

Regularity in Individual Shopping Trips: Implications for Duration Models in Marketing*

Govert E. Bijwaard[†]

Econometric Institute
Erasmus University Rotterdam

Econometric Institute Report EI 2005-07

Abstract

Most models for purchase timing behavior of households do not take into account that many households have regular and non-shopping days. I propose a statistical model for purchase timing that exploits information on the shopping days of households. It delivers forecasts for the number of purchases in the next period and for the timing of the first and consecutive purchases. Purchase occasions are modeled in terms of a counting process, which counts the recurrent purchases for each household as they evolve over time. I illustrate the model for yogurt and detergent purchases and highlight its useful managerial implications.

Key words: Purchase timing; regular shopping days; non-shopping days; counting process; mixed proportional hazard.

*I thank Dennis Fok and Rutger van Oest for data support. Thanks also to Philip-Hans Franses and Ulf Böckenholt for helpful discussions and comments. This research is financially supported by the Netherlands Organisation for Scientific Research (NWO).

[†]Erasmus University Econometric Institute, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands; Phone: (+31) 10 40 81424; Fax: (+31) 10 40 89162; E-mail: bijwaard@few.eur.nl

1 Introduction

There exists a large number of empirical literature that deals with the analysis of purchase timing decisions of households using statistical models. These models characterize households' temporal decisions of when to purchase a particular product over time. Such analysis can be relevant for store managers to understand the influence of the marketing mix on purchase timing. This information can be relevant for active stock management and for the timing of promotional actions. A particular phenomenon first observed by Dunn et al. (1983) for their toilet tissue data that has been consistently noted since is that histograms of repeated purchase times exhibit spikes at multiples of 7 days. Kahn and Schmittlein (1989) found a similar pattern in cracker data, Jain and Vilcassim (1991) for coffee data and Chiang et al. (2001) for a number of other product categories. The consistency of this weekly spike pattern suggests that it may be typical for all frequently purchased consumer goods. Kahn and Schmittlein (1989) show that these weekly patterns may simply be the by-product of consumers' regular shopping trip schedules. I found the same striking pattern in my data on yogurt (see Figure 1) and dry detergent (see Figure 2).

Another aspect of the regularity of shopping days is that some households only shop on a few days of the week. This implies that on the other days of the week they never shop. In my data around 3.5% of the households never shop on Sundays. More than 5% of the households have at least one day of the week they never go shopping. To my knowledge this has not been noticed before. Both the knowledge of the regular and the non-shopping days is valuable information in predicting purchase behavior of households and the timing of marketing actions. When you know a particular household purchased yogurt previous Saturday the probability it purchases yogurt again is higher this Saturday than this Sunday. Similarly, recognizing that a household never goes shopping on Sundays implies that it has zero probability to purchase any product on Sundays.

On the aggregate level, regularity in shopping days is reflected in a weekly distribution of purchases. If for many households Sunday is a non-shopping day and Saturday is a favorite shopping day the number of purchases of yogurt and detergent will be higher on Saturdays (see Figure 3). Another important aggregate regularity in purchase timing is seasonality (see

Figure 4). In my data the monthly purchase pattern for both yogurt and detergent purchases is apparent, but differs between the two products. Accounting for the aggregate regularity in purchase timing behavior also has important implications for the prediction of purchases and the timing of marketing promotions.

The currently applied models for purchase timing behavior usually describe the time until the next purchase by a (mixed) proportional hazard model (MPH) (see, for example, Gupta (1991), Jain and Vilcassim (1991), Vilcassim and Jain (1991) and Helsen and Schmittlein (1993) among others for an application, and see Seetharaman and Chintagunta (2003) for a recent overview). The hazard is the instantaneous probability of making a purchase in a product category conditional on the elapsed time since the household's previous purchase (in the sequel in call this gap-time). An important property of these models is that, after the purchase has been made, time is reset to zero and a new duration spell starts, independent of the previous duration. In fact, each interpurchase duration of a household is modeled separately, using the same type of hazard function.

Recently, Bijwaard et al. (2003) have shown that the gap-time model for interpurchase durations has serious limitations, in particular if one wants to use the model for managerial decisions. First of all, it neglects valuable information about the purchase history of a household's decisions. Second, it is very complicated with the gap-time models to predict whether a household will make two or more purchases in a certain period. Third, it is very cumbersome with these models to account for aggregate regularities in shopping behavior. Finally, these models cannot take information on non-shopping days into account.

If time runs on calendar-time instead of the time since the previous purchase, the hazard (or the intensity) is the instantaneous probability of making a purchase in a product category conditional on all what happened since the start of the observation of the household. Then the hazard can be decomposed into four components, the calendar duration dependence, the gap duration dependence, the covariate function and the unobserved heterogeneity. This amounts to decomposing the baseline hazard into a factor capturing calendar-time effects (such as day-of-the week effects and seasonal effects) and a factor capturing the intrinsic cyclic pattern of purchase timing. In fact, if there is no calendar-time duration dependence (the component is constant)

the model reduces to the gap-time MPH model.

In a calendar-time model it is very easy to implement non-shopping days, in particular when the model is formulated in terms of counting processes. A counting process counts the purchases for each household as they evolve over time. The counting process formulation concerns the situation where a household is at risk of purchasing an item by defining an at-risk-indicator. If, for example, a household is known not to shop on Sundays, the at-risk-indicator is zero for every Sunday in the observation window for that particular household. Bijwaard et al. (2003) do not explicitly take these non-shopping days into account. This paper extends their analysis of the calendar-time model to account for regular and non-shopping days.

Wheat and Morrison (1990) postulate that the "whether or not to buy" specification is superior to the "when to buy" specification in purchase timing models. The calendar-time model simultaneously describes both decisions. This is very important from a managerial point of view. With such a model the manager can answer both the question of how many households purchase the product as what is the probability of a purchase by a household within a given period. A model that accounts for regular and non-shopping days will improve the prediction of both questions, because it incorporates useful information on the shopping patterns of households.

The outline of the paper is as follows. In the next section I present some stylized facts on regular and non-shopping days in the purchase of frequently bought consumer goods. Section 3 describes the counting process view on repeated events and gives an overview of its features and their implications for modeling purchase timing. In Section 4, the model specific for repeated purchase times is described and a discussion on how this can be implemented and which type of forecasts it can deliver is given. In Section 5 the implications of taking regular and non-shopping days into account in the analysis of household purchase timing behavior are illustrated for two product categories: laundry detergent and yogurt. Graphs are used to show how the model outcomes can be used for managerially relevant purposes. Section 6 concludes and provides avenues for further research.

2 Motivation

It has been empirically documented that households tend to shop on a weekly basis (see Dunn et al. (1983), Kahn and Schmittlein (1989), Jain and Vilcassim (1991) and Chiang et al. (2001)). Hence, if most purchases are made at seven-day intervals or multiples thereof, we would expect weekly spikes in the histogram of in the interpurchase times for all households. Another less visible regular pattern is that some households do not shop on particular days of the week. These weekly household shopping trip pattern also has implications for the distribution of purchases over the week. If many households are Saturday-shoppers, the total number of purchases on Saturdays will be higher than on other days of the week. The purchase behavior can also fluctuate over the year, with high purchase months and low purchase months. These issues raise the question whether a typical consumer can be persuaded to make unplanned trips just for promoted products. From a managerial point of view it is even more important to know whether a consumer can be persuaded to purchase early and more often.

Although extensive academic research have been done on purchase timing behavior, the regularity in the timing has hardly been addressed. Only Chiang et al. (2001) incorporate the weekly spike phenomenon in a gap-time model. Bijwaard et al. (2003) consider the seasonal effects on purchase timing. But the issue of non-shopping days has not been investigated at all. I present a model that capture all these regularities in purchase behavior. Before I turn to the model set-up I describe the regular patterns in my data.

The data I use is part of the so-called ERIM database, which is collected by A.C. Nielsen. This scanner data spans the years 1986 to 1988, and the particular subset I use concerns purchases of dry detergent and yogurt by households in Sioux Falls (South Dakota, USA). Figure 1 and Figure 2 depict the histograms of interpurchase times for yogurt and detergent. These pictures clearly show weekly spikes in the interpurchase time for both products.

