University of Groningen

The estimation of pre- and postpromotion dips with store-level scanner data

Heerde, Harald J. van; Leeftang, Peter S.H.; Wittink, Dick R.

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:

1999

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Heerde, H. J. V., Leeftang, P. S. H., & Wittink, D. R. (1999). *The estimation of pre- and postpromotion dips with store-level scanner data*. s.n.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

The Estimation of Pre- and Postpromotion Dips with Store-Level Scanner Data

Harald J. van Heerde

Peter S.H. Leeflang

Dick R. Wittink *

SOM-theme B

Marketing and interactions between firms

Also available (and downloadable) in electronic version: <http://www.eco.rug.nl/SOM/>

* Harald J. van Heerde is Assistant Professor in Marketing and Marketing Research, Department of Marketing and Marketing Research, Faculty of Economics, University of Groningen, P.O. Box 800, 9700 AV Groningen, The Netherlands. Peter S.H. Leeflang is Professor of Marketing and Marketing Research at the same department. Dick. R. Wittink is the General George Rogers Clark Professor of Management and Marketing at the Yale School of Management, New Haven, CT, and Professor of Marketing and Marketing Research, at the University of Groningen. The first author can be contacted by telephone (+31 50 3637093), fax (+31 50 3637202), or by email (h.j.van.heerde@eco.rug.nl). The authors thank ACNielsen for providing the data and Sumas Wongsunopparat for useful suggestions.

Abstract

One of the mysteries of store-level scanner data modeling is the lack of a dip in sales in the week(s) following a promotion. Researchers expect to find a postpromotion dip because analyses of household scanner panel data indicate that consumers tend to accelerate their purchases in response to a promotion – that is, they buy earlier and/or purchase larger quantities than they would in the absence of a promotion. Thus, one should also find a pronounced dip in store-level sales in the week(s) following a promotion. However, researchers find such dips usually neither at the category nor at the brand level.

Several arguments have been proposed for the lack of a postpromotion dip in store-level sales data. These arguments explain why dips may be hidden. Given that dips are difficult to detect by traditional models (and by a visual inspection of the data), we propose models that can account for a multitude of factors which together cause complex pre- and postpromotion dips.

We use three alternative distributed lead- and lag structures: an Almon model, an Unrestricted dynamic effects model, and an Exponential decay model. In each model, we include four types of price discounts: without any support, with display-only support, with feature-only support, and with feature and display support. The models are calibrated on store-level scanner data for two product categories: tuna and toilet tissue. We estimate the dip to be between 4 and 25 percent of the current sales effect, which is consistent with household-level studies.

1. Introduction

One of the key issues in sales promotion research is whether “there is a trough after the deal” (Blattberg, Briesch, and Fox 1995). The evidence from analyses of *household-level panel data* is that consumers accelerate their purchases as a result of sales promotions. For example, Gupta (1988) decomposes the sales effect due to promotion for coffee into brand switching (84 percent), purchase timing acceleration (14 percent), and increased purchase quantity (2 percent). Chiang (1991) obtains similar percentages, while Grover and Srinivasan (1992, pp. 86-87) conclude that “one-fourth of the gain in a week’s product category sales resulting from a promotion is at the expense of the succeeding week’s sales”. Bell, Chiang, and Padmanabhan (2000) decompose the sales effect for thirteen product categories and find that, on average, brand choice accounts for 75 percent of the total elasticity (range 49-94 percent). Thus, the percent attributable to purchase timing acceleration and increases in purchase quantity varies between 6 and 51 percent.

At first sight one might think that the acceleration effects in timing and quantity evident *at the household level* should translate directly in a postpromotion dip in weekly store-level sales data. However, postpromotion dips are rarely detected in visual or (traditional) statistical analyses of *store data*. A resolution of this paradox is important for both researchers and managers. For researchers, a lack of convergent validity between the results from household-level panel data and weekly store-level sales data casts doubt on the nature of the acceleration phenomenon. Also, many managers say they use store-level (or more highly-aggregated) scanner data more frequently for analyses than household-level panel data (Bucklin and Gupta 2000). However, if the household results are accurate, then managers who rely on aggregated data for inferences about promotion effects will obtain incorrect conclusions. Unless stockpiling can be properly accounted for, managers will overestimate the effectiveness of promotions: they tend to classify all of the promotion-based sales spike as incremental (Neslin and Schneider Stone 1996). Since managers rely heavily on aggregated data, the most challenging part of the postpromotion dip paradox is the apparent lack of a dip in store-level data. We show two typical store-level sales graphs in Figure 1.1, from the data used in our empirical application (see below). Neither of these graphs shows any sales dip before or after a sales spike.

As far as we know, only Doyle and Saunders (1985), Leone (1987), Litvack, Calantone, and Warshaw (1985), and Moriarty (1985) studied acceleration based on store data. Doyle and Saunders (1985) demonstrated that *lead* effects of promotions, resulting from the anticipations of promotions by consumers and other economic agents, can be as important as *lagged* effects. They examined monthly gas appliance sales as a function of (a.o.) the commission structure for sales personnel, and analyzed

Figure 1.1: Typical store-level sales graphs

Tuna brand 1, store 28

Tissue brand E, store 21

whether salespeople move some customer purchases to the time period in which their commission rates are higher. Doyle and Saunders calculated that about 7 percent of total sales during an 8-week promotion period consisted of sales that would have taken place prior to the promotion period had commission rates not been increased during the promotion.

Leone (1987) applied intervention analysis to weekly sales data to evaluate a single “5 for \$1.00” sale for wet cat food. The weekly sales data graph showed a clear postpromotion dip and the analyses confirmed this dip. Litvack, Calantone, and

Warshaw (1985) observed the sales of many items before, during, and after a price cut. Interestingly, the authors did not observe a postpromotion dip in sales for the items that could have experienced purchase acceleration. Moriarty (1985) includes one-week lagged promotion variables in a sales response function. He finds significant postpromotion dips for only three of the fifteen cases.

Although a few of these studies obtain evidence of pre- or post-dips, Blattberg, Briesch, and Fox (1995, p. G127) mention that “examination of store-level POS data for frequently purchased goods rarely reveals a trough after a promotion. This anomaly is surprising and needs to be better understood.” Neslin and Schneider Stone (1996) consider eight possible arguments for the apparent lack of postpromotion dips in store-level sales data. Their arguments imply that the dips may be hidden. As a result, dips will be difficult to detect by traditional models or by a visual inspection of the data. Since brand sales are the aggregate of purchases across (heterogeneous) households, both pre- and postpromotion sales data will have complex patterns. Essentially, sales are shifted from multiple future- and past periods into a current, promotion-based sales spike in a nontrivial way.

Neslin and Schneider Stone (1996) suggest that researchers carry out “sophisticated distributed lag analyses of weekly sales data in the hope of measuring the postpromotion dip statistically.” We present a flexible modeling approach, and regress brand-level sales on current-, lead-, and lagged own-brand price indices with three different distributed lead and lag structures: an Almon model, an Unrestricted dynamic effects model, and an Exponential decay model. We distinguish four types of price discounts: ones without any support, ones with feature-only support, ones with display-only support, and price discounts with feature and display support.

The key contributions of our paper are:

- We propose a store-level model specification explicitly based on the *arguments* for the apparent lack of a postpromotion dip in aggregate data;
- We show there is *no postpromotion dip paradox*: we obtain pre- and postpromotion dips that are comparable to those obtained with household data;
- We propose a *new way of modeling* the interaction effects between price cuts and various types of support (feature and/or display), and show differences in the magnitude of pre- and postpromotion effects between price cuts with alternative types of support.

