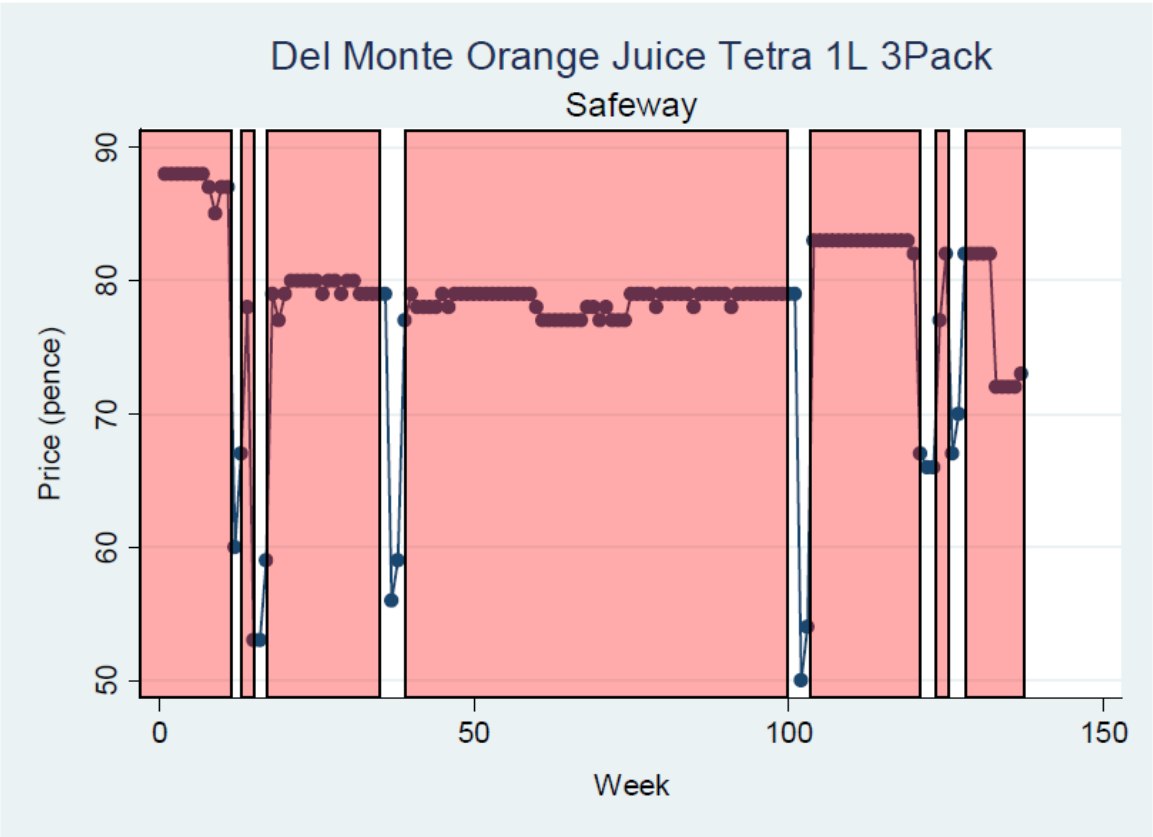
3.2 Censoring

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

7

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 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.

---

[5]Around 90% of the UPCs are available for the full 137 week window, the remainder being 120 weeks long.
[6]One example of censored spells is shown in Appendix A.
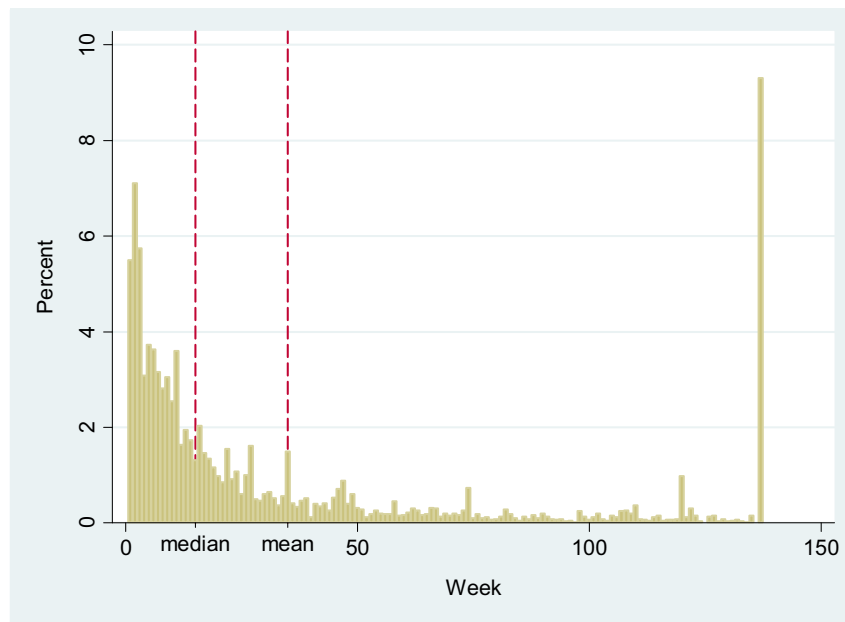
Figure 2: Histogram of regular price spells

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; Fougère, 2008; 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

---

[7]As a robustness check, we have also estimated models using all the regular price spells (i.e. ignoring the

average durations of the full and censored samples. Notice that measures of central tendency are around 40% of those in the full sample, owing primarily to the removal of the double censored spells of long duration. Clearly however the leptokurtic of the duration distribution is largely invariant to the censoring given the wide disparity in mean and median duration of the censored sample which are 20 weeks and 10 weeks respectively.

Table 2: Measures of regular price spell duration

	All data		Censored Data	
	Mean Duration (weeks)	Median Duration (weeks)	Mean Duration (weeks)	Median Duration (weeks)
TOTAL	**35**	**15**	**20**	**10**
TESCO	45	22	24	15
SAINSBURY	39	20	29	17
ASDA	92	110	54	41
SAFEWAY	23	8	14	5
SOMERFIELD	28	13	17	9
KWIK SAVE	28	11	17	10
WAITROSE	42	23	26	16

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

3.3 Heterogeneity in supermarkets, 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 mainstream retailers) resembling ASDA whereas Somerfield and Kwik Save (both soft discounters) more like Safeway. In contrast to the distributions by retailer, there is

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.

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 spells

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. 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.