– **put figure 1 about here**–

– **put figure 2 about here**–

From the yogurt and detergent purchase timing data only we cannot distill the non-shopping days, because households may have gone shopping on days they did not purchase one of these two

products. The non-shopping days can, however, be distilled from the larger data on all purchase occasions of the households. The ERIM data contains purchase behavior of the households on many product goods. I use the data from the purchase occasions of the following 10 product-groups: catsup, dry detergent, ketchup, margarine, peanut butter, dry soup, sugar, tuna, toilet tissue and yogurt. The non-shopping days for the households are then the days of the week they never purchased any of the above 10 products. Table 1 reports the distribution of non-shopping days for the households purchasing yogurt and detergent.

– put table 1 about here–

About 4.2% of the yogurt consumers and 6.4% of the detergent consumers in my sample have non-shopping days. Some households have more than one non-shopping day (one household in the detergent data only shops on Saturdays). Sunday is the most common non-shopping day. About 3% of the yogurt consumers and about 4% of the detergent consumers never goes shopping on Sundays.

The identification of a non-shopping day may be destroyed by an one-time visit on a particular week day. The table therefore also reports the rarely shopping days, where I denote the days on which a household makes less than 5% of its purchases as a rarely shopping day. These rarely shopping days are, of course, more common than non-shopping days. Little more than half of the households do not have a rarely shopping day, about 28% of the households have one rarely shopping day and about 18% of the households have more than one rarely shopping day. Sunday is the day many people rarely go shopping.

Another prominent feature of purchase timing behavior largely overlooked is the purchase pattern within a week and over the year. Figure 3 gives the distribution of the purchases over the week. Both yogurt and detergent purchases have a similar pattern over the week. The purchases are more frequent by the end of the week (Thursday - Saturday) with a culmination on Saturday. The skewed distribution is more pronounced for detergent purchases.

– put figure 3 about here–

The detergent purchases are observed from September 1986 till August 1988 and the yogurt purchases from November 1986 till August 1988. So I have purchase information in winter and

in spring for two years and in autumn only for one year. An additional observational issue is that the observation of the purchase timing behavior for each household only starts after the first purchase of the particular product. For some households, especially for the yogurt purchases, the first purchase occurs beyond halfway the observation window. In calculating the distribution of the purchases over the months I adjust for these observational issues. Figure 4 gives the resulting distribution.

– put figure 4 about here–

The distribution over the year differs between the two products. Yogurt is purchased most frequently in February and March and purchased least frequently in August and November. Detergent is, on the other hand, purchased most frequently in August and January and purchased least frequently in April and May.

3 Basic features of counting processes

The theory of counting processes provides a general framework in which survival and repeated durations may be analyzed. The counting process framework, although dating back to the early 1980's, is not common in marketing. Bijwaard et al. (2003) are the first to show the advantages of the counting process framework for analyzing purchase timing behavior. The seminal book of Andersen et al. (1993) provide an excellent survey of counting process theory. It is, however, rather technical. A less technical survey is given in Klein and Moeschberger (1997). Therneau and Grambsch (2000) provide a practitioner's guide for modeling survival and repeated events in a counting process framework using SAS and S-Plus. The recent book by Kalbfleisch and Prentice (2002) give an update of the current knowledge on counting processes.

To give context to the counting process approach, I first briefly review the familiar or "traditional" description of interpurchase timing. The key variables in such duration analysis are the gap time, the elapsed time since the previous purchase, and the censoring indicator. The gap time for household i for the j^{th} purchase, G_{ij} is a random variable following some p.d.f. $f(g)$ and c.d.f. $F(g)$. In duration analysis it is convenient to consider the hazard rate $\lambda(g) \equiv f(g)/(1 - F(g))$, which is the conditional probability that a household made a purchase at gap time g , given that

it did not make a purchase during the interval $(0, g)$. The censoring indicator is important in duration analysis because the purchase behavior is usually recorded within the boundaries of an observation window. As a result the duration from the last purchase before the end of the observation window till the next purchase after this end is only partially observed. The only piece of information is that the gap time is at least till the end of the observation window. The censoring indicator $\delta_{ij} = 1$ if the the purchase is beyond the observation window.

To estimate a gap time model a particular parametric form for the hazard is chosen. A common choice is the proportional hazard model or Cox model in which the hazard is multiplicative in a baseline hazard and a regression factor capturing the effects of observed characteristics of the households and the marketing mix variables, that is

$$\lambda(g) = \lambda_0^g(g) \exp(\beta' X_i(g)). \quad (1)$$

The first component, $\lambda_0^g(g) > 0$, is the (gap-time) duration dependence that captures the common dependence of the hazard on the elapsed time since the previous purchase. The second component adjusts $\lambda_0^g(g)$ up or down proportionately to reflect the effect of the, possibly time-varying, measured covariates $X_i(g)$.

The counting process formulation replaces the pair of variables (G_{ij}, δ_{ij}) with the pair of functions $(N_i(t), Y_i(t))$. The *counting* process $N_i(t)$ is a stochastic process which describes the number of purchases of household i in the interval $[0, t]$ as (calendar) time proceeds. The *at-risk indicator* $Y_i(t)$ indicates whether a household is under observation, that is at risk of purchasing, at time t . Note that time is in calendar time t and not in gap time. As Bijwaard et al. (2003) point out formulating the repeated purchase behavior in calendar time has many advantages over a gap-time representation. Now the household makes its purchases at the calendar times T_{ij} , with is related to the gap-times by $G_{ij} = T_{ij} - T_{ij-1}$.

The counting process has only jumps of size $+1$. This implies that for each household only one purchase can be made at any point in time. The counting process is governed by its random *intensity* process $\lambda_i(t)$. If we consider a small interval $(t - dt, t]$ of length dt , then $\lambda_i(t)dt$ is the conditional probability that $N_i(t)$ jumps in that interval given all that has happened until just before t . The close connection between the "traditional" framework and the counting process framework is immediately apparent from the similarity between the hazard rate and the intensity

process. In fact these concepts are the same in survival analysis when the event of interest can only occur once.

Let $dN_i(t)$ denote the increment of $N_i(t)$ in the small interval and let \mathcal{H}_{it} denote the information set of household i up to, but not including, t . This *history* process includes a complete specification of the path of the counting process on $[0, t)$ and it includes all other events implicitly or explicitly included in the model which have happened before time t . Hence, the history for household i also includes the occasions when this household was at risk. We can write the conditional probability that household i makes a purchase at time t given its history as

$$\Pr[dN_i(t) = 1 | \mathcal{H}_{it}] = Y_i(t)\lambda_i(t)dt. \quad (2)$$

If the days of the week a household does not shop are known the at-risk indicator is zero for those days. This sets the intensity of purchasing on these "non-shopping" days effectively to zero. Other information on the observation of households, like being away for holidays or hospital visits, can, if known, also be included in the at-risk indicator.

The expected number of purchases of household i at time t from the beginning of the observation window is the *cumulated intensity*

$$\Lambda_i(t) = \int_0^t \lambda_i(s) ds. \quad (3)$$

This implies that,

$$\mathbb{E}[N_i(t) | \mathcal{H}_{it}] = \Lambda_i(t). \quad (4)$$

Using (4) we can derive the expected number of purchases within any given interval, say from t_1 until t_2 (with $t_2 > t_1$), which is equal to $\Lambda_i(t_2) - \Lambda_i(t_1)$.

In a calendar-time set-up forecasting the number of purchases made by all households within a specific period, say 2 weeks, is very easy. That is, we forecast the counting process of purchasing. The time till the first purchase is the number of days or weeks till the household's counting process makes a jump. The total number of forecasted purchases and the percentage of households purchasing more than once within a given period can be derived directly from the forecasted counting processes.