In section two of this paper we briefly review eight arguments for the apparent lack of a postpromotion dip in models of store sales, and derive implications from these arguments for the specification of pre- and postpromotion dips. In section three we

show the model specification and discuss model calibration. We introduce store-level scanner data sets for two product categories in section four, and provide empirical evidence for dynamic promotion effects for both categories in section five. In section six we present our conclusions.

2. Arguments for the Lack of a Postpromotion Dip, and Implications for Model Specification

Neslin and Schneider Stone (1996) provide eight possible arguments for the apparent absence of postpromotion dips in store-level scanner data. Five of their arguments complicate the identification of postpromotion dips in both household- and store data (a-e below):

- a. Consumers Purchase Deal to Deal;
- b. Increased Consumption;
- c. Competitive Promotions Mask the Dip;
- d. Positive Repeat Purchase Effects Cancel the Acceleration Effect;
- e. Retailers Partially Extend Promotions Beyond the First Week;

We refer to Neslin and Schneider Stone (1996) for a discussion of arguments a-d. We expand on argument e, because it has implications for our model specification. Retailers may extend all or part of a promotion –in particular display activity– and thereby increase sales immediately following the initial promotion which will mask the dip (Blattberg and Neslin 1990, p. 358). In principle, this effect can be accounted for with expanded display variables in the model. However, it is possible that extensions are not captured. For example, display activity is measured by weekly store audits, say on Thursday. If a display is extended only during the first three days of a second week (Monday, Tuesday and Wednesday), the value of the display variable will not accurately reflect this second week’s situation. Therefore, we need lagged display variables to capture this extension effect which should affect sales positively. The other common scanner data variables (sales, prices, feature activities) do not have this measurement error problem.

The final three arguments (f-h) provided by Neslin and Schneider Stone (1996) appear to complicate the analysis of store-level data uniquely.

- f. The Combined Effect of Quantity and Timing Acceleration

Time acceleration steals from the weeks immediately following the promotion, depressing sales in those weeks, while the effects of quantity stockpiling are manifested during the next consumer purchase occasion, depressing sales often after

a lag of a few weeks (Blattberg and Neslin 1990, p. 192). In household-level models these effects can be measured separately. Our store-level model should account for flexible, multi-period postpromotion dips to capture these effects jointly.

g. Lack of Consumer Inventory Sensitivity

Inventory may be reduced to its normal level only after an extended period. Hence, the postpromotion dip is dissipated into the future. In household-level models heterogeneity in inventory amounts and sensitivities can be accommodated, while a brand sales model should incorporate multi-period postpromotion dips.

h. Anticipatory effects

The literature suggests that consumers form price expectations (Winer 1986, Kalwani et al. 1990). If consumers expect a significant price reduction in the future they may defer or decelerate their purchases, causing prepromotion dips. In other words, an expected price decrease in period t may decrease sales in period $t - k$, $k = 1, 2, \dots$

We assume (compare Winer 1986 and Kalwani et al. 1990, equation (8)) that price expectations are unbiased, and that actual prices capture price expectations. Hence we propose that sales in period $t - k$ is a function of the actual price in period t . This implies that a retailer who offers a deal in week t will lose sales in week $t - k$ ($k \geq 1$) if customers expect a deal to be offered in week t .

Since consumers are heterogeneous in the length of the period they are willing to defer purchases in anticipation of a price promotion, prepromotion dips are spread out over multiple prepromotion weeks. However, we do not know when a prepromotion dip will be the deepest. Hence the model should also account for flexible, multi-period prepromotion dips. We note that household-level models that include future expected prices implicitly account for prepromotion effects.

To summarize, the store-level model we want to develop should account for factors that can hide postpromotion dips in traditional store-level models and in a visual representation of the data. The model should accommodate:

- i) multiple-week own-brand postpromotion effects, negative (arguments f, g), or positive (arguments d, e) relative to sales under the no-promotion scenario;
- ii) multiple-week own-brand prepromotion dips (argument h);
- iii) flexible pre- and postpromotion effects (arguments f, g, and h);
- iv) current cross-brand effects (argument c).

It is clear that there is very limited opportunity to test the relevance of specific arguments on sales data. For example, several arguments can account for the occurrence of negative postpromotion effects. Thus, we do not propose to test the

relevance of any arguments. Instead, our objective is to present and estimate a model which can accommodate complex pre- and postpromotion effects for managerial use, and to report the nature and magnitude of these effects in store data.¹

3. Model Specification and Calibration

We require a model that accommodates lead- and lagged effects for multiple sales promotion variables over multiple weeks in a flexible manner. Basically, there are two approaches: econometric and time series. The econometric approach includes lead- and lagged sales promotion variables as predictors in a model of brand sales. The time series approach would use transfer function/intervention modeling. In the latter case, we would use ARIMA models for all variables, and estimate a transfer function (if the predictors are continuous) or an intervention model (if the predictors are binary) to relate the criterion variable to the predictors.

The traditional time series approach, used for example by Leone (1987), would have to be modified in four ways. One, we have multiple predictors instead of one used by Leone (1987). Thus, we would have to model dynamic interactions between the predictors as well. Two, for a given predictor we have multiple promotional observations (see Table 4.1 below), whereas Leone evaluated the effect of a single promotion. Three, we need to allow for lead- and lagged effects. While Doyle and Saunders (1985) used time series methods to identify lead- and lagged effects for multiple predictors, their final model (see Doyle and Saunders, 1985, p. 59, equation (3)) is an econometric model. Four, we have time series observations for multiple stores. We are not aware of transfer function/intervention models for lead- and lagged

¹ There may be other, relevant differences between models of store sales and models of household purchases which pertain to the differential observability of postpromotion dips. For example, our model does not accommodate store switching. If the sales effect in a given store due to a promotion for a brand in that store partly reflects shifts in sales between stores, then the current effect that is potentially incremental to the manufacturer of the brand is overstated. Household models of brand choice typically focus on choices conditional upon a product category purchase in any of several stores. If the choice of store is not explicitly modeled, as a function of a.o. promotional activities, then the household models will be subject to similar problems (see Gupta et al. 1996 for a discussion of the conditions under which household choice models generate the same results as store sales models). Thus, we do not believe there is actually a difference between household- and store-level models on this aspect. Of course, households are unlikely to actively switch purchases between stores as a function of promotions for tuna or tissue, the product categories we study. Store switching does tend to occur for such items as soft drinks and disposable diapers (see Kumar and Leone 1988). Other relevant aspects are the following. One, household panel data could represent a nonrepresentative sample from store-level data. Two, household data involve the imputation of causal data for items not purchased whereas store-level data are complete. And three, the purchases for low-penetration categories may be unreliable with household data.

effects of multiple predictors in a pooled data context. Although the development of such models may be a fruitful area for new research, the econometric approach appears to be equally promising and is straightforward to implement.

Our econometric model is a modification of the Scan*Pro model (Wittink et al. 1988, Foekens, Leeflang, and Wittink 1994, Christen et al. 1997). The original model includes own- and cross-brand effects of promotions, and is estimated with store-level scanner data provided by ACNielsen. It was developed for commercial purposes, and the basic model has been used in over 1800 different commercial applications in North America, Europe and Asia. Our modification of the Scan*Pro model is related to: (1) variable choice, and (2) model specification.

