Promotion carryover as a missing data problem

Article

Published Version

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.

To link to this article DOI: http://dx.doi.org/10.2202/1542-0485.1053

www.reading.ac.uk/centaur

CentAUR

Reading's research outputs online

Promotion Carryover as a Missing-Data Problem

Garth J. Holloway[*]   Osman Aydogus[†]

[*]University of Reading, United Kingdom, garth.holloway@reading.ac.uk

[†]Ege University, Turkey, osman.aydogus@ege.edu.tr

Promotion Carryover as a Missing-Data Problem

Garth J. Holloway and Osman Aydogus

Abstract

An important feature of agribusiness promotion programs is their lagged impact on consumption. Efficient investment in advertising requires reliable estimates of these lagged responses and it is desirable from both applied and theoretical standpoints to have a flexible method for estimating them. This note derives an alternative Bayesian methodology for estimating lagged responses when investments occur intermittently within a time series. The method exploits a latent-variable extension of the natural-conjugate, normal-linear model, Gibbs sampling and data augmentation. It is applied to a monthly time series on Turkish pasta consumption (1993:5-1998:3) and three, non-consecutive promotion campaigns (1996:3, 1997:3, 1997:10). The results suggest that responses were greatest to the second campaign, which allocated its entire budget to television media; that its impact peaked in the sixth month following expenditure; and that the rate of return (measured in metric tons additional consumption per thousand dollars expended) was around a factor of 20.

KEYWORDS: promotion carryover, missing data, data augmentation, Gibbs sampling.

1. Introduction

This paper presents an alternative perspective on the estimation of response rates that evolve from the lagged impact of promotion on consumption of a good when promotion expenditures occur intermittently within a time series. Lagged responses are important for agribusiness firms as they seek to establish the true impacts of their promotion investments, the forms that future investments should take and the likely length of return per-unit of funds invested. In this context it is important to have an efficient method for evaluating lagged responses—one that is flexible, reliable and is easy to implement with the data that is usually at our disposal. Such data usually consist of a series of observations on consumption quantities, the prices at which they are purchased, and observations on other factors alleged to influence the consumption decision. One of these is promotion and it is now commonplace to estimate a regression model with consumption dependent on promotion expenditures, usually with some form of a finite distributed lag.

It is unsurprising, therefore, that the issue of modelling lagged responses has received considerable scrutiny in the commodity-promotion literature. Likely, all published studies on commodity promotion recognize the importance of lagged responses. Hence, providing a comprehensive list could easily absorb this Journal article. Nevertheless, a few examples (Liu and Forker, 1988; Liu and Forker, 1990; Pritchett, Liu and Kaiser, 1998; Chang and Kaiser, 1998; Vande Kamp and Kaiser, 1999 among others) are noteworthy in spirit due to their focus on temporal effects, and the notion that advertising effectiveness is likely to change over time.

By and large, these studies employ a 'traditional approach' to the lag problem. By 'traditional' we mean that a (parametric) relationship is proposed between the demand for a commodity in question and a promotion variable (typically the level of expenditure on a particular promotion medium) lagged by some unit of time, and its impacts are then assessed by applying conventional regression methods. One situation in which this approach is unlikely to yield reliable estimates is when promotion expenditures occur intermittently in the time series or, worse, when there is a one-shot campaign in which all funds are disbursed instantaneously. In this situation few observations in the promotion expenditure series tend to imply insufficient variability, multicollinearity, and the inevitable low levels of significance in conventional (frequentist) tests. While this need not be the case in all situations, it seems relevant to pursue an alternative methodology that is relevant to this situation and, thus, is the topic of this paper. One other situation that is potentially more troublesome, and arises frequently in commodity promotion research (all too frequently, we conjecture, but due to an obvious sampling bias we are unable to substantiate) is one in which the researcher (for disclosure reasons or otherwise) does not have access to the promotion expenditure data, but is aware that an expenditure occurred and has interest in measuring its response on some quantity of interest for which price and quantity data are available. In this situation the technique that we propose has an obvious advantage over traditional methods.