4 Modeling the purchase timing

In general, the intensity function will depend on the current and past value of the marketing-mix variables and observed characteristics of the households. Denote these possible time-varying covariates by $x_i(t)$ for household i at time t and let $\bar{X}_i(t) = \{x(s) : s \leq t\}$ denote the complete path of the covariate vector up to time t . This path contains, among other things, the whole set of observed marketing mix variables for each household. For a time-constant covariate, like the household size, we have that $\bar{X}_i(t) = X_i$. The intensity process for the i^{th} household given these covariates is

$$\Pr[dN_i(t) = 1 | \mathcal{H}_{it}] = Y_i(t) \lambda_0(t) \exp(\beta' X_i(t)) dt, \quad (5)$$

where $\lambda_0(t)$ is the *baseline intensity* and β is a parameter vector. The covariates have a multiplicative effect on the intensity through a log-linear regression function $\exp(\beta' X_i(t))$. That is, if $\beta > 0$ an increase in $x_i(t)$ results in a proportional increase in the intensity and therefore leads to a decrease in the expected interpurchase time.

This model is the natural extension to repeated events of the familiar proportional hazard model for survival data. In such a set-up, $\lambda_0(t)$ is the same for all households and it is called the *baseline hazard*. On a calendar timescale, the baseline intensity represents seasonal and day-of-the-week effects. Hence, if on a particular day of the week much more yogurt is purchased than on other days, this will appear in the intensity process as peaks of multiples of 7 days. The difference between the proportional hazard model and this model lies in the definition of the risk indicator $Y_i(t)$. The proportional hazard model for survival data models the hazard of a terminal event. One can experience such an event at most once and therefore the at-risk indicator is set to zero after the event. Here a household remains at risk of repurchasing as long as it has not left the study. Non-shopping days are accounted for by putting the at-risk indicator at zero for those days.

In repeat purchase timing the duration dependence of the time since the previous purchase is considered very important (see e.g. Seetharaman and Chintagunta (2003)). Fortunately, this gap-time duration dependence can also be included in the calendar-time model by pre-multiplying

the intensity by the gap-time duration dependence, i.e.

$$\Pr[dN_i(t) = 1 | \mathcal{H}_{it}] = Y_i(t) \lambda_0(t) \lambda_0^g(t - T_{iN_i(t)-1}) \exp(\beta' X_i(t)) dt, \quad (6)$$

where $T_{iN_i(t)-1}$ is the (calendar) time of the previous purchase and $t - T_{iN_i(t)-1}$ is, thus, the gap-time since the previous purchase. Now the model can exhibit both a calendar-time and a gap-time duration dependence. Note that if $\lambda_0(t) \equiv 1$, a model without calendar-time duration dependence, then we have an alternative representation of a gap-time model given in (1).

I use the restricted version of the Box-Cox formulation in Jain and Vilcassim (1991) to model the gap-duration dependence. One advantage of this formulation is that it nests most of the frequently used gap-time hazard rate formulations that have been used in the literature, like the Weibull, Gompertz and Erlang hazard functions. Another advantage of this formulation is that it allows for non-monotonic gap duration dependence. This gap-time duration dependence is given by

$$\lambda_0^g(g; \alpha) = \exp(\alpha_1 g + \alpha_2 g^2 + \alpha_3 \ln(g) + \alpha_4 I_w(g)). \quad (7)$$

The regular shopping trips are captured by an indicator function $I_w(g)$ that is one if the time since the previous purchase is a multiple of a week and zero otherwise. If $\alpha_4 > 0$ the weekly shopping trip indicator induces upward jumps at weekly intervals in the gap-duration dependence. This is exactly what the frequency distribution of interpurchase times in Figure 1 and Figure 2 suggests.

A more flexible modeling approach is to use the method developed by Cox (1972) (and extended to repeated events by Andersen and Gill (1982)). The Cox method treats the baseline intensity as a nuisance parameter and estimates it non-parametrically. Chiang et al. (2001) advocate this method to accommodate the regular weekly spikes. A major drawback of this method is, however, that the baseline hazard can only be obtained for observed durations. This approach is therefore less suitable for predicting future purchases. For this same reason Bijwaard et al. (2003) use parametric models for prediction, despite that their Cox-type model has a better fit. Furthermore, Seetharaman and Chintagunta (2003) conclude that the additional flexibility afforded by the nonparametric baseline hazard hardly alter the estimated coefficients or the fit of the (gap-time) purchase timing models.

4.1 Unobserved heterogeneity

If there is interdependence of the repeat purchases due to omitted covariates or household-specific effects, like being a heavy-user or not, the parameter estimates may be biased and/or the estimated covariance matrix provides invalid standard errors. To correct for this, one may use a robust covariance matrix estimate, see Lin and Wei (1989). Another approach is to explicitly model the household-specific effects using unobserved heterogeneity. In Cox survival models this kind of model is called the mixed proportional hazard model. The intensity process of household i at time t is now given by

$$Y_i(t)v_i\lambda_0(t)\lambda_0^g(t - T_{iN_i(t)-1}; \alpha) \exp(\beta' X_i(t)), \quad (8)$$

where the $v_i > 0$ are i.i.d. random variables with distribution function $G(v)$ that differ among the households but remain the same for each purchase of a particular household. The Gamma distribution with mean one and variance σ^2 is most often chosen to represent the unobserved heterogeneity. The density function of the gamma distribution is given by

$$G(v) = \frac{v^{(\sigma^{-2}-1)} \exp(-v\sigma^{-2})}{\Gamma(\sigma^{-2})\sigma^{2\sigma^{-2}}}, \quad (9)$$

where $\Gamma(\cdot)$ is the gamma-function (see e.g. Lancaster (1990)). Hence, large values of σ^2 reflect a greater degree of heterogeneity among households and a stronger association within household purchases. The log-likelihood for this model for repeated purchase timing is obtained by integrating out the unobserved heterogeneity and is given by

$$\begin{aligned} L(\beta, \alpha, \theta, \lambda_0, \sigma^2) = & \sum_{i=1}^n \left\{ k_i \ln(\sigma^2) - \ln \Gamma\left(\frac{1}{\sigma^2}\right) + \ln \Gamma\left(\frac{1}{\sigma^2} + k_i\right) \right. \\ & - \left. \left(\frac{1}{\sigma^2} + k_i\right) \ln \left[1 + \sigma^2 \sum_{j=1}^{k_i+1} \int_{T_{ij-1}}^{T_{ij}} \lambda_0(s) \lambda_0^g(s - T_{ij-1}) \exp(\beta' X_i(s)) ds \right] + \right. \\ & \left. \sum_{j=1}^{k_i} [\beta' X_i(T_{ij}) + \ln(\lambda_0(T_{ij})) + \lambda_0^g(T_{ij} - T_{ij-1})] \right\}, \quad (10) \end{aligned}$$

where T_{i0} is the first time household i is observed and $T_{ik_i+1} = T_e$ is the end of the observation period (which the same for all households). Each household may have a different number of purchases, denoted by k_i . Note that we observe all purchase times except for the last one, which

occurs after the observation period has ended. If we assume that the calendar-time duration dependence $\lambda_0(t)$ has some parametric form, we can directly maximize the likelihood function (10) to obtain parameter estimates. Estimates of the variance-covariance matrix are obtained by evaluating the inverse of the information matrix in the parameter estimates.

Estimation of the parameters of the model is straightforward in SAS and STATA if a counting process input style is used. In this input style each household is represented as a set of rows with time intervals that ends in a purchase or a change in one of the time-varying covariates: (entry time, first change], (first change, second change], \dots , (m th change, end of observation window]. The rows of data consist of data observations, each of which contains (fixed) covariate values X , a status indicator $dN_i(t)$ (1 = repeat purchase, 0 = censored; that is, no repeat purchase at the end of the time interval), along with the time intervals for which this information applies.

5 Empirical results

In this section I illustrate the effect of accounting for regular and non-shopping days in a calendar-based purchase timing model on scanner data for yogurt and detergent purchases. The yogurt data covers a period of 91 weeks and the detergent data covers a period of 97 weeks. I select those households who are purchasing only the top brands, that are the brands that are sold frequently enough to build an entire history of the marketing efforts. Furthermore, I restrict the analysis to households who purchased yogurt or detergent at least three times in the observation window.¹ This results in purchase timing information on 598 households for yogurt purchases and on 624 households for detergent purchases. The data contains information on price, in-store display, and newspaper feature advertisements at the brand level for each store and each week. The marketing instruments are constant during a week, where the week is defined from Wednesday to but not including Wednesday.