Variable Choice

In our model the criterion variable is log of unit sales of a brand in a specific store in a given week, as is true for the Scan*Pro model. The Scan*Pro model includes as predictors four current promotional instruments: discount, feature-only, display-only, and feature and display, for the brand and for other brands. We modify this formulation and define: (1) own- and cross-brand discounts without support, (2) own- and cross-brand discounts with feature-only support, (3) own- and cross-brand discounts with display-only support, (4) own- and cross-brand discounts with feature and display support, (5) own-brand feature-only without price cuts, (6) own-brand display-only without price cuts, and (7) own-brand feature and display without price cuts.

We specify the instruments (1)-(4) as log price indices. We use logs in order to have the parameters be elasticities. We use price indices (ratio of actual to regular prices) to capture only the promotional price effects. For the instruments (1) we take the log price index observations and multiply these by one for observations with neither feature nor display activity for the brand, and by zero otherwise. The values of the instruments (2)-(4) are determined in an analogous way. We include instruments (5)-(7) as indicator variables since there is no price discount associated with those observations. Still, these activities can cause sales increases for the brand (Inman, McAlister, and Hoyer 1990).

Our approach has two advantages over the traditional approach of including log price index variables separately from the indicator variables for the non-price promotion variables. In our case, the set of own-brand variables is minimally correlated by definition (as is the set of cross-brand variables), whereas in the traditional approach, the log price index variable is often highly correlated with one or more of the non-price promotion variables. Also, the interpretation of our results is

straightforward: price cuts are the core of sales promotions, whereas feature and display are communication devices. Importantly, our variable definitions capture any interaction effects between price cuts and the various types of support, which is one of the key issues within sales promotion research (Blattberg, Briesch and Fox 1995). Our results show the own- and cross-brand price promotional elasticities for each of the four conditions of support.² Of course, our model also includes variables for dynamic price promotion effects. Specifically, we use variables to capture lead- and lagged own-brand log price index effects under the four conditions of support.

Model Specification

The current own- and cross-brand variables under promoted conditions are included multiplicatively, as in the Scan*Pro model. The lead- and lagged own-brand log price indices for the four different conditions are modeled with three alternative dynamic effects specifications:

- Unrestricted dynamic effects (Judge et al. 1985, pp. 351-356);
- Exponential decay dynamic effects (a finite duration version of the Geometric Lag Model in Judge et al. 1985, p. 388);
- Almon dynamic effects (Judge et al. 1985, pp. 356-364; Stewart 1991, pp. 181-186).

The Unrestricted dynamic effects approach approximates lead- and lagged effects by including the relevant predictors in lagged format $t - 1, t - 2, t - 3, \dots$, as well as in lead format $t + 1, t + 2, t + 3, \dots$. In this specification, all lead and lagged variables have unique parameters. Thus, if criterion variable y_t is explained by current, past and future values of one predictor variable x_t , up to a maximum lag of s periods and a maximum lead of s' periods, then the Unrestricted dynamic effects model is:

$$y_t = \alpha_0 + \alpha_1 x_t + \sum_{u=1}^s \beta_u x_{t-u} + \sum_{v=1}^{s'} \gamma_v x_{t+v} + u_t, \quad t = s + 1, \dots, T - s'. \quad (1)$$

The Exponential decay model imposes a structure on the dynamic effects (see also Blattberg and Wisniewski, 1989): $\beta_u = \lambda^{u-1} \beta$ and $\gamma_v = \mu^{v-1} \gamma$. As a result, the

² Our approach is somewhat similar to the one followed by Papatla and Krishnamurthi (1996, p. 23, equation (2)). They include interaction effects between indicator variables for price cut dummies, feature and display. However, their set of predictors is more correlated than our set, and their approach does not generate different price (promotion) elasticities for the promotion conditions.

model with one predictor x_t is:

$$y_t = \alpha'_0 + \alpha'_1 x_t + \sum_{u=1}^s \lambda^{u-1} \beta x_{t-u} + \sum_{v=1}^{s'} \mu^{v-1} \gamma x_{t+v} + u'_t, \quad t = s+1, \dots, T-s'. \quad (2)$$

The parameters λ and μ are the decay parameters.

The Almon model approximates the dynamic effects in (1) with polynomials. The lagged effect parameters are: $\beta_u = \sum_{m=0}^r \phi_m(u-1)^m$; $r < s$ ($u = 1, \dots, s$).³ The lead effect parameters are: $\gamma_v = \sum_{m=0}^{r'} \theta_m(v-1)^m$; $r' < s'$ ($v = 1, \dots, s'$). In this manner, the Almon model is:

$$y_t = \alpha''_0 + \alpha''_1 x_t + \sum_{u=1}^s \sum_{m=0}^r \phi_m(u-1)^m x_{t-u} + \sum_{v=1}^{s'} \sum_{m=0}^{r'} \theta_m(v-1)^m x_{t+v} + u''_t, \quad t = s+1, \dots, T-s'. \quad (3)$$

Our use of these three alternative dynamic effect specifications differs in three respects from the standard way they are used in the econometric literature. One, rather than only lagged effects, our approach includes lead effects as well. Two, instead of modeling the dynamic effects of just one variable, we model the dynamic effects of multiple variables. And three, the standard way is for researchers to use the Exponential decay- and the Almon models by imposing a structure in which the dynamic effect parameters are linked to the current effect parameter. Our approach relaxes this assumption: i.e., we let the current effect parameter be estimated independently of the lead- and lagged effect parameters. We do this because current price promotion effects are expected to be much larger than (week-specific) lead- and lag effects. In addition, we use separate approaches for lead- and lagged effects for all models since we should incorporate flexible dynamic effects (implication iii). For the Exponential decay model, this means that the lagged effect decay parameter λ may differ from the lead effect decay parameter μ . For the Almon model, this means that the degree of the lagged effect polynomial (r) may be different from the degree of the lead effect polynomial (r').

It is clear that the Unrestricted dynamic effects approach defined in equation (1) offers the highest degree of flexibility. However, it involves many lead- and lagged effect variables which may lead to multicollinearity. At the other end is the Exponential decay model (2) which uses few parameters, but is relatively inflexible. In model (2), the dynamic effect is assumed to be largest in the weeks immediately after (lagged

3 For $u = 1$ the first element of this sum is defined as $\phi_0(0)^0 \equiv \phi_0$.

effect) or before (lead effect) the promotion. This may be restrictive, because of the multitude of factors causing pre- and postpromotion effects (see section two). The Almon approach (3) is between these two approaches: it is more flexible than the Exponential decay model, but more parsimonious than the Unrestricted model.⁴

For brand k , $k = 1, \dots, J$, the full model is:

$$\begin{aligned} \ln S_{ik,t} = & \sum_{j=1}^J \sum_{l=1}^4 \alpha_{jkl} \ln(\text{PI}_{ijl,t}) + \alpha_{Fk} F_{ik,t} + \alpha_{Dk} D_{ik,t} + \alpha_{FDk} FD_{ik,t} \quad (4) \\ & + \sum_{u=1}^s \sum_{l=1}^4 \beta_{kl,u} \ln(\text{PI}_{ikl,t-u}) + \sum_{v=1}^{s'} \sum_{l=1}^4 \gamma_{kl,v} \ln(\text{PI}_{ikl,t+v}) \\ & + \psi_{ik} R_i + \xi_{ik} W_t + u_{ik,t}, \\ & t = s + 1, \dots, T - s' \text{ and } i = 1, \dots, N, \end{aligned}$$

with the lagged effect parameters β and lead effect parameters γ either kept unrestricted as in (1), modeled as an Exponential decay model (2), or modeled as an Almon model (3); and where:

$\ln S_{ik,t}$ is log unit sales of brand k in store i in week t ;

$\ln(\text{PI}_{ijl,t})$ is log price index (ratio of actual to regular price) of brand j in store i in week t ; $l=1$ denotes that the observation is not supported by feature nor display; $l=2$ that it is supported by feature-only; $l=3$: supported by display-only, and $l=4$: supported by feature and display;

$F_{ik,t}$ is a feature-only indicator variable for non-price promotion observations; $= 1$ if brand k is featured, but not displayed nor price promoted, by store i in week t , $= 0$ otherwise;

$D_{ik,t}$ is a display-only indicator variable for non-price promotion observations; $= 1$ if brand k is displayed, but not featured nor price promoted, by store i in week t , $= 0$ otherwise;

$FD_{ik,t}$ is an indicator variable for combined use of feature and display for non-price promotion observations; $= 1$ if brand k is featured and displayed by store i in week t , but not price promoted, $= 0$ otherwise;

R_i is a store indicator variable; $= 1$ if observation is from store i , $= 0$ otherwise;

W_t is a weekly indicator variable; $= 1$ if observation is from week t , $= 0$ otherwise;

α_{jkl} is the elasticity of brand k 's sales with respect to brand j 's price index in the current week; supported by neither feature nor display for $l=1$, by

4 Blattberg and Neslin (1990, p. 190) say: "While there are no published examples of using polynomial lags to measure the lag effects of promotion, the technique appears to be promising."

feature-only for $l = 2$, by display-only for $l = 3$, by feature and display for $l = 4$;

α_{Fk} , α_{Dk} , α_{FDk} are the current-week effects on brand k 's log sales resulting from brand k 's use of feature-only (F), display-only (D), and feature and display (FD); each is the effect in the absence of a discount for brand k ;

$\beta_{kl,u}$ is the elasticity of brand k 's sales in week t relative to brand k 's price index with support l in week $t - u$;

$\gamma_{kl,v}$ is the elasticity of brand k 's sales in week t relative to brand k 's price index with support l in week $t + v$;

ψ_{ik} and ξ_{tk} : store intercept for store i ($i = 1, \dots, N$), brand k , and week intercept for week t ($t = 1, \dots, T$), brand k , respectively;

$u_{ik,t}$ is a disturbance term for brand k in store i in week t .

Weekly indicator variables are included to account for seasonal effects and the effects of missing variables (e.g. manufacturer advertising and coupons).

This model accounts for the factors that can hide dips, presented in section 2. It includes flexible, multiple-week, pre- and post price promotion own-brand variables (implications *i-iii*), and it incorporates current cross-brand price-promotional instruments (implication *iv*). Through use of separate price discount variables for four promotion conditions, the model accommodates a.o. lagged price cut with feature effects, which were found to be significant in prior research (Papatla and Krishnamurthi 1996). In summary, the model includes lead- and lagged variables for temporary price discounts with four types of support: no support, feature-only, display-only, and feature and display. Thus, our model also accounts for dynamic effects for features and displays (documented also by Lattin and Bucklin 1989), to the extent that these promotions were accompanied by price discounts.⁵

We note that the model implicitly accounts for a decrease in promotional effectiveness during a multiple-week price promotion. This occurs if the first week's price promotion causes a dip in the second week's sales, which then takes away from the second week's price promotion effect, etc. In other words, the current effect of the promotion in this second week effect equals the current-week parameter (alpha, negative if $j = k$) plus the one-week postpromotion parameter (beta, generally positive). One rationale is that households who purchased promotional items in week one will be less responsive to the same promotion in a second week. Another

⁵ We also considered the existence of dynamic effects for feature and/or display activity without promotional price cuts. However, as consumers have no monetary incentive to accelerate their purchases in case of "value-less" promotions, we expect little or no dynamic effects. To check, we included dynamic effects for these variables in the models. The effects were significant in only a few cases so that parsimony justifies constraining them to be zero.

rationale is that consumers who engage in deal-to-deal purchasing will have reduced motivation to participate if a promotion is extended. We note that there are other ways of modeling the effect of past and future deals on current price response. For example, Foekens, Leeflang, and Wittink (1999) use a varying-parameter approach, i.e., the alphas are a function of past promotions.

Model Alternatives

We consider various alternatives to approximate the dynamic structure, separately for each of the brands. For example, we have to choose between three alternatives for the dynamic effect specification: the Almon-, Exponential decay-, and Unrestricted models. In addition, for each of these specifications we estimate the durations of the lag period (s) and lead period (s'). Moreover, for the Exponential decay model, we need to find the best values for the lagged effect decay parameter λ and lead effect decay parameter μ . Finally, for the Almon model we have to determine the best degrees for the lagged effect polynomial r and for the lead effect polynomial (r').

To accomplish this, we let the maximum lag period vary from zero to six weeks ($s = 0, 1, 2, \dots, 6$) and also let the maximum lead period vary from zero to six weeks ($s' = 0, 1, 2, \dots, 6$), for the three specifications. Both maxima are six weeks so that they are close to the average interpromotion period in the data sets (see section 4). For the Exponential decay model we let each decay parameter vary independently from 0.1, 0.2, $\dots$ up to 0.9. For the Almon model, we consider polynomials up to degree three for both the lag polynomial ($r = 0, 1, 2, 3$) and, independently, for the lead polynomial ($r' = 0, 1, 2, 3$). We note that the polynomial degree must be smaller than the maximum lead or lag length ($r < s$ and $r' < s'$).

Model Calibration

For a given brand, we proceed as follows:

1. We estimate model (4) by OLS for all alternatives. Since the Almon- and Exponential decay models impose constraints on the parameters of successive periods, we have to impose these constraints during estimation. To accomplish this, we compute linear combinations of both lead- and lagged predictors, using polynomial coefficients as weights for the Almon model and geometrically declining weights for the Exponential decay model (see also Judge et al. 1985, p. 357 and p. 388). The Unrestricted model does not impose constraints, and we estimate it by just including untransformed lead- and lagged predictors. Within each of the three dynamic effect specifications, we vary the lead and lag duration as well as the values of the decay

parameters (for the Exponential decay model) or the polynomial degrees for the lead and lagged effects (for the Almon model). Next, we choose among all alternatives the model that minimizes Akaike's Information Criterion: $AIC = \ln(SSR/n) + 2p/n$, where SSR = Sum of Squared Residuals for a given model, n = number of observations used to estimate this model, and p = number of predictors included in this model. We use AIC because it can be used to compare the nonnested models and it ranked high in a comparison of 11 model selection criteria (Rust et al., 1995). While the Schwarz criterion ranked first overall (Rust et al. 1995, Table 1) and AIC second, for dynamic regression models AIC showed a slightly better performance than the Schwarz criterion. This is important because we are interested in the structural nature of lead- and lag effects. That is, we want to report the best possible estimates of the magnitudes of these effects, based on a model that has the highest possible 'structural' validity. Since the AIC is known to penalize models with extra parameters less heavily than the Schwarz criterion, we expect the models selected on AIC to provide a more flexible representation of dynamic effects (e.g., as in the Unrestricted model), and more flexibility means less bias. For example, the Almon model is known to produce biased parameter estimates if the assumed degree of the polynomial is too low or if the assumed lag length is too short (Judge et al. 1985, p. 358). We acknowledge that the best AIC model may not provide the best possible forecasts, but our focus is not on maximizing forecast accuracy;⁶

2. We test for pooling: we test whether the assumption of parameter homogeneity across the stores is tenable with a Chow test. The unpooled models have store-specific effects for the marketing instruments, but common weekly effects;
3. We test for disturbance-term assumptions: we test the assumptions of homoscedasticity (across the stores) and zero first-order autocorrelation with the Bera-Jarque test.