The proposed method is, in some senses, the antithesis of the usual distributed-lag approach to promotion evaluation. In that method, a set of (time-series constant) regression coefficients are estimated in relation to a set of (sequentially) lagged expenditure levels and the resulting estimates of the coefficients in question are used to infer the effectiveness (or ineffectiveness) of the promotion program. In the methodology we employ, we consider that the expenditures are given and fixed at their respective points in the time series, but that the parameters exemplifying the lagged consumption responses vary at each observation in the post-promotion regime. In view of this innovation our approach is novel, but it is not without a few important precedents. While different conceptually and analytically, they are similar in motivation and focus.

Kinnucan, Chang and Venkateswaran (1993) apply a time-varying parameter model to data on the first fourteen years (1971-84) of a fluid milk-marketing campaign in New York city in order to examine the 'wearout hypothesis;'[1] Reberte *et al.* (1996) estimate a time-varying parameter model to study the evolution of the impact of generic advertising on fluid milk sales during two major generic, fluid-milk advertising campaigns in New York City, 1986-92; and, in an early work, Ward and Myers (1979) study time-varying parameter models to measure advertising responsiveness over time. The main difference between these studies and our approach is the interpretation of the lagged responses. Each of the papers just mentioned postulates a parametric relationship between the (exogenous) lagged expenditure level and the (endogenous) demand quantity. In the present approach, we simply assume that the lag effect is 'missing data.' One immediate advantage of this approach is that the need to specify a parametric form for the carryover relationship is circumvented thereby facilitating lag structures that are more general than the ones that time-varying-parameter or distributed-lag models can handle.

In order to derive efficient estimates of the unknown quantities we assume that the complete set of (data × parameter) responses is contained within a single variable that is latent (or missing data) in the regression. Interpreting the lagged responses in this way has two advantages and one disadvantage, which is the main limitation of the method. The advantages are, first, that this slight perturbation of the conventional set-up allows some very familiar results on the normal linear model to be exploited (Zellner, 1991) in order to devise a numerically intensive, but efficient strategy for retrieving estimates of the parameters of interest; and, second, that this link forges another with an apparently unrelated procedure for Bayes estimates in the Tobit regression using Gibbs sampling and data augmentation (Chib, 1992). The latter methods are numerically intensive, sampling-based approaches to estimating unknown parameters and missing (or latent) data that rely (theoretically) on the properties of Markov-chains and (practically) on the ability to sample efficiently from specific probability distributions. The seminal contribution on Gibbs sampling is Gelfand and Smith (1990) and on data augmentation is Tanner and Wong (1987). A very readable introduction to Gibbs sampling (at an undergraduate statistics level) is Casella and George (1992). For examples in agricultural settings see Dorfman (1996, 1997) and for a closely related application to commodity promotion see Benson, Breidt and Schroeter (2002). Below we show how these methods—underexploited, we feel, by commodity promotion researchers—can be used to derive a robust method for estimating promotion carryover rates.

Section two outlines some basic concepts that provide the basis for the method; section three presents the estimation algorithm; section four presents experiments; section five presents

---

[1] There are three terms in the literature, which some authors use interchangeably to refer to the declining response to an advertising medium, but, in other contexts, have distinguishable features. The terms are 'carryover,' 'decay,' and 'wear-out.' Carryover—the topic of this paper—refers to the (empirically observed) phenomenon that an advertising program will have some effect, even after the program has ceased; decay refers to the lessened impacts of the same phenomenon in future periods; and wearout refers to the declining effectiveness of the same program over time as the promotion program is repeated, but consumers become less responsive to the various media.

background to the empirical application to the Turkish data; section six presents the application; and section seven concludes.

2. **Basic Concepts**

Consider a time series of observations $t = 1, 2, .., T$ on the reduced-form relationship between consumption levels, $y \equiv (y_1, y_2, .., y_T)'$; a set of seasonal effects, the prices of competing products, or any other factors that impact the reduced-form, summarized by the regressor matrix, $x \equiv (x_1, x_2, .., x_K)$, $x_1 \equiv (x_{11}, x_{21}, .., x_{T1})'$, $x_2 \equiv (x_{12}, x_{22}, .., x_{T2})'$, .., $x_K \equiv (x_{1K}, x_{2K}, .., x_{TK})'$; and a vector of unobserved random shocks $u \equiv (u_1, u_2, .., u_T)'$; assume at time $\tau+1$ there is a once-and-for-all promotion expenditure; and assume, in addition, that this expenditure generates a set of non-negligible lagged responses $z_{\tau+1}, z_{\tau+2}, .., z_T$. Then, in terms of the notation just defined, the regression,