For each purchase occasion I know the day and the volume purchased. Furthermore, for each week I know the shelf price (dollars/32oz.) of all brands and which brands are featured or on display. As I do not consider the brand choice the marketing information is aggregated

¹A heavy-user bias is not an issue here because I incorporate the average purchase volume as an explanatory variable. This variable adjusts for high purchase volume.

over stores and brands. I use household-specific weights in this aggregation. Following Gupta (1991), I use household-specific volume brand shares to aggregate over brands. Aggregation over stores is carried out using household-specific store weights. Thus, for each household I only use data on the relevant store and brand options. The feature and display variables now represent the percentages of stores (relevant for the specific household) featuring a brand, or having the brand on display. Next to this information on marketing instruments, I use the household size, household income and the volume purchased at the previous purchase occasion. The latter variable is a proxy for inventory.

For the duration dependence between consecutive purchases, the gap-time duration dependence, the specification in (7) is used. Note that the baseline intensity in the calendar-time model represents the calendar-time duration dependence, such as seasonal- and day-of-the-week effects, and not the duration dependence between re-occurring purchases. I include day- and month dummies to capture these effects. In sum, if Saturday is a favorite shopping day for many households this is reflected by peaks at multiples of one week in the calendar time duration dependence. The possibility that households purchase a particular product on a regular basis, say, every week is reflected in a gap-time duration dependence with indicators at multiples of one week since the previous purchase α_4 .

I consider five different specifications based on the intensity (8) and likelihood function (10) for modeling the purchase timing behavior:

- gap-time model without accounting for regular shopping (Gap 0, $\alpha_4 = 0$ in (7));
- gap-time model with an indicator for regular shopping (Gap 1);
- calendar-time model without accounting for regular shopping (Calendar 0);
- calendar-time model with an indicator for regular shopping (Calendar 1)
- calendar-time model that takes regular and non-shopping days into account (non-shopping, ns).

The gap-time models are included to check the findings of Bijwaard et al. (2003) that the calendar-time models outperform the commonly used gap-time models in the analysis of purchase

timing behavior. The model that accounts for non-shopping days only differs from the Calendar 1 model in the at-risk indicator of purchasing. In the non-shopping model the at-risk indicator is put to zero on the days a household never shops. These non-shopping days are derived from the purchase data on 10 product groups as described in section 2.

5.1 Estimation results

I use the following explanatory variables: household income, household size and the volume purchased at the previous purchase occasion (divided by 32 oz.). The observed volume is decomposed into two variables, that is, the (time-constant) average volume purchased per household and the (time-varying) deviation of this average at each purchase for each household. The first volume indicates whether a household purchases in large or small amounts (a “regular” trip), while the latter volume indicates whether the household makes a “fill-in” trip. Bijwaard et al. (2003) only consider yogurt purchase and because yogurt is only storable for a short period of time they did not find a fill in effect. Detergent can be stored for a long time. I therefore expect that fill in trips will have a significant impact on the time till the next detergent purchase.

The actual reference price is decomposed in a similar way into a (time-constant) household average reference price and the deviation from this average for the actual reference price. The latter is time-varying on a weekly basis and indicates a price cut (or price increase), while the first indicates the preference for a low priced or expensive product. In addition to a price cut two other marketing-mix variables, display and feature, indicate whether brands are on display or are featured in a newspaper. To account for the temporal effect of the marketing instruments, I also add the differences between the current value and the value at the previous purchase for all three marketing-mix variables.

Table 2 presents the estimation results for the five alternative specifications for the yogurt purchase data and Table 3 for the detergent purchase data. The choice among the different alternative specification for the both data lead to the same conclusion that the fit of the model that accounts for regular and non-shopping days is the best. Formal LR-test and the lower values of the AIC for both data indicate that the calendar-time models improve the fit significantly. All the calendar-time models get substantially different results for the effect of many marketing-mix

variables compared to the results of the gap-time models.

Yogurt purchases

From the parameters associated with household characteristics I draw the conclusion that for all models both the average volume purchased by a household and the deviation from this average has a significant impact on predicting the yogurt repurchase behavior. Households that purchase yogurt in larger quantities purchase it more often. Purchasing an increased quantity compared to the average volume leads to a longer time till the next yogurt purchase. Neither household income nor household size has a significant impact on the timing of yogurt purchases. Households with a preference for expensive yogurt purchase it more frequently than households with a preference for less expensive yogurt.

The marketing mix variables are also important for explaining the timing of yogurt purchases. The signs of these three included variables are as one would expect. A price cut, a display or feature of yogurt all lead to an increase in the purchase probability. The effect of these marketing variables differ between the gap-time and the calendar-time models. In the gap-time models the effect of a price cut is larger, while in the calendar-time models the effect of a feature or a display is larger. We would expect that a promotion at the previous purchase increase the time till the next purchase. These intertemporal effects of all three marketing mix variables have the expected sign (except for the price effect in gap-time models) but only the temporal effect of featuring yogurt is significant. For the calendar time models the temporal effect of featuring yogurt is larger than the direct effect. This would imply a huge post-promotion dip in purchases. In the simulation of marketing actions in section 5.3 the long-run effects of featuring yogurt is indeed negative.

– put table 2 about here–

Detergent purchases

Household income and size are important determinants for explaining the purchase timing of detergent. The larger the households and higher the household income the more prone a household is to purchase detergent. Detergent can be stored easily. It is therefore not a surprise that for

detergent purchases the volume bought and the interpurchase time are inversely related. Households purchasing in larger quantities purchase less often and when a household was induced to purchase a larger than average volume the next purchase time will be postponed. Price is also negatively related to the interpurchase time of detergent purchases. Households that prefer expensive detergent purchase it less often than household that prefer less expensive detergent.

All the parameters associated with the marketing mix variables indicate that detergent purchases are very sensitive to a promotion. Both putting detergent on display and featuring in a newspaper increases the purchase probability 5 to 7 times. The price 1 parameter, that denotes the effect of a price cut, is a negative sign. That is, if a price cut of detergent occurs, a household's probability of purchasing detergent increases. If the price at the previous purchase is lower than the current price, we get as expected a positive effect on the repeat purchase probability. The temporal effects of displaying detergent on the purchase timing is large and significant. The temporal effect of featuring detergent has the wrong sign but is insignificant.

– put table 3 about here–

Gap-time duration dependence

All the estimated parameters of the gap-time duration dependence are significant. From these parameters I derive the implied gap-time duration dependence. This duration dependence is common for all households. They are shown in Figure 5 for the yogurt data and in Figure 6 for the detergent data.

– put figure 5 about here–

The shape of the implied gap-time duration dependence in yogurt purchases is similar for the different models. The probability to purchase yogurt again increases fast (positive duration dependence) right after the last purchase to reach a maximum (apart from the spikes) at 2 weeks since the previous purchase. Then it shows negative duration dependence (for 40 weeks). The spikes at regular intervals of a week implied by the relative large α_4 are in close connection with the frequency distribution in Figure 1. Neglecting the regular shopping days (models gap 0 and calendar 0) lead to higher probability to purchase on all days which are not multiples of a week.

– put figure 6 about here–

For detergent purchases the gap-time duration dependence has a different shape. It shows positive duration dependence (apart from the spikes) in the first 12 weeks after the previous detergent purchase. Then it decreases slowly for the next 22 weeks to increase again. The effect of regular shopping is also apparent for detergent purchases from the spikes in the figures. The effect of regular shopping is also for the detergent data large and significant and seems to capture the pattern in Figure 2 well.

The estimated variance of the unobserved heterogeneity, σ^2 , is around 0.30 ($= \exp(-1.18)$) for the yogurt data and around 0.48 ($= \exp(-0.73)$) for the detergent data and it differs significantly from zero, which indicates that a model without unobserved heterogeneity is not correctly specified.