4. Data

We use weekly store-level scanner data from ACNielsen for two product categories to calibrate the models. We use pooled data to estimate the effects across all stores. The first data set of 52 weeks pertains to the three largest national brands (items) in the 6.5

⁶ A third alternative, the F-test, cannot be applied since this test requires that all models are nested. Since we choose AIC as the model selection criterion, we do not report Schwarz criterion values or F-test statistics in the results section.

oz. canned tuna fish product category in the U.S. ($T=52$, $J=3$). The data are from 28 stores belonging to one supermarket chain in a metropolitan area ($N=28$). We show descriptive statistics based on 1456 observations for each of these tuna brands in the first three columns of Table 4.1. The average interpromotion time varies from 3.2 to 9.9 weeks across the brands. Each brand is promoted frequently.

Table 4.1: Descriptive statistics for tuna and tissue categories

	Tuna brand			Tissue brand					
	1	2	3	A	B	C	D	E	F
Number of stores	28	28	28	24	24	24	24	24	24
Number of weeks	52	52	52	52	52	52	52	52	52
Number of observations	1456	1456	1456	1248	1248	1248	1248	1248	1248
Brand share percentage	46.6	30.8	22.5	5.2	12.9	10.8	14.9	46.9	9.2
Average regular price	0.86	0.87	0.86	1.15	1.23	1.16	1.17	0.62	1.44
Average interpromotion period in weeks	3.2	9.9	5.9	6.2	5.7	4.8	5.4	4.9	8.2
(standard deviation)	(2.5)	(10.7)	(5.0)	(5.5)	(3.5)	(3.9)	(4.0)	(3.9)	(8.2)
# Price promotions w.o. support	245	101	220	103	48	53	36	7	104
# Price promotions with feature-only	65	70	16	46	20	18	29	49	9
# Price promotions with display-only	59	18	77	42	14	17	44	6	5
# Price promotions with feature and display	241	79	112	76	136	149	130	105	81
# Non-price promotions with feature-only	8	8	2	11	59	43	58	42	14
# Non-price promotions with display-only	44	23	47	9	6	13	21	16	9
# Non-price promotions with feature and display	22	11	10	1	2	2	16	12	5

The second data set (also 52 weeks) pertains to the six largest national brands in the toilet tissue product category in the U.S. ($T=52$, $J=6$). Each brand is composed of a number of SKU's, representing different sizes and forms. The database contains data aggregated across the SKU's to define brand-level variables. Unit sales is defined as the total number of sheets per package sold, and unit price is the price per thousand sheets. We developed a procedure to impute the weekly regular prices required for the price index variables, since these were not included in this data set.⁷ The data are from 24 stores of different chains in one region of the US ($N=24$). We show descriptive statistics based on 1248 observations for each of the toilet tissue brands in the last six columns in Table 4.1. The average interpromotion period for the toilet tissue brands varies from 4.8 to 8.2 weeks. The frequencies with which the brands are promoted is somewhat lower for this category than for tuna.

We note that Narasimhan, Neslin, and Sen (1996) obtained consumer-based ratings of “ability to stockpile” for 108 product categories. In their listing tuna fish is rated first and toilet tissue rated fourth. Although their measure is incomplete with regard to actual stockpiling of products, and is subject to an unknown degree of error,

⁷ Details about this procedure are available from the first author.

these two product categories should provide excellent opportunity to examine the postpromotion dip controversy.

There is an important difference between the two data sets, however. Whereas the tuna data are at the SKU level, the tissue data are at the brand level (multiple SKU's). For example, there may be within-brand switching between SKU's due to heterogeneity in the timing of promotions at the SKU level. As a result, the current price index elasticity at the brand level may be smaller (alpha less negative) than the current price index elasticity at the SKU level. At the SKU level, the model does not accommodate dynamic effects for other SKU's belonging to the same brand. However, at the brand level, the dynamic effects represent the (net) effect across the SKU's. If consumers are just as inclined to stockpile one SKU as another, for a given brand, then the dynamic effects may be larger (e.g., beta and/or gamma more positive) at the brand level. This suggests that the dip, in a relative sense, may be larger for tissue. However, tuna seems to be especially easy to stockpile.

5. Results

We estimate the models with data pooled across stores. Since promotional effects may differ between stores, we test the null hypothesis of parameter homogeneity. To do this, we (also) estimate the best models with store-specific current and dynamic effect parameters, but homogenous weekly effects. We report the p-values of the Chow test in the panel headed "Pooling test" in the lower part of Table 5.1. We do not reject the null hypothesis for any of the nine brands.

To test the error assumptions, we use the Bera-Jarque test (Körösi, Mátyás, and Székely, 1992, pp. 173-180). We first test the error assumptions simultaneously, and report the p-values in the panel headed "Error assumption tests" of Table 5.1. If the outcome is a rejection (i.e., a p-value lower than 5 percent), the hypotheses of zero autocorrelation and homoscedasticity are also tested separately. We use a significance level of 2.5 percent for these tests (based on Bonferroni's rule). We report the p-values for these tests also in Table 5.1. We find four cases of non-zero autocorrelation and seven cases of heteroscedasticity.

Given some error-term assumption violations, we re-estimate the models where necessary, using Iterative GLS (IGLS). Kmenta (1986, pp. 609-622) describes IGLS accounting for non-zero autocorrelation (IGLS-A), for heteroscedasticity (IGLS-H), or both (IGLS-AH). We iterate the GLS-procedure until convergence. In the panel headed "Re-estimation results" in Table 5.1 we show the estimation procedure used for each brand. We note, however, that the IGLS parameter estimates are very close to

Table 5.1: Model calibration results

	Tuna brand					Tissue brand			
	1	2	3	A	B	C	D	E	F
<i>Model selection</i>									
AIC Almon model (A)	<u>-2.008</u>	<u>-1.967</u>	<u>-1.703</u>	<u>-1.285</u>	<u>-2.203</u>	<u>-1.917</u>	<u>-2.360</u>	<u>-2.851</u>	<u>-2.6616</u>
AIC Exponential decay model (E)	-1.994	<u>-1.967</u>	-1.693	-1.268	-2.166	<u>-1.917</u>	-2.338	-2.842	<u>-2.6622</u>
AIC Unrestricted model (U)	-2.002	<u>-1.967</u>	-1.693	-1.283	-2.195	<u>-1.917</u>	<u>-2.369</u>	<u>-2.855</u>	<u>-2.630</u>
Model choice	A	A/E/U ^a	A	A	A	A/E/U ^a	U	U	E
Lead period s' (weeks)	4	0	6	0	5	0	5	5	3
Lag period s (weeks)	3	0	1	6	5	1	6	1	6
<i>Re-Estimation results</i>									
estimation procedure ^b	IGLS-AH	IGLS-AH	IGLS-AH	OLS	OLS	IGLS-H	IGLS-H	IGLS-H	IGLS-AH
number of observations for estimation	1260	1456	1260	1104	1008	1224	1008	1104	1032
number of parameters	104	91	103	111	116	105	132	120	103
R ²	0.941	0.834	0.915	0.883	0.929	0.957	0.975	0.987	0.947
<i>Pooling test</i>									
p-values Chow test	0.986	1.000	0.996	1.000	1.000	0.995	1.000	1.000	1.000
<i>Error assumption tests</i>									
p-value simultaneous test ($\alpha = 0.05$)	0.000	0.000	0.000	0.074	0.020	0.000	0.000	0.000	0.000
p-value zero autocorrelation ($\alpha = 0.025$)	0.000	0.000	0.000	-	0.057	0.657	0.060	0.846	0.005
p-value homoscedasticity ($\alpha = 0.025$)	0.000	0.000	0.000	-	0.035	0.000	0.000	0.000	0.001

^a The models are exactly the same in this situation.