$$y \quad = \quad x\,\beta + u, \tag{1}$$

is miss-specified by measure $z \equiv (0, 0, .., 0, z_{\tau+1}, z_{\tau+2}, .., z_T)'$, because the data evolve according to

$$y \quad = \quad z + x\,\beta + u; \tag{2}$$

or, in a notation that will prove convenient later,

$$y^* \quad = \quad x\,\beta + u, \tag{3}$$

where $y^* \equiv y - z$. In this context, the problem of estimating promotion responses reduces to one of estimating the nonzero elements of the vector z, which vary across the post-promotion component of the time series and depict the lagged impacts of the promotion expenditure on the observed consumption quantities, y. In practice, z will depend on two factors. One is the actual level of promotion investment (that is, real expenditure on advertising) and the other is consumers' responsiveness to the promotion expenditure, its timing, its composition, and a host of random factors that the model in (1) leaves unexplained. Accordingly, define $\alpha \equiv (0 .. 0, \alpha_{\tau+1}, .., \alpha_T)'$, as the (time-varying) set of lagged responses and assume that the once-and-for-all expenditure is the level $\kappa_{\tau+1}$ so that $z \equiv \kappa_{\tau+1}\,\alpha \equiv \kappa_{\tau+1}(0, 0, .., 0, \alpha_{\tau+1}, \alpha_{\tau+2}, .., \alpha_T)'$. Our task in the remainder of the paper is to derive a robust method for estimating $\alpha \equiv (0, 0, .., 0, \alpha_{\tau+1}, \alpha_{\tau+2}, .., \alpha_T)'$.

This task is made easy by exploiting the redefinition of the (true, data-generating) model in (3) and some fairly recent advances in Bayesian estimation of regression relations that depart slightly from the normal-linear model. Before outlining the estimation strategy it is useful to articulate in terms of (1)-(3), the observed and the estimable quantities. The quantities y and x are observed data; z and therefore y^* are unobserved data; , of course, and u are unobserved but u is assumed to have a normal distribution with mean vector 0 and variance-covariance $\sigma^2 I_T$.

3. **Estimation**

In terms of equation (3) and the particular notation we have adopted, the task of estimating the response rates, α, is now made very easy. Suppressing dependence on the observed data $(x,$

y), the joint posterior distribution for the unobserved quantities (α, , σ), using a non-informative prior, has the following component conditional distributions

$$\alpha_t \mid \beta, \sigma \quad \sim \quad \text{Normal}(\ \hat{\alpha}_t\ ,\ V_{\alpha_t}\ ), \qquad t = \tau+1, \tau+2, .., T;$$

$$\beta \mid \sigma, \alpha \quad \sim \quad \text{Normal}(\ E_\beta\ ,\ V_\beta\ ); \tag{4}$$

$$\sigma \mid \alpha, \beta \quad \sim \quad \text{Inverted-Gamma}(\ \upsilon\ ,\ s^2\ ),$$

where $\hat{\alpha}_t \equiv (y_t - x_t\ )/\kappa_t$, $V_{\alpha_t} \equiv \sigma^2/\kappa_t^2$, $t = \tau+1, \tau+2, .., T$; $E \equiv (x'x)^{-1} x'y^*$, $V \equiv \sigma^2 (x'x)^{-1}$; $\nu \equiv T\text{-}K$ and $s^2 \equiv (y^* - x\ )'(y^* - x\ )/\nu$; and the definitions and results follow easily from applying the simple manipulations in Zellner (equations (3.26)-(3.29), pp. 65-66) with the model in (3) in place of the model in (2) (which is Zellner's description of the normal multiple regression model, equation (3.25), p. 65). It follows immediately that a Gibbs-sampling, data-augmentation algorithm for retrieving estimates of the unknown quantities proceeds as follows:

Step 1: Select starting values $\beta^{(s)}$ and $\sigma^{(s)}$.

Step 2: For each $t = \tau+1, \tau+2, .., T$, draw $\alpha_t^{(s)} \sim \text{Normal}(\ \hat{\alpha}_t\ ,\ V_{\alpha_t}\ )$, where $\hat{\alpha}_t$ and V_{α_t} are defined with respect to $\beta^{(s)}$ and $\sigma^{(s)}$, above.