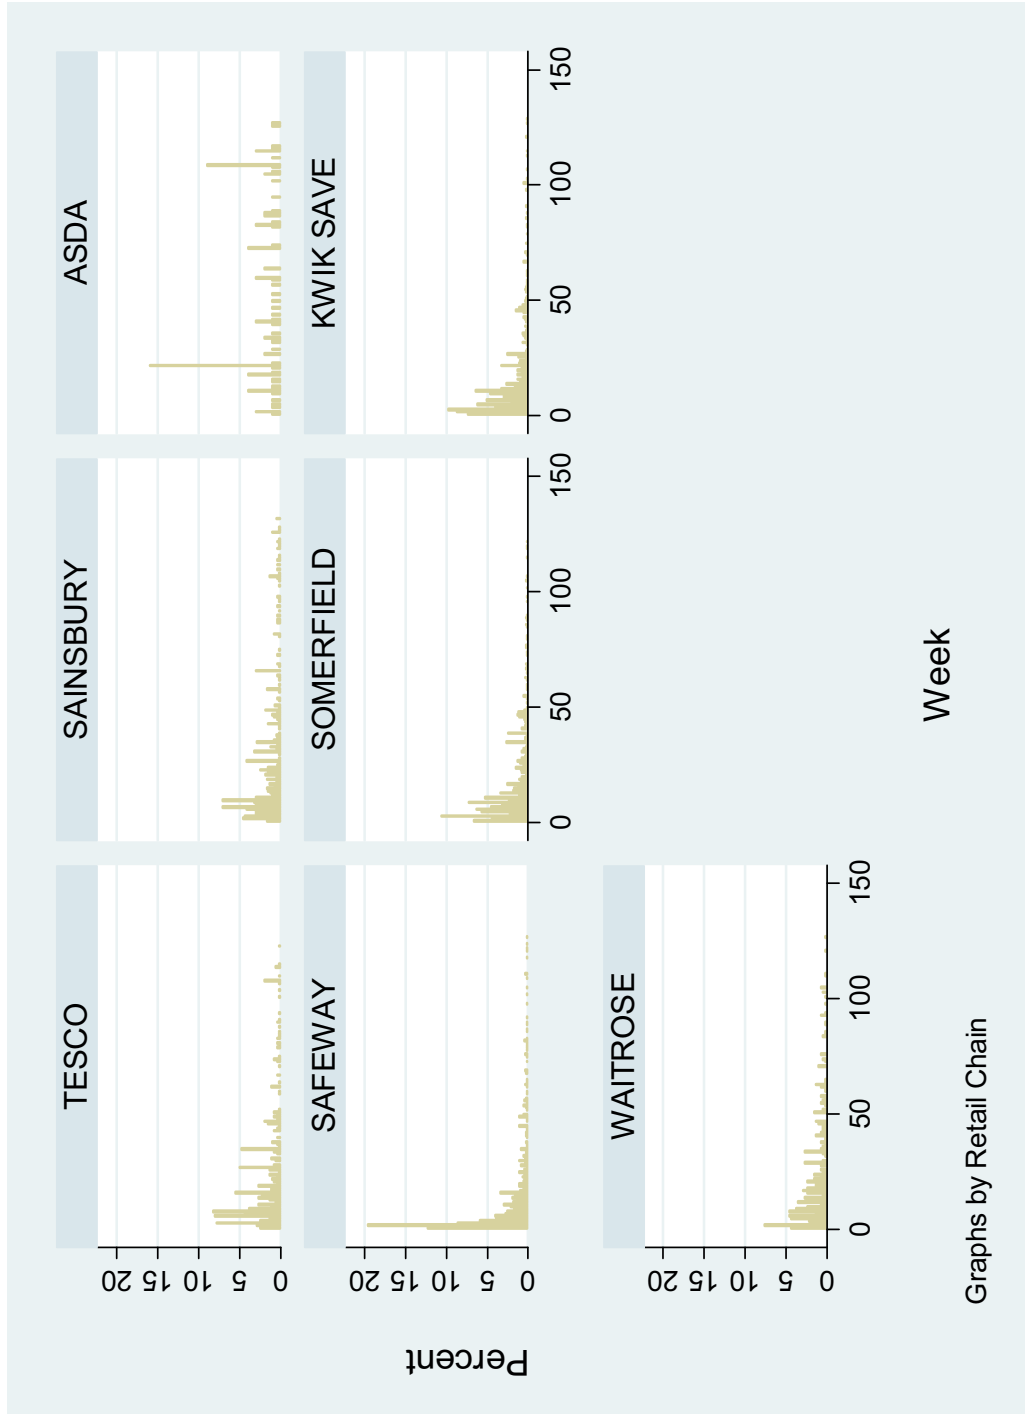Figure 3: Histogram of regular price spell duration by supermarket
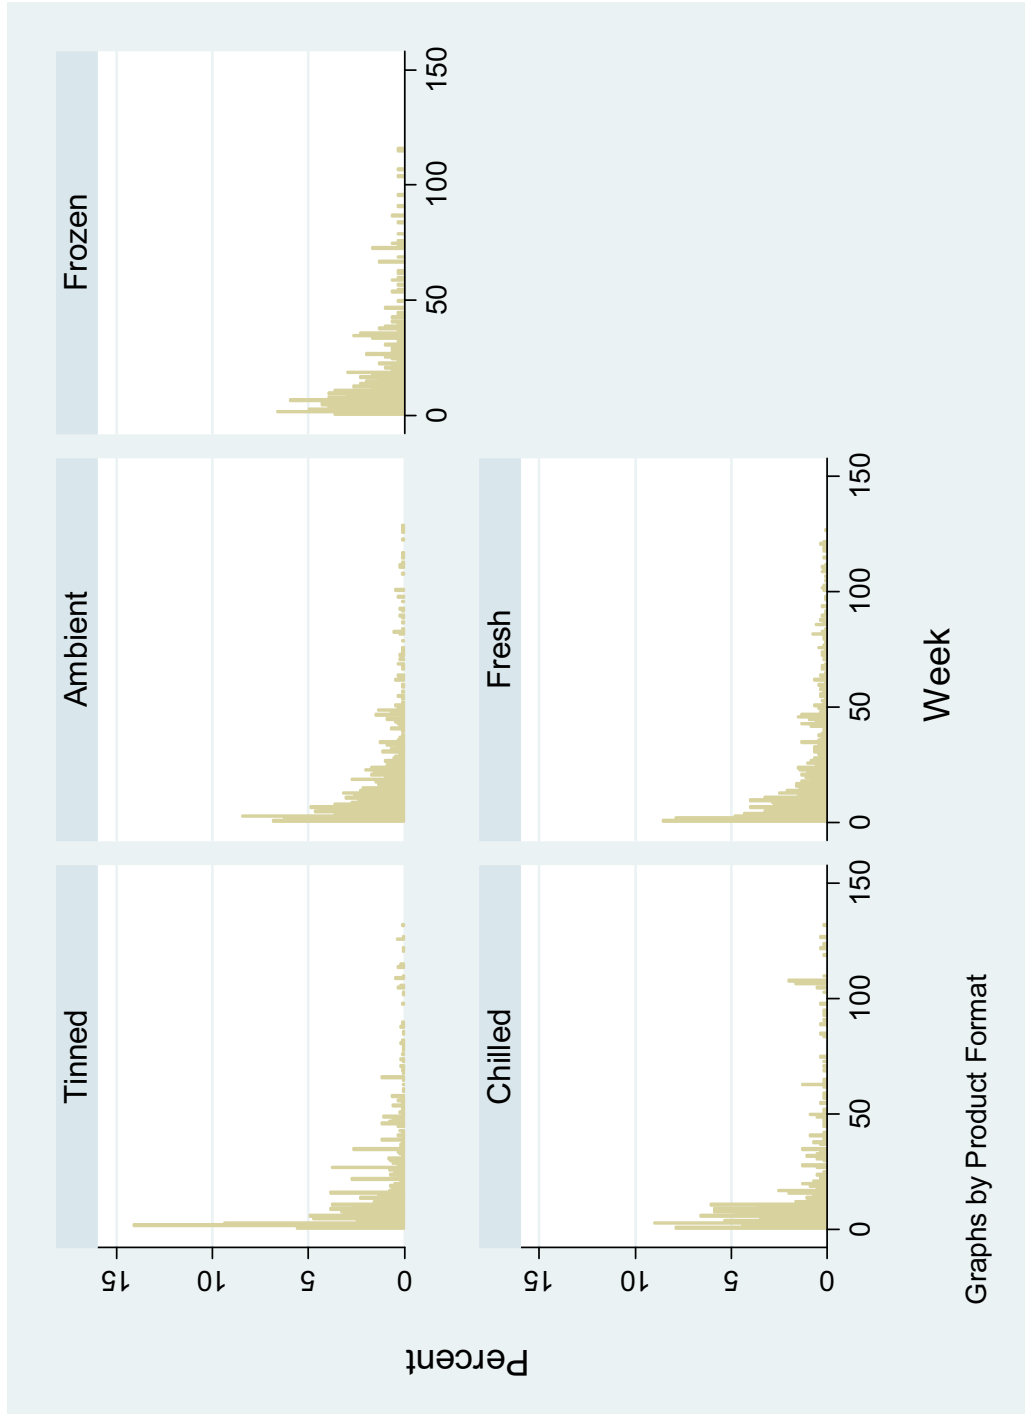
Figure 4: Histogram of regular price spells across format