^b IGLS-A = Iterative GLS that accounts for non-zero autocorrelation, IGLS-H = IGLS that accounts for heteroscedasticity, and IGLS-AH accounts for both.

the OLS-estimates (the primary benefit of IGLS lies in obtaining more valid estimated standard errors). In the same panel we report the number of observations used for estimation and the number of parameters. The number of observations used differs from the original number, due to the inclusion of lead- and lagged variables. For example, tuna brand 1 has a lead effect of four weeks and a lagged effect of three weeks. Hence we can only use weeks 4 through 48 of the 52 weeks for each store. The number of observations is 28 (stores) times 45 weeks or 1260 for this brand. The number of parameters includes the indicator variables for stores and weeks. Below these numbers we show the R² values which are between 0.834 and 0.941 for tuna, and between 0.883 and 0.987 for toilet tissue.⁸

We provide summary statistics for the OLS estimation results of (4) for the nine brands in the panel headed “Model selection” in Table 5.1. Specifically, we report the

⁸ We also performed multicollinearity analyses. We computed the condition indices from the predictor matrices for the “best models”. They were higher than 30 for four brands. Hence there is a substantial amount of multicollinearity in the data, for those brands. This multicollinearity stems from the allowance for multiple and flexible lead- and lagged promotion effects. Therefore, individual parameter estimates may not be reliable, so that we base our conclusions on total dynamic effects, summed over pre- and postpromotion periods.

AIC for the AIC-minimizing alternative of each of the three model specifications. For each brand, we then choose the model specification that minimizes AIC (underlined values), and find dynamic price promotion effects for eight out of nine brands (tuna brand 2 being the exception). For the tuna brands, the preferred model on average has a longer lead- than lag length. For tissue brands, the lag period tends to be longer than the lead period. We also see in Table 5.1 that the AIC-values for the Almon model are lowest for four out of nine brands. The Unrestricted model has the lowest AIC for two brands, and the Exponential decay model for one brand. For the two remaining brands, the three models are exactly the same: tuna brand 2 (no dynamic effects), and tissue brand C (one-week lagged effect). In the former case, the model without dynamic effects is best, and in the latter case, the lagged effect lasts one week only, so that there is no opportunity to consider alternative ways to describe dynamic patterns.

The lead- and lag lengths vary across the brands. However, for a given brand, the three dynamic-effect specifications often yield the same lead and lag lengths (not shown). Across the nine brands, 83 percent of the lead and lag lengths are exactly the same for the three AIC-minimizing specifications. For all brands we find that if one AIC-minimizing specification yields a non-zero lead- or lag length, the other specifications also do.

We present *average* estimated effects from the best individual models, which have been re-estimated with IGLS if appropriate, in Tables 5.2 and 5.3. We report four current price index elasticities for each category, averaged across the brands, in Tables 5.2.a and 5.3.a: price cut without support, price cut with feature-only support, price cut with display-only support, and price cut with feature and display support. For both product categories, the price elasticity is the lowest for unsupported price cuts and the highest for price cuts with feature and display support, exactly as we would expect. Except for price cuts without support, the average own-brand price elasticities are larger for the tissue category than for the tuna category. However, these results are not comparable, because there are differences in the recency of the data (the tuna data are from 1986/1987, whereas the tissue data are from 1992), the regions represented (the data are from different US regions), the store types, etc.

In Tables 5.2.b and 5.3.b we report the results from a simulation exercise. We use the brand-specific current-, lead- and lagged effect parameter estimates from the best individual models to calculate the current sales effect and the pre- and postpromotion sales effects for a 20 percent promotional price cut for each of the four types of support. A 20 percent price cut is typical in both product categories. We show in units the average *current sales effect* as well as the average *dynamic sales effect*, which is the sum of the pre- and postpromotion sales effects. We then combine the current and dynamic effects into a *net sales effect*, which is the total promotion effect over time.

Table 5.2: Tuna data: average current-, dynamic-, and net sales promotion effects

Table 5.2.a: average parameter estimates for best models				
	support type ^a			
	no F, no D	F-only	D-only	F and D
Average current own-brand price elasticity	−3.1	−4.0	−4.5	−5.2
Table 5.2.b: simulated sales effects using best models				
	sales impact of a 20 % price cut supported by			
	no F, no D	F-only	D-only	F and D
Current sales effect	126	190	222	284
Dynamic sales effect	−28	−24	−8	−21
Net sales effect (=current+dynamic)	98	165	214	263
Table 5.2.c: percent net gains for all models				
	support type ^a			
	no F, no D	F-only	D-only	F and D
Best model ^b	78%	87%	96%	93%
Almon model	78%	87%	96%	93%
Exponential decay model	86%	84%	102%	91%
Unrestricted model	79%	88%	98%	93%

^a no F, no D = neither feature nor display, F-only = feature-only, D-only = display-only, F and D = feature and display.

^b The percent for the best model is the ratio of the net sales effect over the current sales effect, reported in Table 5.2.b.

⁹ The *percent net gain* is obtained by the ratio of the net sales effect to the current sales effect. The results in Tables 5.2.b and 5.3.b show that the difference between current- and net sales effects can be quite large, so that the profit implications may change greatly if dynamic effects are taken into account.

Tables 5.2.b and 5.3.b show that more (feature/display) support produces larger own-brand *current sales effects*, for both product categories. The *dynamic sales*

⁹ We do not report separate lead- and lagged effects, because they may be confounded. For example, one postpromotion period may interfere with the prepromotion period of the next promotion, if these promotions are close in time. In addition, lagged terms can capture prepromotion dips since these dips are due to anticipatory responses that consumers may base on current- and lagged prices. We do not exclude lead effects from our model, however. First, since the length of the interpromotion period varies, pre- and postpromotion periods do not necessarily interfere. Second, lag terms capturing prepromotion dips should be modeled differently from lag terms capturing the effects of purchase acceleration and lack of consumer inventory sensitivity. For example, the sales pattern including pre- and postpromotion dips may look as follows: promotional sales spike (week 1) - postpromotion dip (weeks 2-3) - regular sales (week 4) - prepromotion dip (weeks 5-6) - promotional sales spike (week 7). We prefer to model such a pattern with a lagged effect specification (for weeks 2 and 3) and a lead effect specification (for weeks 5 and 6) over a model with a highly irregular five-week lagged effect. However, for the purpose of understanding the magnitudes of dynamic effects, it is sufficient to combine these effects into a “total dynamic sales effect”.

Table 5.3: Tissue data: Average current-, dynamic-, and net sales promotion effects

Table 5.3.a: average parameter estimates for best models				
	support type ^a			
	no F, no D	F-only	D-only	F and D
Average current own-brand price elasticity	-2.9	-5.4	-5.2	-5.7
Table 5.3.b: simulated sales effects using best models				
	sales impact of a 20% price cut supported by			
	no F, no D	F-only	D-only	F and D
Current sales effect	227	542	529	689
Dynamic sales effect	-57	-72	+50	-84
Net sales effect (=current+dynamic)	170	470	579	605
Table 5.3.c: percent net gains for all models				
	support type ^a			
	no F, no D	F-only	D-only	F and D
Best model ^b	75%	87%	110%	88%
Almon model	70%	89%	111%	88%
Exponential decay model	110%	96%	117%	93%
Unrestricted model	75%	86%	110%	87%

^a no F, no D = neither feature nor display, F-only = feature-only, D-only = display-only, F and D = feature and display.