Calendar-time effects

As indicated by figure 3 and figure 4 calendar time effect play a role in the purchase timing behavior of households for these two products. The estimated calendar-time effects are shown in Table 4. Both yogurt and detergent are purchased more often on Saturdays, than on other days of the week. Friday is the second important shopping day. The month effects are all relative to January. These seasonal effects differ between the two products. Yogurt is purchased more frequently in February and March and less frequently in May, August, November and, December, while Detergent is purchased less frequently in March, April, May, July and, November.

– put table 4 about here–

5.2 Predictive performance

In the previous subsection I have discussed the impact of household characteristics and marketing mix variables on the purchase timing behavior. I showed that the model that accounts for both regular and non-shopping days has the best fit. For store managers the predictive performance of the models is even more important, because that shows the value of the model in real setting. In most marketing research the predictive performance is calculated by testing whether a model predicts well for new households by splitting the data in two by household. For a part of the

households, the in-sample population, the whole purchase history is used to (re-)estimate the model while for the remaining households, the out-of sample population, the purchase history is predicted using the estimated parameters. Then the predictive power the model for these new households is calculated.

To my opinion this is not the most relevant predictive performance of a model for purchase timing behavior. A store manager wants to know how much his regular costumers will shop in the future. For frequently bought consumer goods new costumers only arise if people move residence or supermarket. I think this is a minor issue compared to the shopping volume of the households already present. I therefore asses the performance of the models in predicting future purchases. As already indicated, the calendar-time models allow us to forecast purchases beyond the next purchase. For this purpose the data for both the yogurt purchases and the detergent purchases is split in two. The first 78 weeks of observation are used for estimating the model parameters. Then I use the estimation results of the five models based on these first 78 weeks of observation to predict the purchase behavior of the households in the remaining weeks (13 for the yogurt data and 19 for the detergent data). In the yogurt data 3 households were not observed before week 78 and are therefore not included in the analysis. This leaves me with 595 households purchasing 6874 time yogurt and 624 households purchasing 6202 times detergent to base the estimation of the models on.²

The inherent non-linear feature of the models requires simulation based forecasting. Forecasting the time till the first purchase is easy for all models. However, as mentioned by Bijwaard et al. (2003), forecasting the number of purchases made by all households within a specific period, 13 and 19 weeks in our validation, is very difficult with a gap-time model because we need to forecast two or more purchases ahead. This implies resetting all time-varying covariates after each purchase. For each household this may occur at a different duration and, therefore, adjusting the input data for the simulation becomes very complicated.

In a calendar-time set-up we can simulate per day for all households simultaneously whether a purchase is made. Thus, we simulate the counting process of purchasing. The time till the first

²The estimation results do not differ much from the results in Table 2, Table 3 and, Table 4 and can be obtained from the author.

purchase is the number of days till the simulated household counting process makes a jump. The data is adjusted in a natural way with the additional simulated forecast after each day we simulate. Then, the total number of forecasted purchases and the percentage of households purchasing more than once within a given period can be directly derived from the simulated counting processes. Hence, if we have n households and the horizon is w weeks, each simulation round provides us with a matrix of size $n \times (7w)$ of zeros and ones, where a 1 corresponds with a purchase of a household on that particular day.

The household specific unobserved heterogeneity is, as is in the name, not observed but come from a random (gamma) distribution. The intensity is therefore unknown. Two solutions to this problem exist. The first solution is to estimate the implied contributions from the conditional distribution of v_i given the estimated parameters. This is related to the EM-algorithm applied by Bijwaard et al. (2003). They are forced to apply this method, because they (also) consider a semi-parametric model in which the (calendar-time) duration dependence is left unspecified. However, the semi-parametric model is not very useful in predicting beyond the estimation sample, because then the duration dependence is not available out of sample. The second solution of integrating out the unobserved heterogeneity is very easy with a gamma distributed unobserved heterogeneity. The resulting intensity used in the simulations is

$$\lambda_i^v(t) = \frac{1 + \sigma^2 N_i(t^-)}{1 + \sigma^2 \Lambda_i(t^-)} Y_i(t) \lambda_i(t|x_i(t)), \quad (11)$$

where t^- is the time just before t . Thus, if household i makes it's 10th purchase at t then $N_i(t) = 10$ and $N_i(t^-) = 9$. $\lambda_i(t|x_i(t))$ is the intensity given in (8) for $v_i \equiv 1$ and $\Lambda_i(t)$ is the related cumulative intensity.

The simulation process averages 500 of the simulation rounds. Then, for each day I calculate the average (simulated) number of purchases across households and the percentage of households making a first or second purchase. For each of the five models considered I compare the simulated number of purchases, percentage of first repeat purchase and percentage of second repeat purchase both on a daily and on an, aggregated, weekly basis. The percentage of the first repeat purchase is of interest if we like to know how many households make at least one purchase within the validation period. The percentage of the second repeat purchase is of interest if we like to know how many households make at least two purchases within the validation period. I

present the results for a MSE-validation test in Table 5. The value of the MSE-test is equal to $\sum_t (P_t^s - P_t^o)^2$, where P_t^s and P_t^o are the simulated and observed values (# of purchases or % of first and second repeat purchases) at time t (day or week), respectively. A model performs best if it has the lowest value for that test. Therefore, we compute the ratio to the lowest value of the test among the models. These ratios are shown in Table 5.

– put table 5 about here–

For both the yogurt and the detergent purchase timing data the validation results indicate that the calendar-time models outperform the gap-time models. Accounting for regular shopping days in the gap-time model increases the predictive performance. The comparison among the calendar-time models is somewhat ambiguous. The model that accounts for regular and non-shopping days outperforms the other calendar-time models half of the time. The difference in MSE among the models is about 1% to 9%. For the yogurt purchases the number of purchases and the percentage of households purchasing at least twice yogurt in the 13 weeks of validation are better predicted when regular and non-shopping days are accounted for. While for the detergent purchases only the percentage of households purchasing at least twice detergent in the 19 weeks of validation are better predicted with such a model. These ambiguous results for the model are probably related to the relative small number of households in my sample for which I could identify the non-shopping days. With more information on shopping-days and holidays the model would improve and provide a better tool for prediction the purchase timing behavior of households. Note that the model with non-shopping days always predicted less purchases than the models without. Thus, if the calendar-time model already underestimates the number of purchases accounting for non-shopping days would decrease the prediction fit.

5.3 Simulating marketing policies

Accounting for non-shopping days in the model not only leads to a change in the estimated parameters but also in a change of days that the households are at-risk of purchasing in the future. With only a limited non-shopping days identified the effect of this may be small in the short run, but with the accumulation of non-shopping days it is non negligible in the long run. To illustrate the long run implications of accounting for regular and non-shopping days I examine

in this section the short- and long-run effects on purchase timing behavior of three different promotion scenarios. I analyze the effect of a promotion in a single week, which I choose to be week 78.

To assess the dynamic impact of the promotions I rely on simulation. I use the calendar-time models to simulate purchases for the next 13 weeks (yogurt) or 19 weeks (detergent), starting in week 78. All explanatory variables are set at their average value, except for the time-constant household-specific variables. In the first scenario, I introduce a display promotion without price-cut and/or feature support in week 78. In week 79 and beyond I set all the marketing-mix variables at their non-promotional values. This implies that after week 78 no display or feature takes place and that the price is equal to the average household-specific price. Note that if a household purchases yogurt or detergent in week 78 the difference between the value of the marketing instrument after week 78 and the value at the previous purchase differs from zero until the next purchase. In the second and third scenario I analyze the effect of a feature advertisement and a price cut of 33% of the household-specific referencer price in week 78, respectively.

Again I simulate 500 times the counting process of purchasing. Then, for each day the average (simulated) number of purchases across households is calculated. All scenarios are compared to a baseline scenario in which the marketing-mix variables are put at their non-promotional value (zero) in week 78 and beyond. Figure 7 and Figure 9 show the cumulative effect on the number of purchases in the weeks after the promotion for the non-shopping model. In the promotion week the purchases are increased substantially. However, in the weeks after the promotion the purchases are lower. Hence, we observe the well-known post-promotion dip in purchases.