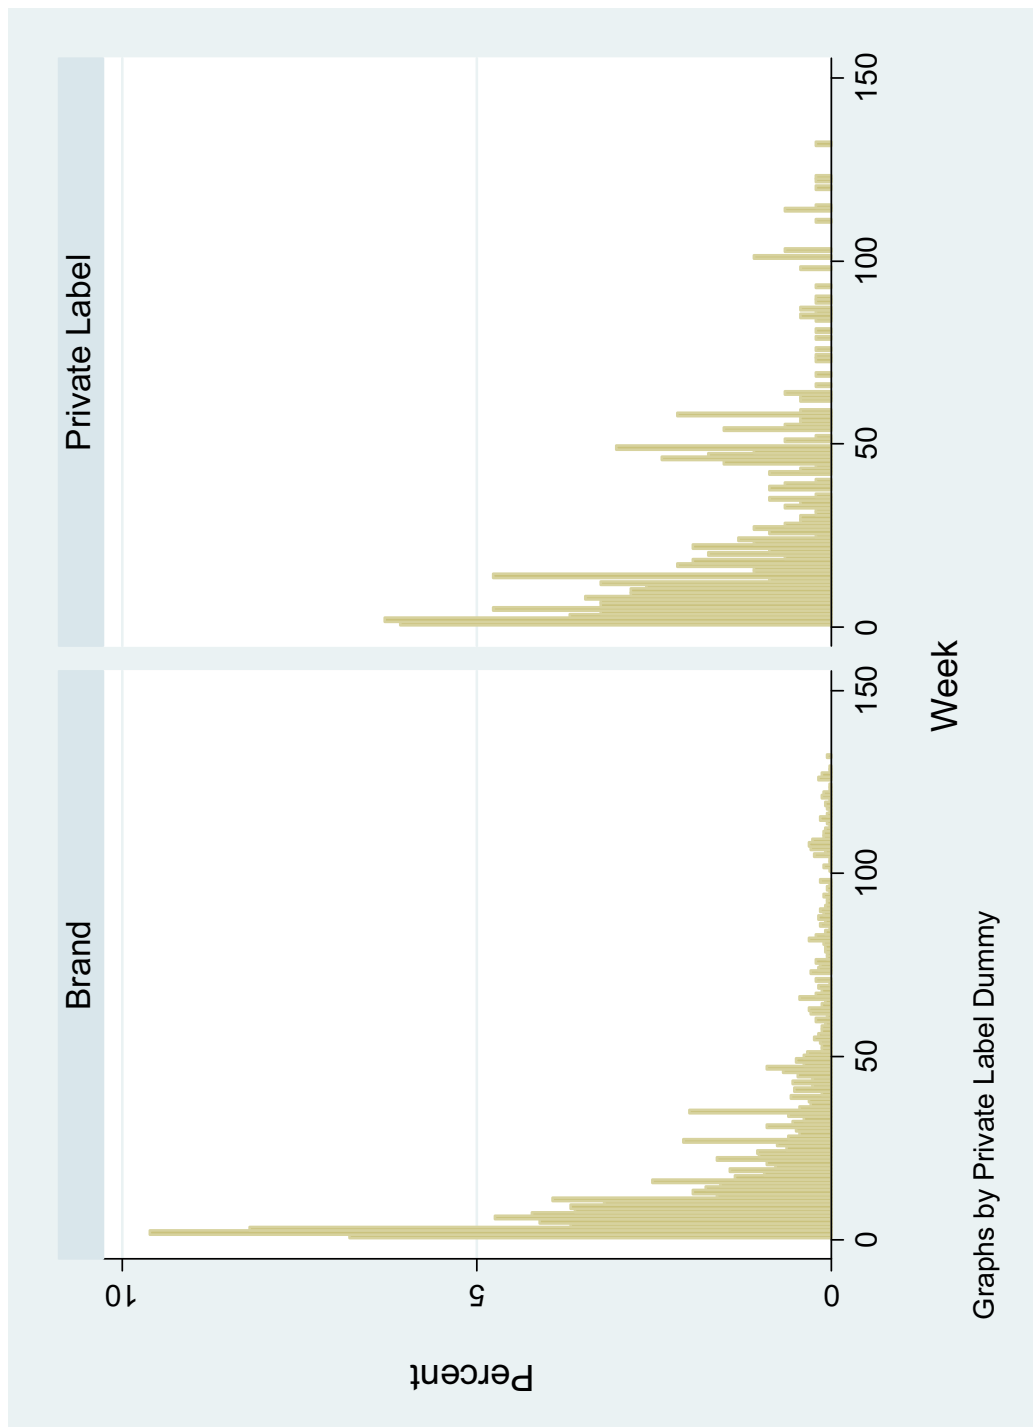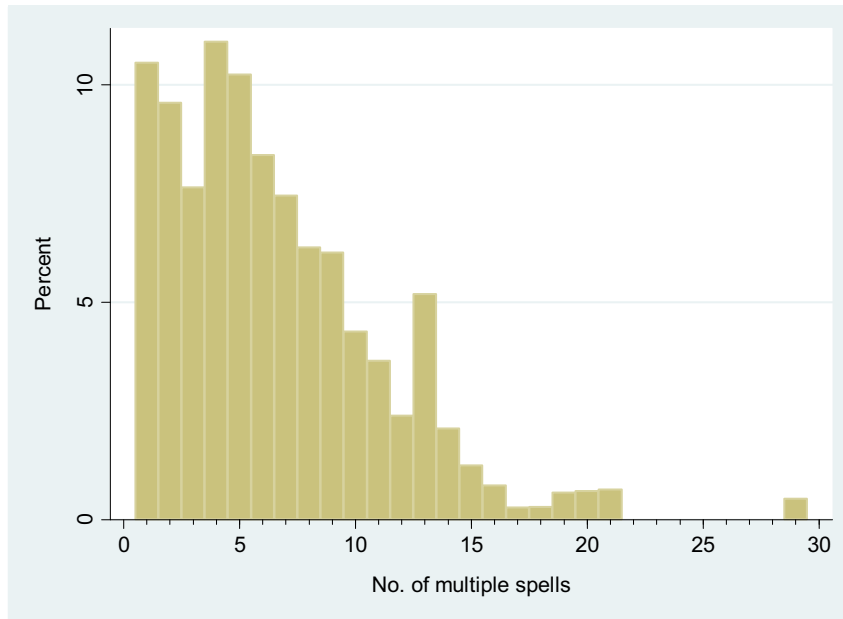Figure 5: Histogram of regular price spells across brands

Figure 6: Distribution of regular price spells

3.5 Seasonality

Finally, one last facet of the data worthy of note is seasonality. Figure 7 shows the month of the year that regular price spells are terminated in. Given that by definition sales demarcate spells of regular prices the figure also sheds some light on the seasonality of promotional behaviour. We find that more than 10% of regular price spells are terminated in November and March, twice as many as in the summer months of May, June and July. While these winter and spring peaks could reflect the Christmas and Easter festivities at these times, it is interesting that the peaks occur in the month preceding each festival, rather than the festival month itself, a feature also noted by Nakamura and Steinsson (2008) in their analysis of price changes in US retailing.

Figure 7: Seasonality of regular prices
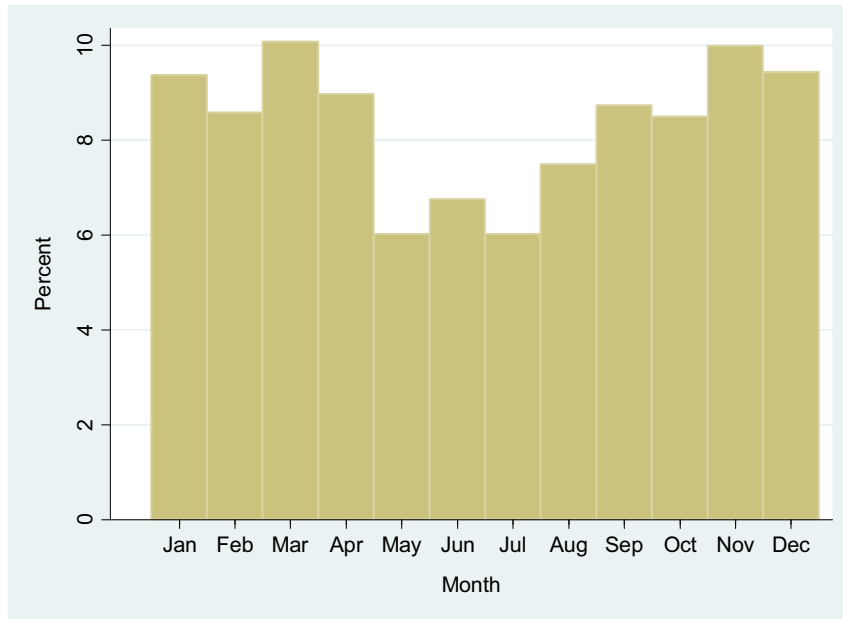
4 The Methodology

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, Fougère et al. 2007, Ikeda & 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 call 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 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.

More formally, let T be a non-negative random variable measuring the duration of a regular price spell in duration time t, with density function $f(t) = \lim_{\Delta \to 0} P(t < T < t + \Delta)/\Delta$. This describes the probability of a sale terminating regular price spells of length T .[9] Accumulating those probabilities over all spell durations gives the cumulative density $F(t) = P(T < t)$, the probability that a sale will have terminated regular price spells lasting up to t periods long. It then follows that the probability that regular prices persist longer that t, what is called the

---

[8]For example, where the risk of a sale is constant over time, as in the Varian model, time to a sale is exponentially distributed.

[9]Duration time differs from actual time in that it excludes periods of sales and left censored spells.

survival 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 of the regular price spell $h(t)$ 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. The continuous hazard function is written as:

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

It is important to recognise that unlike the density function describing the duration of regular prices $f(t)$, the hazard function define conditional on the regular price spell having lasted for t periods. If the spell is right-censored, the hazard function is given by the survival rate of regular price spells so is written as $h(t) = P(T > t) = S(t)$. Finally, we define the *cumulative hazard function* as

$$H(t) = \int_0^t h(t)dt = -\ln S(t) \tag{2}$$

which accumulates the conditional probability of the occurrence of sales over time and clearly shows the inverse relationship between the accumulated risk of experiencing a sale and the probability of the regular price surviving in to the future.

Various methods to calculate the hazard function have been developed in the literature, the simplest of which is the non-parametric approach. This simply involves plotting the unconditional hazard rate for each duration time. However, semi-parametric analysis allows the hazard to vary according to covariates (conditional hazard function). In the non-parametric analysis, the cumulative hazard rate for each duration time is estimated using the discrete cumulative hazard function:

$$H(t_j) = \sum_{j|t_j > t} \frac{k_j}{n_j} \tag{3}$$

where k_j is the number of complete regular price spells at week t_j; n_j is the total number of regular price spells at week t_j including both complete and right-censored spells. As a result, the ratio k_j/n_j is an estimate for the cumulative conditional probability that a sale occurs at week t_j. Then, the hazard rates are calculated using $\Delta H(t_j) = H(t_j) - H(t_{j-1})$.