^b The percent for the best model is the ratio of the net sales effect over the current sales effect, reported in Table 5.3.b.

effects, however, do not show a similar pattern. For example, a 20 percent price cut with display has a relatively small negative dynamic effect for tuna (-8 units), and even a positive dynamic effect for tissue (+50 units). A display-extension effect (argument e) is perhaps the most plausible explanation for this phenomenon. In any event, the *net sales effect* is a managerially more relevant measure than the current sales effect. For instance, if we use the current sales effects only for the tissue category, a manager may believe a price cut with feature-only has a larger sales effect than a price cut with display-only support (542 vs. 529 units). However, the net sales effect indicates that the contrary is true (470 vs. 579 units).

If we had to recommend one of the three model specifications to a manager, we would choose the Unrestricted model, although Table 5.1 indicates no single specification can be considered best. One reason is that the substantive results from the Unrestricted model are the closest to those of the best models. This can be inferred from the results in Tables 5.2.c and 5.3.c, in which we report average percentages net gain for the price cuts with different types of support. The first line shows these percentages based on the results for the best model in Tables 5.2.b and 5.3.b. The second line shows the percentages for the Almon model, the third line for the Exponential decay model, and the fourth line for the Unrestricted model. Although

the Almon model and the Unrestricted model give very similar results, the average absolute percent deviation in the percent net gains (from the results for the best model) is slightly smaller for the Unrestricted model than for the Almon model: 1.2 percent (tuna) and 0.6 percent (tissue) for the Unrestricted model, versus 0.0 percent (tuna) and 2.6 percent (tissue). For the Exponential decay model, the deviations in net percent gain from the best model are much larger: 5.1 percent (tuna) and 13.4 percent (tissue). Hence, the Exponential decay model gives very different conclusions about the percent net gain from these promotions. Another reason is the ease of implementation of the Unrestricted model. Whereas the Almon- and Exponential decay model require transformations of lead- or lagged predictor variables, the Unrestricted model only requires that the user provides lead- and lagged predictor variables. In addition, the Almon- and Exponential decay model require a grid search, across polynomial degrees (Almon), or across decay parameters (Exponential decay), on top of a search for lead- and lag lengths. The Unrestricted model only requires the latter search.¹⁰

We see that the *percent net gain* figures for the best model are quite consistent across the two categories. For an unsupported price cut we have 78 percent (tuna) and 75 percent (tissue), for a feature-only supported price cut 87 percent for both categories, for a display-only supported price cut 96 percent (tuna) and 110 percent (tissue), and for a feature-and-display supported price cut it is 93 percent (tuna) and 88 percent (tissue). The percent net gains lower than 100 reflect pre- and/or postpromotion dips. For these cases, the purchase acceleration effects vary between 4 and 25 percent. These numbers appear to be consistent with the results from household-level studies.

We illustrate the performance (of the best model) by plotting actual and fitted sales. Instead of showing the graphs for all brands in all stores for all weeks (28 stores with 3 tuna brands with 52 observations each, and 24 stores with 6 tissue brands with 52 observations each), we select two illustrative examples. We take the same brands and stores as in Figure 1.1. Figure 5.1 shows actual and fitted sales for tuna brand 1 in store 28 for weeks 35-45. Fitted sales are obtained by using the model that includes current and dynamic effects (model (4)), and separately a model that includes current effects only. The latter “no-dynamic-effects model” was rejected in favor of the dynamic-effects model (see Table 5.1). Both models track actual sales quite well up to week 42. However, the dynamic-effects model shows a much smaller effect for the price promotion in week 43 than the no-dynamic-effects model. This store

10 Alternatively, we could use other model criteria (such as the Schwarz criterion) to decide which of the three specifications is best overall. However, this would be inconsistent, since we obtain the best model within each specification based on AIC. Thus, we focus on a comparison of effect sizes if one specification were applied for all brands, with the effects sizes for the best model.

had a price promotion for week 41 as well, causing negative postpromotion effects which are captured by the dynamic-effects model but not by the no-dynamic-effects model.¹¹

Figure 5.1: Tuna data: actual sales and predicted sales

The second example is from tissue brand E sold in store 21, shown in Figure 5.2. In this panel we see that the dynamic-effects model approximates the postpromotion dip better than the no-dynamic-effects model does (see weeks 37 and 38). Again, the no-dynamic-effects model was rejected (Table 5.1). This graph illustrates the subtleness of dynamic promotion effects, and it is more representative of brand sales graphs than the one in Figure 5.1.

6. Conclusions

We investigated one of the mysteries of sales promotion research: the lack of postpromotion dips in store data. From studies of household panel data it is known that consumers often accelerate their purchases in time and/or quantity due to promotions, which should result in a dip in purchases in the weeks following a promotion. This dip, however, is rarely observed in sales data. Extant arguments for the apparent lack of postpromotion dips imply that the dips may be difficult to detect by traditional models. Since brand sales are the aggregate of purchases across (heterogeneous) households, both pre- and postpromotion sales data may have

¹¹ The decreased effectiveness of the price promotion in week 43 is not due to prepromotion effects, since the brand was not price promoted after week 43.

Figure 5.2: Tissue data: actual sales and predicted sales

complex patterns. Essentially, sales are shifted from multiple future- and past periods into a current, promotion-based sales spike in a nontrivial way. Neslin and Schneider Stone (1996, p. 92) suggested that researchers “... conduct sophisticated distributed lag analyses on weekly sales data in the hope of measuring the postpromotion dip statistically”. Our modeling approach reflects the multitude of factors pertaining to dynamic promotion effects, and we obtain postpromotion dips convincingly.

We use an econometric model to regress brand-level sales on current-, lagged-, and lead price discount variables (price indices) for three different distributed lead-and lag structures: an Almon model, an Unrestricted dynamic effects model, and an Exponential decay model. The Unrestricted dynamic model is very flexible but not parsimonious. The Exponential decay model is the least flexible and the most parsimonious one, while the Almon polynomial model is in between these extremes. Importantly, we distinguish the effects of four types of price discounts: without support, with feature-only support, with display-only support, and with feature and display support.

We applied the models to nine brands in two product categories: tuna fish and toilet tissue. Within each of these three models, we varied lead and lag lengths as well as the parameters describing the lag structure. For each brand, we selected the model specification that minimizes Akaike’s Information Criterion. We tested the assumption of parameter homogeneity among stores (no evidence of heterogeneity), and re-estimated models by accounting for nonzero autocorrelation or heteroscedasticity where necessary.

Our main findings across two product categories are:

- Significant dynamic promotion effects exist (for eight of the nine brands);
- The dynamic effects can be substantial. Negative dynamic effects are indicative of acceleration effects which vary between 4 and 25 percent of the current sales effect, across the two categories and across different support activities for discounts. These numbers are consistent with the results from household-level studies which have found the acceleration effect to vary between 6 and 51 percent;
- The conclusion for researchers is that the postpromotion dip paradox does not have to exist: household-level studies and this store-level study find acceleration effects of comparable sizes;
- For managers, our results suggest that the results from models that accommodate only current sales effects from a promotion may be quite misleading. Managers should insist on obtaining the sum of the current and dynamic effects from a model that accounts for (1) purchase acceleration effects and (2) display extension effects.
- Given the complexity of dynamic sales promotion effects, it is advisable to use a flexible specification, such as the Unrestricted model or the Almon model. We find that the percent net gains (Tables 5.2.c and 5.3.c) are very similar for the Almon-, and Unrestricted models (while the Exponential decay model produces very different results). Overall, the Unrestricted model is the closest to the “best model” results. It is also the model that is easiest to implement.