For both yogurt and detergent purchases putting the product on display has a strong and long lasting effect on the number of purchases. The detergent purchases more than double in the promotion week and the yogurt purchases increase with 25%. The effect of a display in the long run is negative for detergent. However, the effect is positive for the first 20 weeks after the promotion. Putting detergent on display again within this period would remove this negative effect. For a price-cut in yogurt and for featuring yogurt, the post-promotion dip lead to negative long-run effects of these promotions. For the long-run effect of featuring yogurt this negative long run effect is implied directly by the size of the estimated parameters of the temporal and direct

effect the feature percentage on yogurt purchases in table 2. This would suggest that the store manager should promote yogurt again within 4 weeks after a price cut to get positive effects of the promotion. Combining the price-cut with a in-store display will extend the time till the next necessary promotion. The effect of featuring the product or a price cut has a smaller initial effect on the purchases.

– put figure 7 about here–

– put figure 9 about here–

Accounting for non-shopping days leads to less predicted purchases in both the baseline and in the promotion scenarios. My interest lays in the relative effect of the promotions on the number of purchases. I therefore calculate the difference between the promotion impact of the three calendar-time models, where the promotion impact is the average number of additional purchases since week 78 due to the promotion activity in week 78 per household. The results for the calendar-time model without regular nor non-shopping days are taken as the reference results. The difference of the cumulative effect of the promotions in the other two calendar-time models compared to the reference results over time is shown in Figure 8 and Figure 10.

– put figure 8 about here–

– put figure 10 about here–

Both figures clearly depict that the impact of neglecting regular and non-shopping days leads to overestimation of the effect of promotions on the purchases (up to 8%). Only account for regular shopping days and not for non-shopping days leads to an overestimation of the promotion effects (about 2%). These two consequences of the neglecting regular and non-shopping behavior increase with the time since the promotion week. This is most clearly for the yogurt purchases. For the yogurt purchases the consequence of neglecting regular and non-shopping days is the largest for a price cut in yogurt and featuring yogurt. For the detergent purchase timing the overestimation of the effect of promotion is the largest for display.

6 Conclusion

One of the most important questions related to purchase timing behavior of consumers is whether promotions are effective in changing the consumers' purchase timing. Chiang et al. (2001) postulate that marketing mix variables have little impact on this timing when adjusting for regular shopping days. I reject their conclusion if I analyze yogurt and detergent purchase timing behavior from households in the ERIM data with a purchase timing model that accounts for regular and non-shopping days. I show in my simulations, however, that the long run effects of some of the marketing promotion may be negative. Therefore, for a store manager it is important to get the timing of the promotions right.

Bijwaard, Franses, and Paap (2003) introduced the calendar-time model for analyzing the purchase timing behavior of consumer good of households. They used a counting process formulation of the model and showed that their approach has many advantages over the conventional gap-time models of purchase timing. The contribution of this paper is that I extend their model to include regular and non-shopping days. Regular shopping at weekly intervals are captured in the model by a weekly indicator in the gap-time duration dependence. The non-shopping days are captured in the model by setting the at-risk indicator to zero for the days a household never shops. This effectively sets the probability to purchase on those days to zero.

In an empirical application on yogurt and detergent purchase timing, I find that the model that accounts for both regular and non-shopping days yields the best fit. I also provide a validation comparison on predicting future purchases and a marketing action simulation. The validation results are rather mixed. The calendar-time models all outperform the gap-time models, but accounting for regular and non-shopping days do not always lead to improvement. This is probably related to the small number of households for which I could identify their non-shopping days from the data. In many countries in Europe the shops are closed on Sundays. Therefore, it would be interesting to estimate the model that accounts for non-shopping days on European purchase timing behavior.

Accounting for non-shopping days in the model not only leads to a change in the estimated parameters but also in a change of days that the households are at-risk of purchasing in the future. With only a limited non-shopping days identified the effect of this may be small in

the short run, but with the accumulation of non-shopping days it is non negligible in the long run. This is shown in a simulation of promotional actions. Not accounting for regular and non-shopping days would lead to overestimate the effect of promotions on the purchases. This bias is increasing over time.

Table 1: Percentage of non- and rarely ($< 5\%$) shopping days for yogurt and detergent purchase data

	Non-shopping days						
	Distribution over the week						
	Su	Mo	Tu	We	Th	Fr	Sa
yogurt	3.0	0.5	0.7	0.5	0.7	0.2	0.3
detergent	4.0	1.1	1.4	0.8	0.5	1.1	0.8
	Frequency distribution						
	0	1	2	3	4	5	6
yogurt	95.8	3.2	0.5	0.3	0.2	0.0	0.0
detergent	93.6	4.5	1.3	0.2	0.3	0.0	0.2
	Rarely shopping days						
	Distribution over the week						
	Su	Mo	Tu	We	Th	Fr	Sa
yogurt	23.0	12.7	15.5	11.7	9.2	5.7	3.5
detergent	28.7	12.5	15.7	11.1	7.1	6.2	6.1
	Frequency distribution						
	0	1	2	3	4	5	6
yogurt	53.7	27.9	8.7	4.8	3.2	1.2	0.5
detergent	53.5	27.7	6.9	5.8	3.0	1.9	1.1

Table 2: Parameter estimates for yogurt purchase data
(t-values in parentheses)

	Gap0	Gap 1	Calendar 0	Calendar 1	Non-shopping
income	-0.0138 (-1.40)	-0.0144 (-1.47)	-0.0145 (-1.46)	-0.0152 (-1.54)	-0.0160 (-1.61)
Household size	0.0221 (1.12)	0.0242 (1.24)	0.0225 (1.14)	0.0243 (1.23)	0.0244 (1.24)
volume 1 ^b	-0.0798 (-2.37)	-0.0832 (-2.47)	-0.0769 (-2.28)	-0.0799 (-2.37)	-0.0810 (-2.40)
volume 2 ^b	0.3929 (3.80)	0.3842 (3.74)	0.3944 (3.79)	0.3892 (3.76)	0.3888 (3.75)
price 1 ^b	-0.5285 (-4.92)	-0.5304 (-4.93)	-0.4526 (-4.10)	-0.4530 (-4.10)	-0.4508 (-4.08)
price 2 ^b	0.6097 (8.89)	0.6087 (8.93)	0.6041 (8.74)	0.6044 (8.79)	0.5989 (8.71)
display	0.8099 (3.87)	0.8089 (3.86)	0.9375 (4.42)	0.9302 (4.38)	0.9319 (4.39)
feature	0.4000 (2.12)	0.4095 (2.16)	0.2447 (1.26)	0.2578 (1.32)	0.2636 (1.35)
price difference ^c	0.0292 (0.28)	0.0355 (0.34)	-0.0370 (-0.35)	-0.0340 (-0.32)	-0.0356 (-0.34)
display difference ^c	0.2356 (1.26)	0.2239 (1.20)	0.1988 (1.06)	0.1882 (1.01)	0.1866 (1.00)
feature difference ^c	0.3154 (1.87)	0.3031 (1.81)	0.3116 (1.84)	0.2998 (1.78)	0.2967 (1.76)
α_1	-0.1550 (-25.73)	-0.1399 (-23.09)	-0.1550 (-25.60)	-0.1401 (-23.01)	-0.1402 (-23.02)
α_2	0.0018 (19.72)	0.0017 (17.64)	0.0019 (19.57)	0.0017 (17.53)	0.0017 (17.56)
α_3	0.2722 (14.61)	0.1946 (10.22)	0.2786 (14.90)	0.2034 (10.63)	0.2037 (10.66)
α_4	. (.)	0.8338 (31.05)	. (.)	0.8079 (29.94)	0.7989 (29.60)
constant	-1.3584 (-43.31)	-1.5061 (-47.25)	. (.)	. (.)	. (.)
$\ln(\sigma^2)$	-1.1854 (-16.37)	-1.2016 (-16.53)	-1.1685 (-16.18)	-1.1837 (-16.33)	-1.1820 (-16.32)
log-likelihood	-18733.3	-18314.8	-18600.8	-18208.8	-18161.3
AIC	37498.6	36663.6	37267.6	36485.6	36390.6

^a Estimation based on 598 households and 7942 purchases. ^b Volume 1 is the volume bought at the previous purchase compared to household average and Volume 2 is the household average volume bought. Price 1 is the current price compared to household average and Price 2 is the household average price. ^c Price/feature/display difference are the differences of the current price/feature/display compared to the value at the previous purchase.