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

$$h_i(t) = h_0(t) \exp(\beta_0 + \mathbf{x}_i \beta) \tag{4}$$

where $h_0(t)$ is a baseline hazard common to all i as the duration of regular price spell increases; $\mathbf{x}_i$ is a vector of covariates used to control for the heterogeneity in the duration data and β is a vector of coefficients to be estimated with the sample of observations. In this specification the exp() function ensures that the hazard is non-negative and the covariates have a multiplicative or proportional effect, shifting the baseline hazard up or down according to characteristics contained in $\mathbf{x}_i$. The attraction of the proportional hazards model lies in its flexibility (many functional forms can be accommodated, some of which can be directly translated in to time-metric representation), tractability (it is easily estimated by maximum likelihood in modern software) and because it allows state-dependent influences ($\mathbf{x}_i$) to be augmented by a time-dependent baseline hazard, which proxies for a potentially large number of omitted variables that may be correlated with the passing of time (time itself not having any 'causal' effect). (4) is commonly estimated using semi-parametric methods developed by Cox (1972, 1975) whereby the parameters on the covariates of the model are estimated by 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

19

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, for positive coefficients $\exp(\hat{\beta}) > 1$ the covariate increases the hazard by $\hat{\beta}\%$ *ceteris paribus*; when $\hat{\beta}$<0, $\exp(\hat{\beta}) < 1$ and the hazard is reduced by $(1 - \hat{\beta})\%$ *ceteris paribus*.

4.2 Accounting for Multiple spells

Unlike many (time-to-death) biomedical applications of duration analysis, the time-to-sale for many UPCs is not a once-and-for-all occurrence. As noted in Section 3 promotions are widely used and in many products, repeatedly so. This gives rise to the existence of multiple spells and taking proper account of this aspect of our duration data is key to the calculation of an appropriate hazard rate since the hazard of a spell of length T is likely to differ given that some products are more frequently promoted than others. In the literature, one response is to pick a single spell at random from the multiple spells for each UPC (Fougère *et al*, 2007). Given the vast dataset (2.3 million price observations) at their disposal and a focus on price changes (i.e. price stickiness) rather than the occurrence of sales, this approach fits the bill. Ikeda and Nishioka (2007) use a finite mixture model to control for multiple spells in which spells are clustered by product groups and then the product-specific weights are attached to each UPC in the likelihood function. While the method serves their purpose, it is computationally complex. In this paper we propose an altogether more tractable approach that directly accounts for a UPC's sales profile. We set out the approach below.

In the duration analysis literature, 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 delayed entry (complete) spell is[10]

---

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

$$h^d(t_0 + t) = \frac{P(T = t_0 + t)}{P(T > t_0 + t \mid T > t_0)} = \frac{f(t_0 + t)}{\frac{S(t_0 + t)}{S(t_0)}} = \frac{h(t_0 + t)}{S(t_0)} \tag{5}$$

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. (5) 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 (5) says nothing about the order of the spell in the sequence or the number spells the UPC experiences, it does recognise that the spell occurred with delayed entry, which is always the case for any multiple spell. 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. Treating spells as independent thus introduces a bias in the estimate of 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 15, 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 15 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 for the UPC $t_0 = 0$ 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 8: Effects of multiple spells: an example
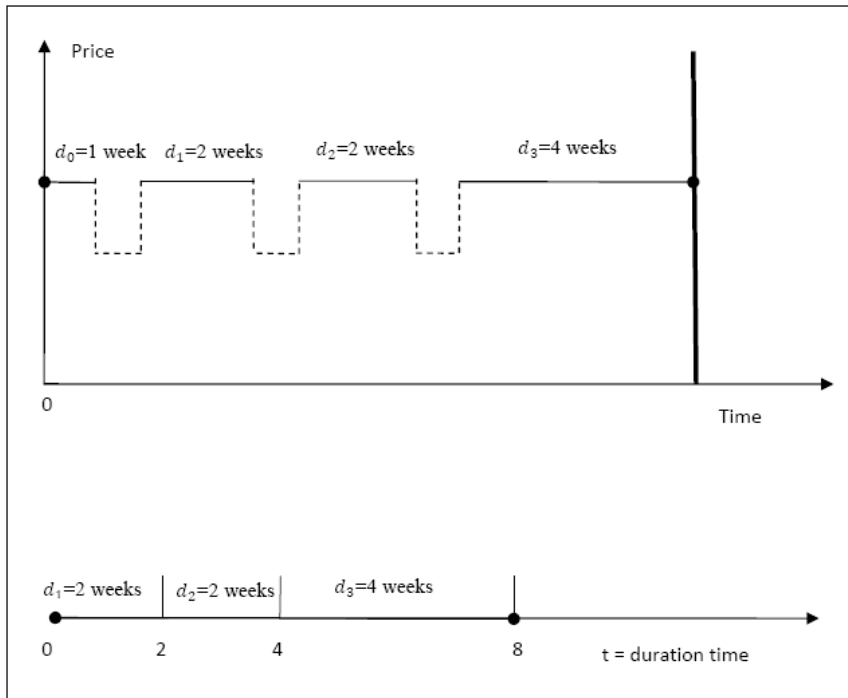
According to (5), the hazard rate of spell d_1 is

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

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 duarion. In contrast, d_2 and d_3 are delayed entry spells with associated hazards given by

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

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 simple the ratio of survival functions. In contrast, if we ignore the effect of delayed entry so that each of the spells is treated as if they were single spells, then the hazard of d_1, d_2 and d_3 is :

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

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

$$h(4) = P(T > 4) = S(4)$$

Table 3 summarises these results.

Table 3: Effects of multiple spells: Delayed entry

No.	Censoring	No delayed entry			Delayed entry		
		t_0	t	Hazard	t_0	t	Hazard
d_0	left-censored	-	-	-	-	-	-
d_1	complete	0	2	$\frac{f(2)}{S(2)}$	0	2	$\frac{f(2)}{S(2)}$
d_2	complete	0	2	$\frac{f(2)}{S(2)}$	2	2	$\frac{f(4)}{S(4)/S(2)}$
d_3	right-censored	0	4	$S(4)$	4	4	$\frac{S(8)}{S(4)}$

As noted previously, the delayed entry effect is only one albeit important aspect of multiple spells, and ideally we wish to capture all the information contained in a UPC's regular price profile, such as the number and order of regular price episodes, both of which may affect the hazard. For example, in the same way that sales may be more likely in those UPCs that have experienced sales in the past, the hazard may vary with time or the number of previous sales. Nakamura and Steinsson (2008) attempt to estimate these dimensions of the UPC-specific profile using a random effects term in the hazard function. They account for the multiple spell effects as unobserved heterogeneity in the hazard function.

An alternative is to include the UPC-specific aspects of each regular price episode as fixed effects in the hazard function. Fixed multiple effects terms are more useful to describe the behaviour of multiple spells as well as their effects on the hazard function than random effects model. Specifically, we include two covariates, namely: the number of spells experienced by that UPC; and the order of the spell in the sequence; both of which are summarised in Table 4. In general, we expect the number and order of spells to increase the hazard.

Table 4: Description of covariates

Covariate	Details of variables	Within a UPC
Number of spells	Total number of spells per UPC	Time-invariant
Order of spells	Number of spells per UPC at week t	Time-varying

To motivate the use of these covariates recall the UPC illustrated in Figure 15(a), which we reproduce in Figure 9 as UPC1 along with another denoted named UPC2. Spells of UPC1 are denoted as d_{1i} and spells of UPC2 are d_{2i}. Without taking account of UPC specific differences the hazard of the spells d_{11} and d_{21} are the same, despite UPC2 being promoted less frequently. To distinguish the two we therefore include the number of regular price spells experienced by each UPC.

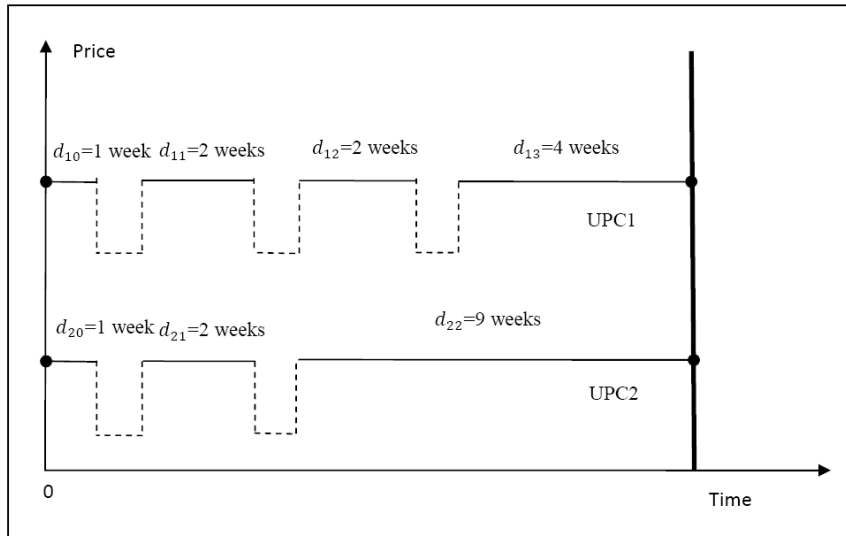Figure 9: Fixed effects covariates: Example of total number of spells

Table 5 details the values of the two covariates for these two hypothetical UPC price profiles. While the number of spells per UPC is feature that remains constant for each of the UPC, regardless of duration time, the order of spells provides a time-varying dimension in the hazard. In all they produce a relatively rich set of conditioning variables in the hazard function despite seemingly similar price profiles.