Step 3: Draw $\beta^{(s+1)} \sim \text{Normal}(\ E_\beta\ ,\ V_\beta\ )$, where E_β and V_β are defined with   (5) respect to $\alpha^{(s)}$ and $\sigma^{(s)}$, above.

Step 4: Draw $\sigma^{(s+1)} \sim\ $ Inverted-Gamma$(\ \nu\ ,\ s^2\ )$, where s^2 is defined with respect to $\alpha^{(s)}$ and $\beta^{(s+1)}$, above.

Step 5: Repeat steps 2-4 until s equals some predetermined limit, say, S^1, at which convergence is achieved.

Step 6: Resetting s = 0, repeat steps 2-4 until s equals some predetermined limit, say, S^2, and collect output { $\alpha^{(s)}$ s = 1, 2, .., S^2 }, { $\beta^{(s)}$ s = 1, 2, .., S^2 } and { $\sigma^{(s)}$ s = 1, 2, .., S^2 }.

Because the draws in this algorithm form a Markov chain with desirable convergence characteristics, its output can be used to obtain reliable estimates of the lagged promotion responses, the effects of other factors on consumption, and the variance parameter. The algorithm for obtaining estimates of the parameters and σ is known as Gibbs sampling and the step for obtaining the (latent) lagged responses, α, is known as data augmentation. Although the two techniques are in wide application in applied Bayesian statistics and econometrics, they have

not, until recently (Benson, Breidt and Schroeter (2002)), been applied to evaluate agribusiness promotion programs. Our interest lies principally in their use in estimating the lagged responses to three non-consecutive promotion campaigns in the Turkish pasta industry (1996:3, 1997:3, 1997:10) using a consistent set of time series observations (1993:5-1998:2). However, before turning to consider these estimates we examine the results of some experiments that illustrate well the power of the method and its ease of implementation.

4. Experimental Evidence

We consider four experiments. Consistent across each experiment we draw four vectors of covariates from a multivariate-normal distribution with mean 0 and variance $^2 I_\tau$; set each element of the coefficient vector $\equiv (\beta_1, \beta_2, \beta_3, \beta_4)'$ equal to one; and set the error variance constant at $\sigma = 1$. In addition, and for illustrative purposes, we assume a single promotion expenditure at the sixth-last observation in the time series, $\kappa_{T-5} = 100$, where T denotes the last observation in the time series; and adopt the carryover profile consisting of the three responses $(\alpha_{T-5}, \alpha_{T-4}, \alpha_{T-3}) = (3/3,2/3,1/3)$. Responses at all other points in the time series are set equal to zero so that the (latent) promotion responses satisfy $z \equiv 100 \times (0, 0, .., 0, 3/3, 2/3, 1/3, 0, 0, 0)'$. We focus on the precision with which each of the nonzero elements is estimated, given that the researcher knows that a promotion injection occurred in the sixth last period, but does not know its impact and, in particular, is unaware that the last three lagged responses are zero. We also keep consistent across the experiments the number of iterations in the convergence sample (commonly referred to as the 'burn-in' sample), as well as in the Gibbs-sample, from which the parameters are estimated; and we select (conservative) levels $S^1 = 5,000$ and $S^2 = 10,000$, respectively. Finally, the dependent variables in the respective experiments are defined according to the model in equation (2), with the unknown error vector replaced by a vector of pseudo-random numbers,

$$\hat{\mathbf{u}} \sim \text{Normal}(0 , \sigma^2 I_T). \tag{6}$$

The first experiment sets T = 50, the second sets T = 100, the third sets T = 500, and the fourth sets T = 1000. It is useful to keep these quantities in mind in the context of examining the empirical application that follows. That application consists of a 58-observation time series.

Table 1 presents the results of the respective experiments. The first column reports parameter definitions and columns two through five report parameter and missing-data estimates (posterior means) obtained from output of the Gibbs samples. The numbers in parentheses are the ratios of posterior-means to estimates of posterior standard deviations as measured by standard deviations in the Gibbs sample. Although they do not correspond to a well-defined sampling density, these measures are useful as a summary of 'spread' relative to their means and, hence, whether the corresponding highest-posterior-density regions are likely to contain zero. Henceforth, we use the term 'significant' to indicate 95% highest posterior densities that do not contain zero.