Table 3: Parameter estimates for detergent purchase data
(t-values in parentheses)

	Gap 0	Gap 1	Calendar 0	Calendar 1	Non-shopping
income	0.2617 (9.93)	0.2598 (9.90)	0.2605 (9.92)	0.2572 (9.86)	0.2519 (9.69)
Household size	0.0250 (2.06)	0.0249 (2.06)	0.0249 (2.06)	0.0249 (2.08)	0.0235 (1.96)
volume 1 ^b	-0.1492 (-15.10)	-0.1486 (-15.05)	-0.1481 (-14.95)	-0.1475 (-14.89)	-0.1475 (-14.89)
volume 2 ^b	-0.1580 (-5.53)	-0.1558 (-5.47)	-0.1566 (-5.50)	-0.1547 (-5.45)	-0.1572 (-5.55)
price 1 ^b	-0.9198 (-4.57)	-0.9155 (-4.54)	-0.7672 (-3.72)	-0.7589 (-3.68)	-0.7630 (-3.69)
price 2 ^b	-0.7635 (-4.02)	-0.7596 (-4.02)	-0.7556 (-3.99)	-0.7546 (-4.02)	-0.7378 (-3.93)
display	1.9276 (3.43)	1.9050 (3.39)	1.8949 (3.32)	1.8914 (3.32)	1.8796 (3.30)
feature	1.6303 (2.09)	1.6474 (2.11)	1.9553 (2.47)	1.9444 (2.45)	1.9647 (2.47)
price difference ^c	-0.3108 (-1.70)	-0.3056 (-1.67)	-0.3314 (-1.80)	-0.3208 (-1.74)	-0.3154 (-1.71)
display difference ^c	0.8744 (1.96)	0.8650 (1.94)	0.7134 (1.60)	0.7098 (1.59)	0.7201 (1.61)
feature difference ^c	-0.0183 (-0.03)	-0.0389 (-0.07)	-0.0821 (-0.14)	-0.0805 (-0.14)	-0.0992 (-0.17)
α_1	-0.1193 (-15.14)	-0.1138 (-14.31)	-0.1180 (-14.96)	-0.1134 (-14.23)	-0.1133 (-14.23)
α_2	0.0013 (9.75)	0.0012 (9.20)	0.0013 (9.66)	0.0012 (9.18)	0.0012 (9.17)
α_3	1.0379 (30.64)	1.0056 (29.21)	1.0285 (30.31)	1.0004 (29.02)	1.0006 (29.04)
α_4	.	0.8635 (33.26)	.	0.8154 (31.20)	0.8021 (30.67)
constant	-2.8892 (-71.66)	-3.0510 (-74.62)	.	.	.
$\ln(\sigma^2)$	-0.7144 (-10.78)	-0.7253 (-10.92)	-0.7225 (-10.88)	-0.7347 (-11.03)	-0.7400 (-11.10)
log-likelihood	-20248.5	-19767.9	-19997.5	-19569.85	-19506.6
AIC	40529.0	39569.8	40061.0	39207.7	39081.2

^a Estimation based on 624 households and 7290 purchases. ^b Volume 1 is the volume bought at the previous purchase compared to household average and Volume 2 is the household average volume bought. Price 1 is the current price compared to household average and Price 2 is the household average price. ^c Price/feature/display difference are the differences of the current price/feature/display compared to the value at the previous purchase.

Table 4: Estimated calendar-time effects
(t-values in parentheses)

	yogurt			detergent		
	Calendar 0	Calendar 1	non-shopping	Calendar 0	Calendar 1	non-shopping
Monday	-1.4184 (-25.01)	-1.5436 (-27.11)	-1.5379 (-27.00)	-2.9323 (-47.46)	-3.0666 (-49.38)	-3.0560 (-49.22)
Tuesday	-1.4467 (-25.39)	-1.5604 (-27.29)	-1.5534 (-27.16)	-3.0269 (-48.29)	-3.1443 (-49.97)	-3.1281 (-49.72)
Wednesday	-1.4507 (-25.47)	-1.5708 (-27.50)	-1.5629 (-27.35)	-2.9678 (-47.83)	-3.0989 (-49.76)	-3.0877 (-49.58)
Thursday	-1.3042 (-23.40)	-1.4651 (-26.08)	-1.4579 (-25.94)	-2.8533 (-46.87)	-3.0135 (-49.16)	-3.0050 (-49.03)
Friday	-1.2638 (-22.94)	-1.4285 (-25.71)	-1.4253 (-25.65)	-2.6396 (-44.80)	-2.8183 (-47.40)	-2.8042 (-47.17)
Saturday	-1.0573 (-19.71)	-1.2402 (-22.89)	-1.2327 (-22.74)	-2.3418 (-40.98)	-2.5420 (-43.95)	-2.5324 (-43.80)
Sunday	-1.4631 (-25.73)	-1.5924 (-27.88)	-1.5644 (-27.37)	-3.0246 (-48.39)	-3.1445 (-50.11)	-3.1106 (-49.54)
February	0.1627 (3.16)	0.1582 (3.07)	0.1582 (3.07)	-0.0763 (-1.41)	-0.0791 (-1.46)	-0.0792 (-1.46)
March	0.1316 (2.63)	0.1304 (2.61)	0.1306 (2.61)	-0.1105 (-2.06)	-0.1114 (-2.08)	-0.1113 (-2.07)
April	0.0214 (0.42)	0.0206 (0.40)	0.0206 (0.40)	-0.2140 (-3.90)	-0.2136 (-3.89)	-0.2136 (-3.89)
May	-0.1342 (-2.52)	-0.1324 (-2.49)	-0.1323 (-2.48)	-0.1975 (-3.65)	-0.2017 (-3.73)	-0.2018 (-3.73)
June	-0.0046 (-0.09)	-0.0045 (-0.09)	-0.0043 (-0.08)	-0.0518 (-0.98)	-0.0536 (-1.01)	-0.0535 (-1.01)
July	-0.0450 (-0.86)	-0.0427 (-0.82)	-0.0427 (-0.82)	-0.1183 (-2.20)	-0.1174 (-2.18)	-0.1175 (-2.18)
August	-0.2014 (-3.53)	-0.1973 (-3.46)	-0.1971 (-3.45)	0.0408 (0.66)	0.0417 (0.68)	0.0412 (0.67)
September	0.0417 (0.66)	0.0402 (0.63)	0.0402 (0.63)	0.0045 (0.07)	0.0037 (0.06)	0.0038 (0.06)
October	0.0101 (0.16)	0.0110 (0.18)	0.0110 (0.18)	-0.1068 (-1.84)	-0.1054 (-1.81)	-0.1053 (-1.81)
November	-0.1738 (-2.60)	-0.1750 (-2.62)	-0.1747 (-2.61)	-0.1366 (-2.44)	-0.1368 (-2.44)	-0.1369 (-2.44)
December	-0.1972 (-3.23)	-0.1931 (-3.16)	-0.1928 (-3.15)	-0.0397 (-0.75)	-0.0400 (-0.75)	-0.0402 (-0.75)

Table 5: Validation results (MSE ratio to best)

	<i>yogurt</i>				
	Gap 0	Gap 1	Calendar 0	Calendar 1	non-shopping
<i>on daily basis</i>					
# of purchases	1.36	1.31	1.03	1.01	1
% 1st purchase	1.26	1.18	1.01	1	1.01
% 2nd purchase	1.05	1.03	1.02	1.00	1
<i>on weekly basis</i>					
# of purchases	1.42	1.40	1.09	1.06	1
% 1st purchase	1.17	1.20	1	1.03	1.04
% 2nd purchase	1.04	1.01	1.04	1.01	1
	<i>detergent</i>				
	Gap 0	Gap 1	Calendar 0	Calendar 1	non-shopping
<i>on daily basis</i>					
# of purchases	1.58	1.51	1	1.00	1.01
% 1st purchase	1.39	1.29	1	1.00	1.00
% 2nd purchase	1.19	1.17	1.01	1.01	1
<i>on weekly basis</i>					
# of purchases	1.07	1.07	1.00	1	1.03
% 1st purchase	1.24	1.23	1	1.02	1.02
% 2nd purchase	1.13	1.14	1.01	1.00	1