Table 5: Fixed effects covariates in the hazard function

Duration time	Number of spells	Order of spells
UPC1		
2	3	1
4	3	2
8	3	3
UPC2		
2	2	1
9	2	2

In sum, this discussion underlines the importance of dependence when multiple spell data is used to estimate the hazard function. Given that most spells of regular prices are from products

that have experienced a number of sales during the sample period they are both numerous in the data and because sales behaviour tends to be UPC specific, in that some UPCs are promoted frequently while others are not, treating spells of regular prices as independent events is likely to be inappropriate, producing estimates of the hazard that suffer from potentially serious bias.

5   Empirical Results and Discussion

In the analysis, both non-parametric and semi-parametric estimations are implemented to analyse the sales behaviour over time and across supermarkets. In non-parametric analysis, we estimate the unconditional hazard functions of sales for overall sales behaviour by graphing all the regular price spells together and then examine the state-dependent heterogeneity in formats, supermarkets and private-label products further by graphing spells across different characteristic groups (conditional hazard functions). In the semi-parametric analysis, we include the covariates in the hazard function to obtain a more detailed picture of sales behaviour observed in UK food retailing. 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   Non-parametric results

Using non-parametric methods, we obtain unconditional hazard function of sales for the entire dataset given in Figure 10. Time dependence is apparent suggesting that sales are not random events according to this measure: the probability of a sale initially rising as the duration of regular prices increases, peaking at 17 weeks and falling thereafter. Hence for products that remain without sale for more than four months, the less likely the product will be promoted.

Figure 10: Unconditional hazard function of sales (ignoring delayed entry)
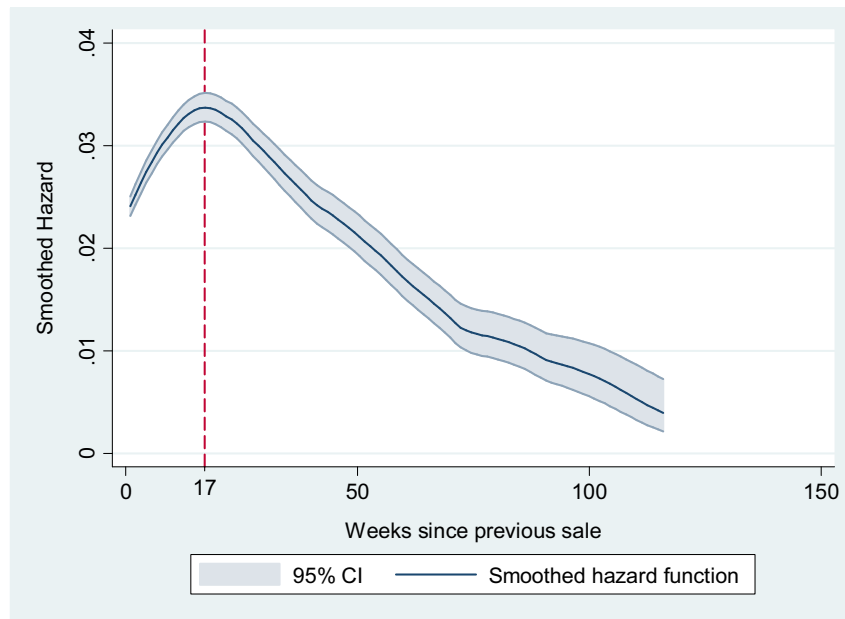
Of course, the unconditional hazard function of sales portrayed in Figure 10 ignores any heterogeneity (by supermarket or brand status for example) and treats all regular price spells as independent events unrelated to the sales profile of the UPC in question. To relax the last of these two assumptions, we cluster the 4,303 spells in 1,703 UPCs to account for delayed entry spells and this gives rise to the hazard function illustrated in Figure 11, which now rises as the duration time increases until roughly one year (50 weeks) has elapsed since the last sale; then it falls, only to peak again at two years since the previous sale, possibly suggesting that there are annual cycles in the sales behaviour; the second being stronger than the first. 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. Hence time dependence is also apparent here, and in a somewhat more complicated form. This pattern may arise due to heterogeneity between sub-populations of the data, a possibility that we explore in Figures 12, 13 and 14 which portray the non-parametric hazard functions by retailer, format and brand status respectively.

Figure 11: Unconditional hazard function of sales (with delayed entry)
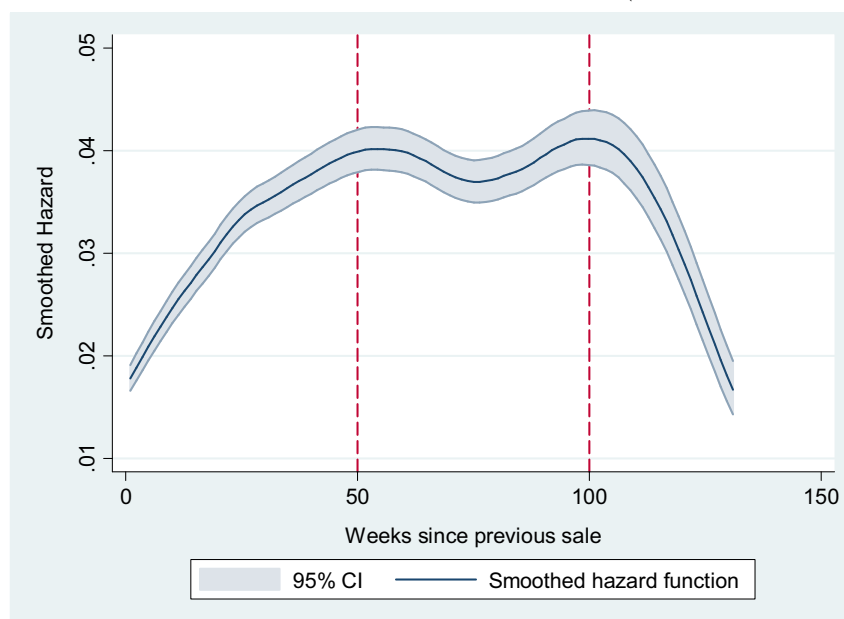
Comparing the figures with the hazard for the data as a whole, the importance of heterogeneity becomes clear in that the complicated shape of the aggregate hazard is an amalgam of category-specific characteristics rather than a pattern that is common to all categoriastions of the data. For example, as, the hazard functions by retailer show, most hazards are unimodal, albeit at either 50 or 100 weeks, so that the bimodal pattern observed in uses all the data is most likely a combined effect of those individual patterns of retailers, rather than a cycle common to all retailers. Also one retailer, ASDA, is something of an outlier, showing little time dependence in the occurrence of sales at all. Heterogeneity by format and brand is also observed: whereas tinned products display a twin-peak other formats do not. Finally, brands are more likely to be promoted than private label products.[12]

---

[12] We also estimate the price spells generated using different sales thresholds (25% and 35%) to test the sensitivity of the results. Although the scale of the hazard patterns are significantly reduced when the 25% and 35% sales definitions are used, the shape of the hazards are unaffected by the sales thresholds. Therefore, we only report the parametric results from the 10% sale in the following for simplicity. Full details available upon request.

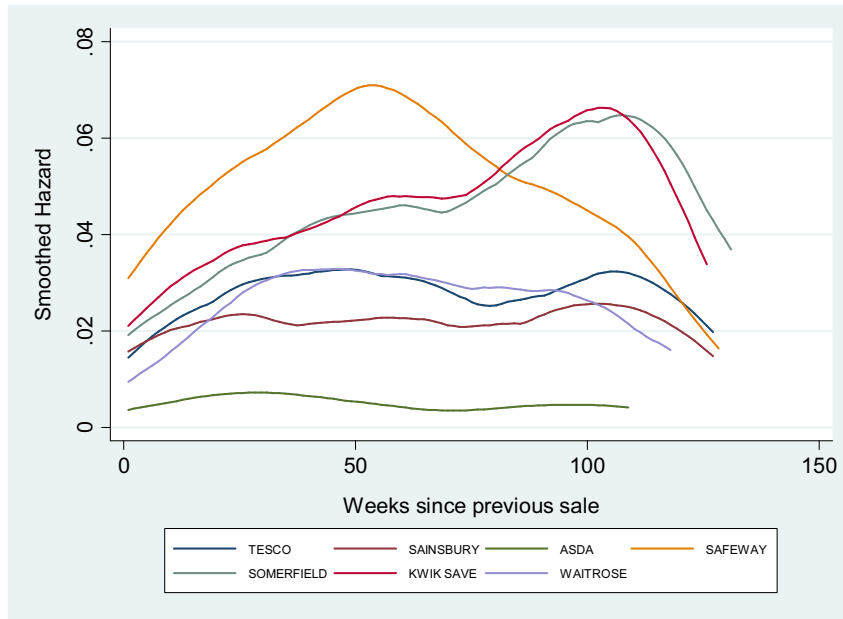Figure 12: Unconditional hazard function of sales by retailer