An interesting future research issue in this context is the accommodation of within-store and between-store heterogeneity in models of store sales. Within-store heterogeneity can occur due to changes in the composition of the set of households that purchase items from the product category over time. The use of time-varying response parameters is one way to account for such effects. Between-store heterogeneity may result from customer-, assortment-, and other differences between stores. Even though we found no evidence in favor of this type of heterogeneity, it may be relevant for other categories. Hsiao, Appelbe, and Dineen (1993) provide a general framework for varying parameter panel data models.

From a substantive perspective, it will be useful to apply distributed lead- and lag models to other product categories to discover commonalities and idiosyncracies. With results on a much larger number of items it should be instructive to explore the effects of product category and brand measures, marketing activities (e.g., interpromotion periods and depth of price cuts) and consumer characteristics on the observed dynamics.

References

- Bell, David R., Jeongweng Chiang, and V. Padmanabhan (1999), "The Decomposition of Promotional Response: An Empirical Generalization," *Marketing Science*, forthcoming.
- Blattberg, Robert C., Richard Briesch, and Edward J. Fox (1995), "How Promotions Work," *Marketing Science* 14 (Part 2 of 2), G122-32.
- Blattberg, Robert C. and Scott A. Neslin (1990), *Sales Promotion, Concepts, Methods and Strategies*. Englewood Cliffs: Prentice Hall.
- Blattberg, Robert C. and Kenneth J. Wisniewski (1989), "Price-Induced Patterns of Competition," *Marketing Science* 8 (Fall), 291-309.
- Bucklin, Randolph E. and Sunil Gupta (2000), "Commercial Adoption of Advances in the Analysis of Scanner Data," *Marketing Science*, forthcoming.
- Chiang, Jeongwen (1991), "A Simultaneous Approach to the Whether, What and How Much to Buy Questions," *Marketing Science* 10 (Fall), 297-315.
- Christen, Markus, Sachin Gupta, John C. Porter, Richard Staelin, and Dick R. Wittink (1997), "Using Market-Level Data to Understand the Promotion Effects in a Nonlinear Model," *Journal of Marketing Research* 36 (August), 322-34.
- Doyle, Peter and John Saunders (1985), "The Lead Effect of Marketing Decisions," *Journal of Marketing Research* 22 (February), 54-65.
- Foekens, Eijte W., Peter S.H. Leeflang, and Dick R. Wittink (1994), "A Comparison and an Exploration of the Forecasting Accuracy of a Nonlinear Model at Different Levels of Aggregation," *International Journal of Forecasting* 10, 245-61.
- Foekens, Eijte W., Peter S.H. Leeflang, and Dick R. Wittink (1999), "Varying Parameter Models to Accommodate Dynamic Promotion Effects," *Journal of Econometrics* 89, 249-68.
- Grover, Rajiv and V. Srinivasan (1992), "Evaluating the Multiple Effects of Retail Promotions on Brand Loyal and Brand Switching Segments," *Journal of Marketing Research* 29 (February), 76-89.
- Gupta, Sachin, Pradeep K. Chintagunta, Anil Kaul, and Dick R. Wittink (1996), "Do Household Scanner Data Provide Representative Inferences From Brand Choices: A Comparison With Store Data," *Journal of Marketing Research* 33 (November), 383-98.

- Gupta, Sunil (1988), "Impact of Sales Promotion on When, What and How Much to Buy," *Journal of Marketing Research* 25 (November), 342-55.
- Hsiao, Cheng, Trent W. Appelbe, and Christopher Dineen (1993), "A General Framework for Panel Data Models With an Application to Canadian Customer-Dialed Long Distance Telephone Service," *Journal of Econometrics* 59, 63-86.
- Inman, Jeffrey, Leigh McAlister, and Wayne D. Hoyer (1990), "Promotion Signal: Proxy for Price Cut?," *Journal of Consumer Research* 17 (June), 74-81.
- Judge, George G., W.E. Griffiths, R. Carter Hill, Helmut Lütkepohl, and Tsoung-Chao Lee (1985), *The Theory and Practice of Econometrics*, 2nd edition. Singapore: John Wiley and Sons.
- Kalwani, Manohar U., Chi Kin Yim, Heikke J. Rinne, and Yoshi Sugita (1990), "A Price Expectation Model of Customer Brand Choice," *Journal of Marketing Research* 27 (August), 251-62.
- Kennedy, P.B. (1981), "Estimation with Correctly Interpreted Dummy Variables in Semilogarithmic Equations," *American Economic Review* 71, 801.
- Kőrösi, Gábor, László Mátyás, and István P. Székely (1992), *Practical Econometrics*, Aldershot: Avebury.
- Kmenta, Jan (1986), *Elements of Econometrics*, New York: Macmillan Publishing Company.
- Kumar, V. and Robert P. Leone (1988), "Measuring the Effect of Retail Store Promotion on Brand and Store Substitution," *Journal of Marketing Research* 25 (May), 178-85.
- Lattin, James M. and Randolph E. Bucklin (1989), "Reference Effects of Price and Promotion on Brand Choice Behavior," *Journal of Marketing Research* 26 (August), 299-310.
- Leone, Robert P. (1987), "Forecasting the Effect of an Environmental Change on Market Performance: An Intervention Time Series Approach," *International Journal of Forecasting* 3, 463-78.
- Litvack, David S., Roger J. Calantone, and Paul R. Warshaw (1985), "An Examination of Short-Term Retail Grocery Price Effects," *Journal of Retailing* 61 (Fall), 9-26.
- Moriarty, Mark M. (1985), "Retail Promotional Effects on Intra- and Interbrand Sales Performance," *Journal of Retailing* 61 (Fall), 27- 47.

- Narasimhan, Chakravarthi, Scott A. Neslin, and Subrata K. Sen (1996), "Promotional Elasticities and Category Characteristics," *Journal of Marketing* 60 (April), 17-30.
- Neslin, Scott A. and Linda G. Schneider Stone (1996), "Consumer Inventory Sensitivity and the Postpromotion Dip," *Marketing Letters* 7, 77-94.
- Papatla, Purushottam and Lakshman Krishnamurthi (1996), "Measuring the Dynamic Effects of Promotions on Brand Choice," *Journal of Marketing Research* 33 (February), 20-35.
- Rust, Roland T., Duncan Simester, Roderick J. Brodie, and V. Nilikant (1995), "Model Selection Criteria: An Investigation of Relative Accuracy, Posterior Probabilities, and Combinations of Criteria," *Management Science* 41 (February), 322-33.
- Stewart, Jon (1991), *Econometrics*. Hemel Hemstead: Philip Allan.
- Winer, Russell S. (1986), "A Reference Price Model of Brand Choice for Frequently Purchased Products," *Journal of Consumer Research* 13 (September), 250-6.
- Wittink, Dick R., Michael J. Addona, William J. Hawkes, and John C. Porter (1988), "The Estimation, Validation, and Use of Promotional Effects Based on Scanner Data," Working Paper, Johnson Graduate School of Management, Cornell University, Ithaca.