Figure 1: Frequency of time since previous yogurt purchase (in first 13 weeks)




Figure 2: Frequency of time since previous detergent purchase (in first 13 weeks)

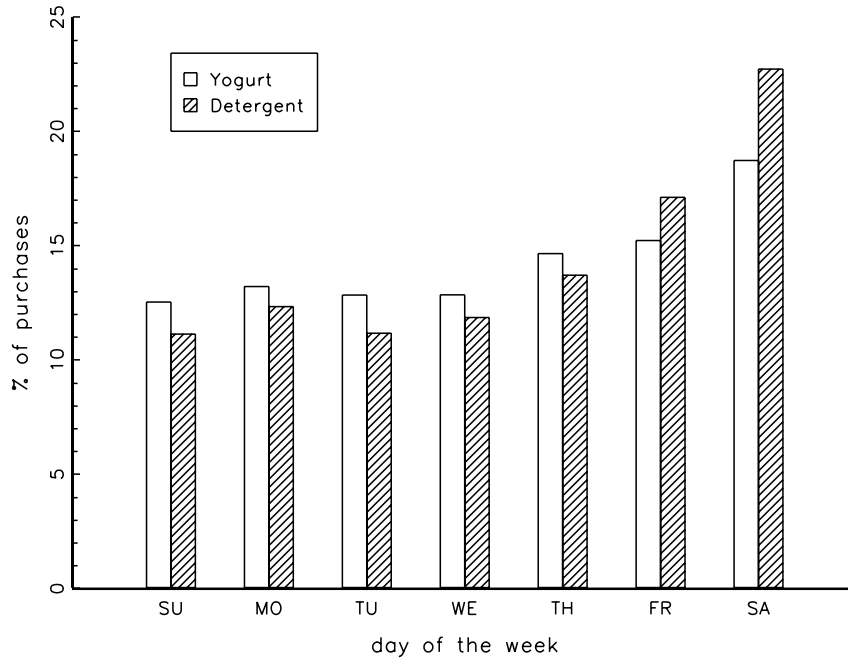


Figure 3: Distribution of purchases over the week




Figure 4: Distribution of purchases over the year (adjusted for observation)

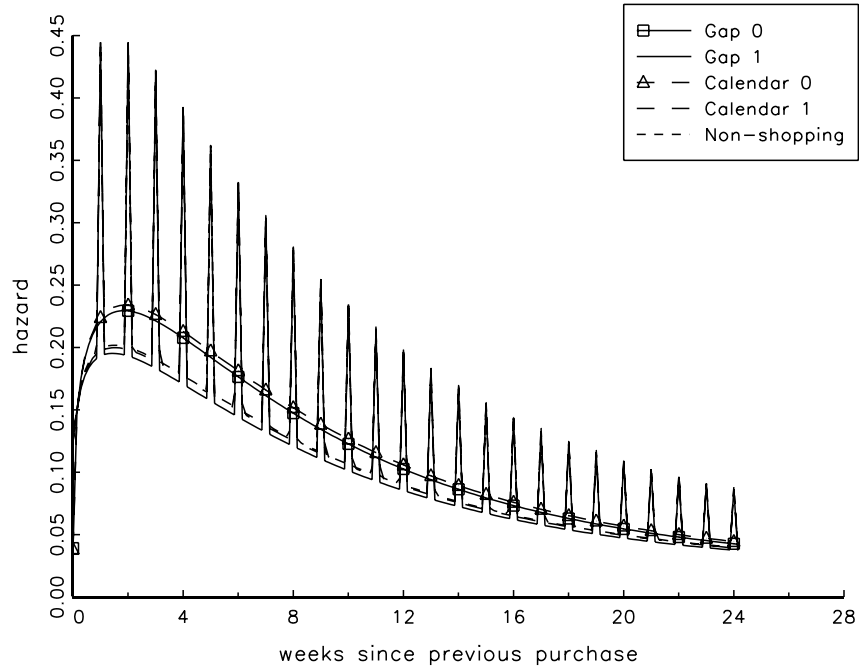


Figure 5: Estimated Gap-time duration dependence (yogurt)




Figure 6: Estimated Gap-time duration dependence (detergent)




Figure 7: Cumulative effect of marketing promotion in week 78 on number of yogurt purchases with model accounting for regular and non-shopping days




Figure 8: Effect of accounting for regular (1) and non-shopping days (ns) on the effect of 1 week yogurt promotion in week 78 on the cumulative yogurt purchases




Figure 9: Cumulative effect of marketing promotion in week 78 on number of detergent purchases with model accounting for regular and non-shopping days




Figure 10: Effect of accounting for regular (1) and non-shopping days (ns) on the effect of 1 week detergent promotion in week 78 on the cumulative laundry detergent purchases

References

- Andersen, P. K., O. Borgan, R. D. Gill, and N. Keiding (1993). *Statistical Models Based on Counting Processes*. New York: Springer–Verlag.
- Andersen, P. K. and R. D. Gill (1982). Cox’s regression model for counting processes: A large sample study. *Annals of Statistics* 10, 1100–1120.
- Bijwaard, G. E., P. H. Franses, and R. Paap (2003). Modeling purchases as repeated events. Working paper, Econometric Institute Erasmus University, EI 2003-45.
- Chiang, J., C. F. Chung, and E. T. Cremers (2001). Promotions and the pattern of grocery shopping time. *Journal of Applied Statistics* 28, 801–819.
- Cox, D. R. (1972). Regression models and life–tables (with discussion). *Journal of the Royal Statistical Society: Series B* 34, 187–220.
- Dunn, R., S. Reader, and N. Wrigley (1983). An investigation of the assumptions of the NBD model as applied to purchasing at individual stores. *Applied Statistics* 32, 249–259.
- Gupta, S. (1991). Stochastic models of interpurchase time with time dependent covariates. *Journal of Marketing Research* 28, 1–15.
- Helsen, K. and D. C. Schmittlein (1993). Analyzing duration times in marketing: Evidence for the effectiveness of hazard rate models. *Marketing Science* 12, 395–414.
- Jain, D. C. and N. J. Vilcassim (1991). Investigating household purchase timing decisions: A conditional hazard function approach. *Marketing Science* 10, 1–23.
- Kahn, B. and D. C. Schmittlein (1989). Shopping trip behavior: An empirical investigation. *Marketing Letters* 1, 55–69.
- Kalbfleisch, J. D. and R. L. Prentice (2002). *The Statistical Analysis of Failure Time Data, second edition*. John Wiley and Sons.
- Klein, J. P. and M. L. Moeschberger (1997). *Survival Analysis: Techniques for Censored and Truncated Data*. New York: Springer–Verlag.
- Lancaster, T. (1990). *The Econometric Analysis of Transition Data*. Cambridge University Press.

- Lin, D. Y. and L. J. Wei (1989). The robust inference for the Cox proportional hazards model. *Journal of the American Statistical Association* 84, 1074–1078.
- Seetharaman, P. B. and P. K. Chintagunta (2003). The proportional hazard model for purchase timing: A comparison of alternative specifications. *Journal of Business & Economic Statistics* 21, 368–382.
- Therneau, T. and P. Grambsch (2000). *Modeling Survival Data: Extending the Cox Model*. Springer-Verlag.
- Vilcassim, N. J. and D. C. Jain (1991). Modeling purchase-timing and brand-switching behavior incorporating explanatory variables and unobserved heterogeneity. *Journal of Marketing Research* 28, 29–41.
- Wheat, R. D. and D. G. Morrison (1990). Assessing purchase timing models: Whether or not is preferable to when. *Marketing Science* 9, 162–170.