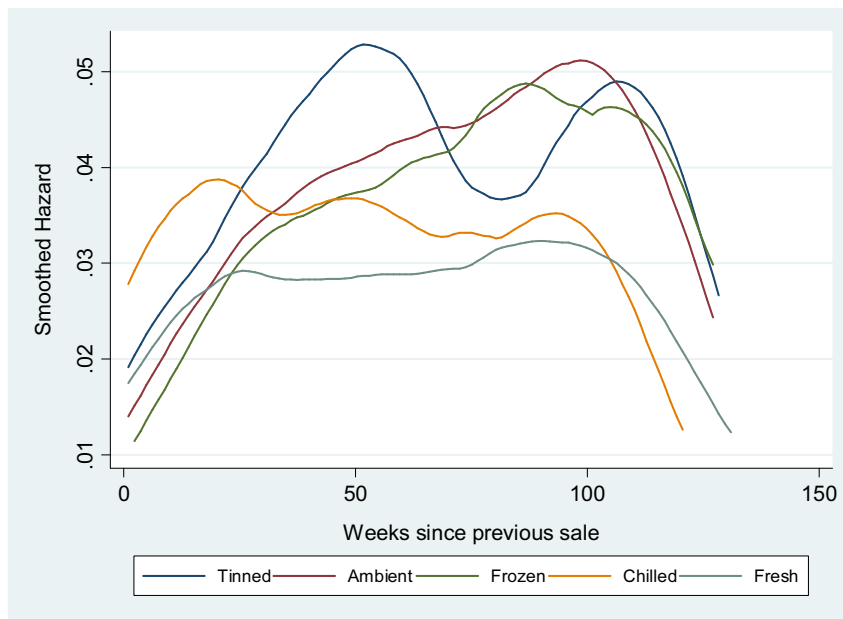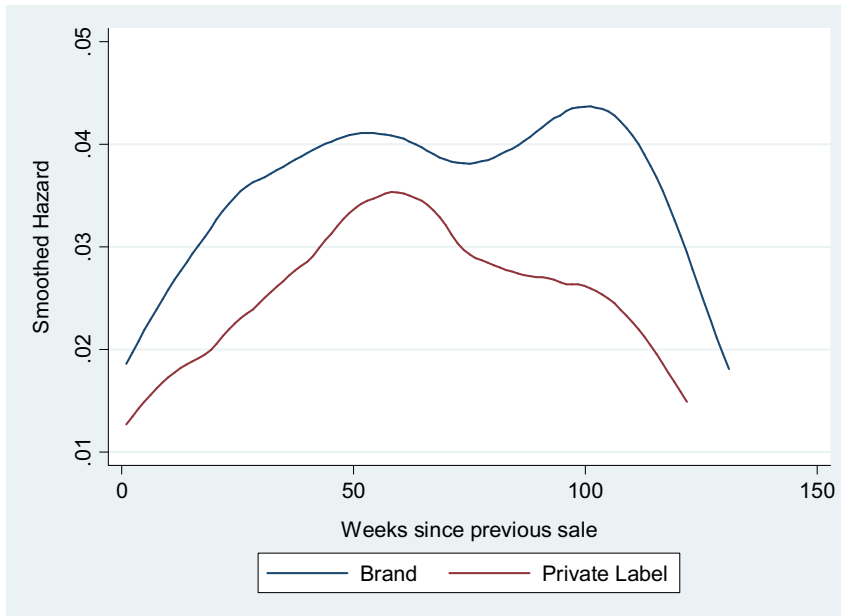Figure 13: Unconditional hazard function of sales by format

Figure 14: Unconditional hazard function of sales by brand status

While hardly necessary given the clear differences visible in the figures we compute log-rank tests to test the heterogeneity more formally. Table 6 shows that, in all cases, the null hypothesis that survival or hazard functions across characteristics are not statistically significant is rejected at 1% significant level using χ^2 statistics, confirming that the differences in hazard functions are statistically significant across retailers, product durability and brand status.

Table 6: Log-rank test of heterogeneity

Characteristic	Log-rank test
Retailer (supermarket dummies)	$\chi^2(6) = 646.90^{***}$
Product durability (format dummies)	$\chi^2(6)= 169.62^{***}$
Private-label Product (brand dummy)	$\chi^2(1) =79.21^{***}$

Note: *** signifies significance at the 1% level

5.2 Semi-parametric results: overall sales pattern

The results of the preceding section have an important bearing on the semi-parametric mod-

elling. They suggest: first, treating regular price spells as independent (i.e. single spells) may distort the hazard function and second, while account needs to taken of the heterogeneity that characterises the data, there is sufficient similarity across the various classifications of the data to suggest that the proportional hazard specification may not be unreasonable specification to adopt. To begin we assess the role of multiple spells uses two regression models, both of which condition on dummies for supermarket, format and brand status, augmented by monthly seasonals, but they differ in their treatment of multiple spells. The first of these which treats sales as independent events across UPCs produces the negatively sloped baseline hazard presented in Figure 15; the second, which takes account of multiple spells using delayed entry and covariates for UPC sales profile is presented in Figure 16. As with the non-parametric results, taking account of multiple spells reverses the slope of the baseline function, emphasising just how important it is to account for multiple spells. In effect, the change induces a reversal of the hazard, underlining that the assumption of spell independence is far from innocuous regarding inference on the hazard of sales.

Before we consider the estimated coefficients from the estimated models, it is worthwhile to consider why the treatment of multiple spells imparts such a decisive impact on the hazard. The explanation is that, because the multiple spell model controls for key aspects of sale history, namely delayed entry, average spell length and spell order in the sales profile of each UPC, whereas the single spell model does not, the baseline hazard of the latter becomes dominated by spells of longer duration as spells from frequently promoted products drop out of the risk set. Initially, spells from all UPCs are in the duration data so the shape of the baseline hazard from the single spell model mimics that of the multiple spell model: the longer the duration of the regular spell the more likely it is that a sale will occur. However, as the spells from frequently promoted products drop out of the duration data, only longer spell durations remain hence the single spell hazard generally declines, in an analogous manner to the unconditional hazard presented in Figure 10. In contrast, by controlling for some the factors that affect the frequency of sales the baseline hazard from the conditional model reveals the underlying tendency that is

common across all UPCs, and this displays a generally upward sloping hazard, implying that longer the product remains without a sale, the more likely it is to be promoted.

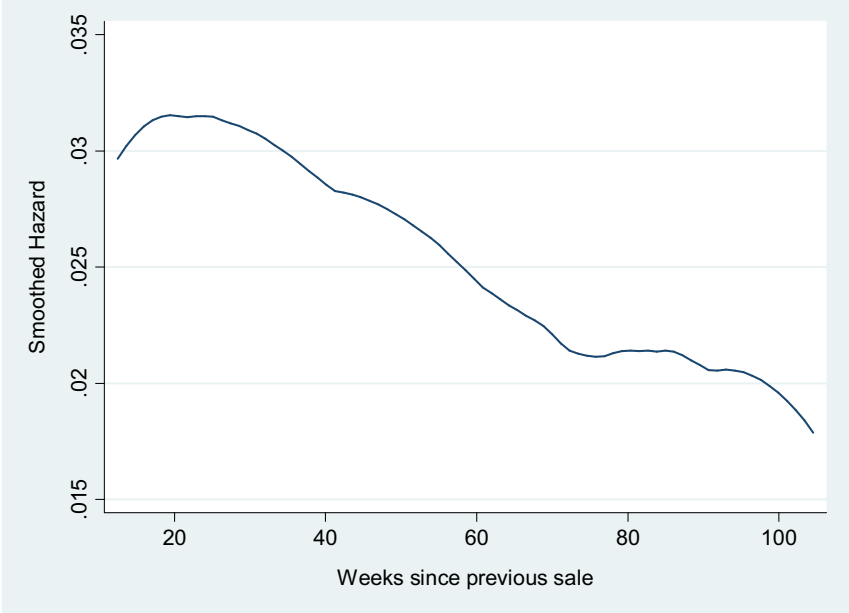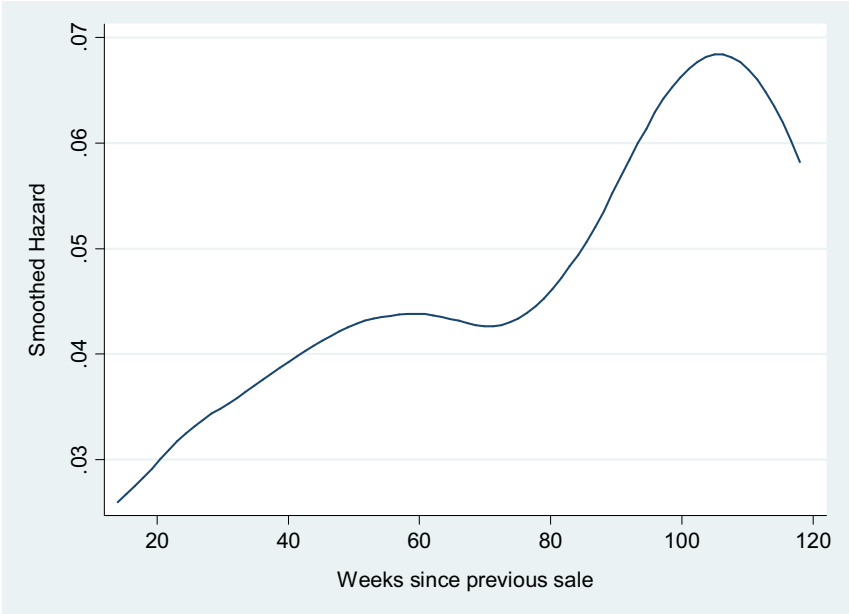Figure 15: Single spell conditional baseline hazard function of sales

Figure 16: Multiple spell conditional baseline hazard function of sales

Thus, we find an upward-sloping pattern to the hazard rate of sales, indicating that in general the longer the product is on the shelves, the more likely it is being on sale (positive time-dependence). The key advantage of the semi-parametric regression is that it provides quantification of the heterogeneity (by retailer, format and brand status) observed using the non-parametric methods.

Results of conditional hazard functions are reported in Table 7. In the interest of brevity we focus here on the estimates of multiple spell model but report estimates from the single spell model for comparison. Suffice to say results are most different for the discounters, since it is they who have the most highly promoted products. 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 those comparators. Results suggests that two of the soft discounters (Somerfield and Kwik Save) behave similarly to Tesco with regards to sales, only Safeway displaying a (22%) higher likelihood of a sale. Both of the other mainstream retailers have a significantly lower hazard rate of sales than Tesco; ASDA being half as likely to have a sale than Tesco, other things reaming 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. A more surprising finding from Table 7 is that Waitrose, the luxury food retailer of the seven in the sample, has a hazard rate greater than all retailers except Safeway, and 15% 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 22%. Recall Hosken and Reiffen (2007) propose a state-dependent model of where sales behaviour increases with perishability. While our results suggest differences by format, it appears that perishable goods are less - not more - likely to go on sale than more durable formats. Note that since our sample does not

Table 7: Semi-parametric estimation results

REGULAR PRICE SPELLS

	Single Spell Conditional hazard function		Multiple Spell conditional hazard function	
	Hazard Ratio	Robust Standard Error	Hazard Ratio	Robust Standard Error
SAINSBURY	0.905	(0.063)	0.853***	(0.049)
ASDA	0.417***	(0.070)	0.502***	(0.082)
SAFEWAY	1.869***	(0.140)	1.222**	(0.098)
SOMERFIELD	1.390***	(0.099)	1.009	(0.059)
KWIK SAVE	1.372***	(0.105)	1.001	(0.068)
WAITROSE	1.116	(0.092)	1.152**	(0.072)
AMBIENT	0.933	(0.061)	0.861***	(0.050)
FROZEN	0.698***	(0.059)	0.854**	(0.060)
CHILLED	1.071	(0.089)	0.924	(0.055)
FRESH	0.778***	(0.057)	0.779***	(0.052)
PRIVATE LABEL	0.713***	(0.058)	0.828***	(0.047)
NUMBER OF SPELLS	-		1.162***	(0.018)
ORDER OF SPELLS	-		0.973*	(0.015)
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. Single spell conditional hazard function estimates the regular price spells conditional on retailer, format, brand status and seasonal dummies; multiple spell conditional hazard function is augmented by accounting for multiple spell effects using covariates describing the UPC's sales profile and delayed entry spells. Estimated monthly seasonal dummies are not reported for simplicity.

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. The results do however confirm the theoretical prediction of Lal (1990), that posits that sales are less likely to occur in private labels than brands. The estimate in Table 7 suggests private labels are 17% less likely to be promoted compared to national brands.

Turning to the effect of multiple spells on the hazard of sales results that while the number of spells is positively related to sales and highly (1% level) significant and effect of spell order is slightly negative and weakly significant (10% level). These results confirm that sales are indeed not randomly assigned but tend to be concentrated in certain UPCs (reflecting factors such as brand competition) and 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).

5.3 Hazard models by supermarket

One of the key findings from the previous section is that differences in the hazard of sales is most keenly affected by retailer. As a result we drill down in to the data and estimate a separate hazard function for each retail chain using the same multiple spell model as previously. [13] Baseline hazard functions for each are graphed in Figure 17. 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 hazard rate 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 'hi-lo' sales strategy whereas the flat baseline hazard of ASDA is consistent with an EDLP pricing

Figure 17: Semi-parametric result: Sales patterns across retailers
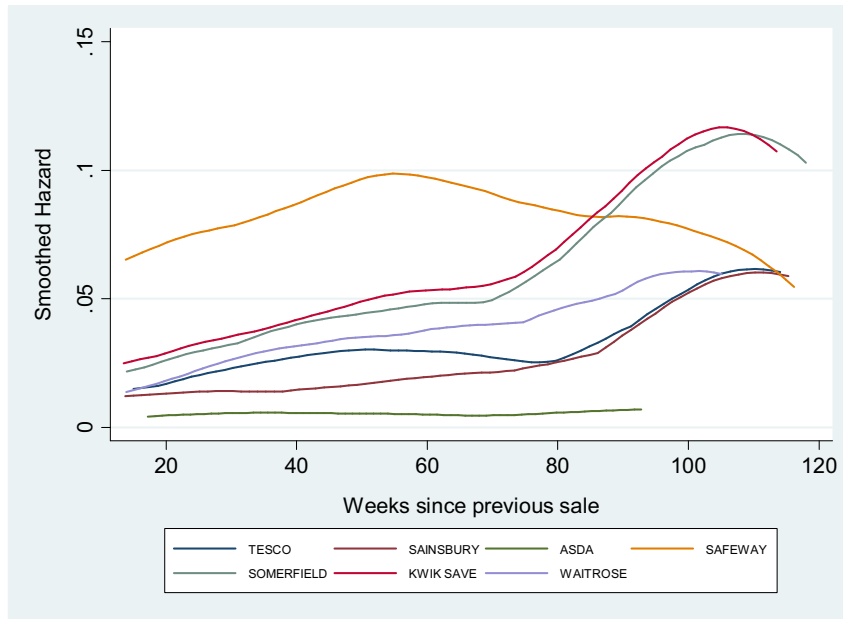
strategy.

Estimates of the parameters from the duration models associated with these baseline hazards are reported in Table 8. There are three 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.[14] Second, and in stark 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. Tesco) show little difference across the formats, others (e.g. Safeway) a lot. Whereas frozen products are the most promoted of all formats in Waitrose they are least promoted in Safeway. 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

---

[13]Due to no-convergence of the ML estimator we do not control seasonal effects when we estimate the hazard function in Asda.

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

in (Safeway and Somerfield). Third, differences between the frequency of sales between brand and private label products are apparent in most retail chains; all statistically significant relationships suggest that private label products are less likely to be promoted than brands.

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, 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 nonparametric and 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 (2000). While this result relates

Table 8: Semi-parametric estimation: Sales patterns across supermarkets

	(1) TESCO	(2) SAINSBURY	(3) ASDA	(4) SAFEWAY	(5) SOMERFIELD	(6) KWIK SAVE	(7) WAITROSE
AMBIENT	1.060	1.038	0.729	0.582***	0.790**	0.856	1.374**
	(0.125)	(0.129)	(0.415)	(0.063)	(0.093)	(0.127)	(0.172)
FROZEN	1.123	0.886	1.191	0.680***	1.213	0.671**	1.965***
	(0.193)	(0.126)	(0.691)	(0.079)	(0.234)	(0.117)	(0.368)
CHILLED	1.189	1.196	0.955	0.632***	1.166	1.154	1.270
	(0.307)	(0.150)	(0.677)	(0.056)	(0.117)	(0.206)	(0.407)
FRESH	1.299	1.369***	1.115	0.594***	0.781**	0.845	1.590***
	(0.222)	(0.147)	(0.416)	(0.084)	(0.089)	(0.106)	(0.238)
PRIVATE LABEL	0.727**	0.921	2.066	0.855	0.647***	0.755*	0.747*
	(0.116)	(0.084)	(1.464)	(0.093)	(0.096)	(0.109)	(0.114)
NUMBER OF SPELLS	1.307***	1.478***	3.651***	1.150***	1.177***	1.226***	1.378***
	(0.036)	(0.057)	(0.607)	(0.016)	(0.033)	(0.026)	(0.054)
ORDER OF SPELLS	0.934	0.747***	0.337***	1.060***	0.885***	0.886***	0.861**
	(0.050)	(0.043)	(0.086)	(0.023)	(0.039)	(0.029)	(0.063)
N	527	600	101	1151	820	739	365

Note: hazard ratio and robustness standard error are reported here. Robustness standard error is presented in ( ). ***, ** and * denote that the null hypothesis (of unity in the hazard ratio) is rejected at the 1% , 5% and 10% level. Seasonal dummies are not reported for simplicity.

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.

Our analysis also addresses the issue of heterogeneity in terms of supermarket, product format (perishability) and brand status. 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., Burriel, P. and Hernando, I. (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", *Journal of Business*, vol. 53, no. 3, pp. S51–S67

Baumgartner, J., Glatzer, E., Rumler, F. & Stiglbauer, A. (2005), "How frequently do consumer prices change in Austria? Evidence from micro CPI data", *European Central Bank Working Paper Series, No.523*

Berck, P., Brown, J., Perloff, J.M. & Villas-Boas, S.B. (2008), "Sales: Tests of Theories on Causality and Timing", *International Journal of Industrial Organization*, vol. 26, no. 6, pp. 1257-1273

Bunn, P. & Ellis, C. (2009), "Price-setting behaviour in the United Kingdom: a microdata approach", *Quarterly Bulletin 2009 Q1, Bank of England*, vol. 49, no. 1, pp. 28-35

Cavallo, A. (2009), "Scraped Data and Sticky Prices: Frequency, Hazards, and Synchronization", *Working paper in Harvard University*

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

Cox, D. R. (1972), "Regression models and life tables", *Journal of the Royal Statistical Society. Series B (Methodological),* vol. 34, no. 2, pp. 187-220.

Cox, D. R. (1975). "Partial likelihood", *Biometrika*, vol. 62, no.2, pp. 269-276.

Degraba, P. (2006), "The Loss Leader is A Turkey: Targeted Discounts From Multi-product Competitors", *International Journal of Industrial Organization*, vol. 24, no. pp. 613-628

Dias, M., Dias, D. & Neves, P. (2004), "Stylised Features of Price Setting Behaviour in Portugal: 1992-2001", *European Central Bank working paper series No. 332*

Fougère, D., Le Bihan, H. & Sevestre, P. (2007), "Heterogeneity in Consumer Price Stickiness", *Journal of Business and Economic Statistics*, vol. 25, no. 3, pp. 247-264

Guimaraes, B. & Sheedy, K. D. (2011) "Sales and monetary policy", *American Economic Review*, vol. 101, no. 2, pp. 844-876.

Hosken, D. & Reiffen, D. (2004), "Patterns of Retail Price Variation", *The RAND Journal of Economics*, vol. 35, no. 1, pp. 128-146

Hosken, D., Reiffen, D. & (2007), "Pricing Behaviour of Multi-product Retailers", *The B.E. Journal of Theoretical Economics*, vol. 7, no. 1, pp. 1-43

Ikeda, D. & Nishioka, S. (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.* Available from: University of Essex, Institute for Social & Economic Research,
Web site: http://www.iser.essex.ac.uk/files/teaching/stephenj/ec968/pdfs/ec968lnotesv6.pdf
[Accessed: March 23, 2010]

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

Kiefer, N.M. (1988), "Economic Duration Data and Hazard Functions", *Journal of Economic Literature*, vol. 26, no. 2, pp. 646-679

Klenow, P. J. & Kryvtsov, O. (2008) "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent US Inflation?", *The Quarterly Journal of Economics*, vol. 73, no. 3, pp. 863-903

Klenow, P. J. & Malin, B. A. (2010), "Microeconomic Evidence on Price-Setting", *NBER Working Paper No. 15826*

Lal, R. (1990), "Price Promotions: Limiting Competitive Encroachment", *Marketing Science*, vol. 9, no. 3, pp. 247-262

Lloyd, T.A., Morgan, C.W., McCorriston, S. & Zgovu, E. (2011), "Do Sales Matter? Evidence from UK Food Retailing", *The University of Nottingham Discussion Paper in Economics, No. 11/01*

Meyer, B.D. (1990), "Unemployment Insurance and Unemployment Spells", *Econometrica*, vol. 58, no. 4, pp. 757-782

Nakamura, E. & Steinsson, J. (2008), "Five Facts about Prices: A Reevaluation of Menu Cost Models", *Quarterly Journal of Economics*, vol. 123, no. 4, pp. 1415-1464

Nakamura, E. & Zerom, D. (2010), "Accounting for Incomplete Pass-Through", Review of of Economic Studies, vol. 77, no. 3, pp. 1192–1230

Pesendorfer, M. (2000), "Retail Sales: A Study of Pricing Behaviour in Supermarkets", *Journal of Business*, vol. 75, no. 1, pp. 33-66

Richards, T.J. (2006) "Sales by Multi-product Retailers." *Managerial and Decision Economics*, vol. 27, pp. 261-277

Shilony, Y. (1977), "Mixed Pricing in Oligopoly", *Journal of Economic Theory*, vol. 14, no. pp. 373-388

Sobel, J. (1984), "The Timing of Sales", *The Review of Economic Studies*, vol. 51, no. 3, pp. 353-368

Varian, H. (1980), "A Model of Sales", *The American Economic Review*, vol. 70, no. 4, pp. 651-659

Veronese, G., Fabiani, S., Gattulli, A., & Sabbatini, R. (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 18: Examples of types of censored spells
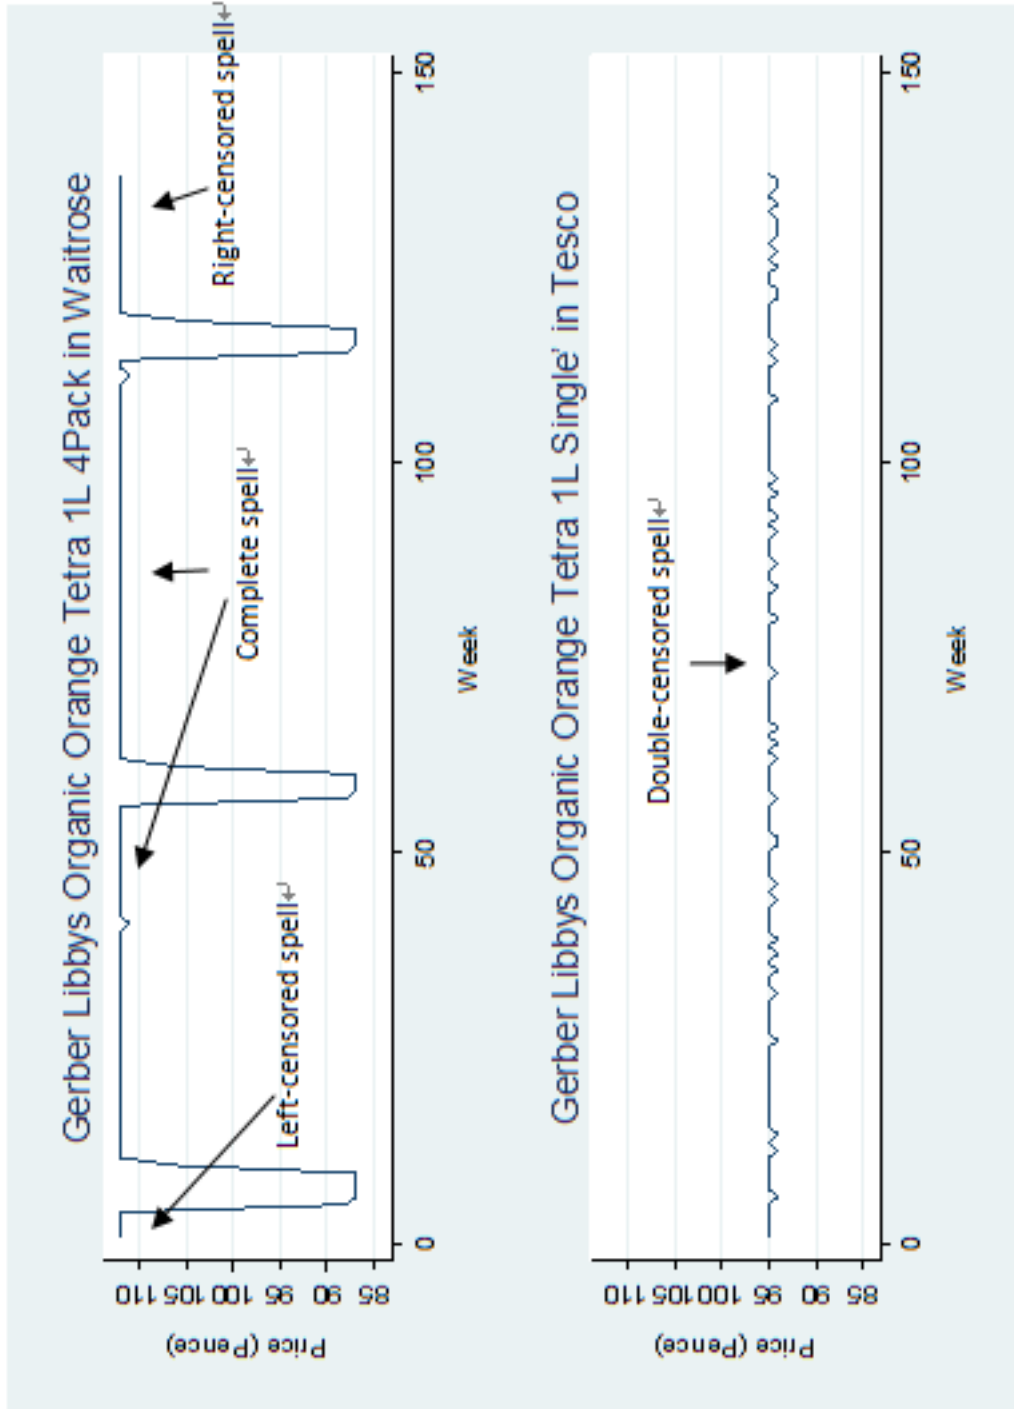
Note: the top price series illustrates the left, right-censored and complete spells; the bottom one demonstrate a double-censored spell.