Summary statistics are reported at the base of each column. They include the squared correlations between the observed consumption quantities and the means of the predictive densities (R^2); the condition number of the covariate matrix (Condition(x'x)); parameter values defining the experiment (T, σ, S^1 and S^2); and the minutes of real time (Execution Time) taken to complete the experiment (on a Dell Desktop Running MATLAB version 5.1, using Windows

'98, an Intel 330mhz co-processor and 160 megabytes of RAM). These and other programs required for implementation are available upon request.

Table 1. Experimental Evidence, β = (1.0, 1.0, 1.0, 1.0), = 1.0, = (1.0, 0.67, 0.33, 0.0, 0.0, 0.0), x ~ MVNormal(0,I$_4$).

Effect	Expt. 1	Expt. 2	Expt. 3	Expt. 4
$\hat{\beta}_1$	1.00	1.00	1.00	1.00
	(618.18)	(783.70)	(2205.41)	(3014.04)
$\hat{\beta}_2$	1.00	1.00	1.00	1.00
	(639.29)	(839.88)	(2255.06)	(3006.76)
$\hat{\beta}_3$	1.00	1.00	1.00	1.00
	(534.20)	(952.44)	(2340.60)	(3078.81)
$\hat{\beta}_4$	1.00	1.00	1.00	1.00
	(628.23)	(938.51)	(2244.70)	(3066.32)
$\hat{\sigma}$	1.14	1.22	0.99	1.05
	(4.32)	(6.64)	(15.54)	(22.12)
$\hat{\alpha}_{T-5}$	1.01	1.00	1.00	1.00
	(90.30)	(89.31)	(99.60)	(96.21)
$\hat{\alpha}_{T-4}$	0.67	0.65	0.66	0.67
	(62.82)	(57.76)	(66.98)	(65.37)
$\hat{\alpha}_{T-3}$	0.32	0.34	0.33	0.32
	(29.35)	(30.53)	(33.60)	(31.35)
$\hat{\alpha}_{T-2}$	-0.01	0.01	0.01	-0.02
	(-0.76)	(0.78)	(0.72)	(-1.46)
$\hat{\alpha}_{T-1}$	0.00	0.00	-0.01	0.00
	(-0.14)	(-0.43)	(-0.97)	(-0.42)
$\hat{\alpha}_T$	-0.01	0.00	0.01	0.02
	(-0.81)	(-0.23)	(0.98)	(1.86)
R^2	0.99	1.00	1.00	1.00
κ_{t-5}	100.00	100.00	100.00	100.00
Condition(**x'x**)	1.92	1.45	1.31	1.14
T	50	100	500	1000
σ	1.00	1.00	1.00	1.00
S^1	5000	5000	5000	5000
S^2	10000	10000	10000	10000
Execution Time	2.24	4.03	67.13	294.58

Note: Numbers in parentheses are parameter estimates divided by their standard deviations in the Gibbs sample. Execution Time is real time in minutes.

The results in general—but the estimates on the impacts of the lagged effects on the last three observations in particular—serve to motivate several conclusions. First, the methodology provides fairly accurate estimates of all parameters of interest. Second, the method is capable of distinguishing lagged impacts that are positive from those that are zero. Third, the usual small-sample biases that one expects from the smaller data sets (T = 50 and T = 100) appear to be small in comparison to the results generated from the larger samples (T = 500 and T = 1000). Fourth, the technique can be implemented and executed with a minimum of cost of real time, particularly for smaller data sets like the ones we frequently encounter in time-series applications to agribusiness promotion. Fifth, these results are obtained from models in which the ratio of systematic to unsystematic model variation (approximately 4:1 in the present case) compares reasonably to many empirical applications in the literature. Therefore, we expect that the procedure is likely to perform reasonably well in the empirical application that follows.

Before turning to the empirical application, it is worth mentioning one variation in the estimation that arises due to the hypothesis that the lagged effects should be non-negative. This variation is to sample the missing data from normal distributions that are truncated at zero, instead of from the complete normal distribution appearing in the first line in (4). Sampling from the truncated-normal distribution can be achieved by using a simple accept-reject mechanism. Alternatively, a more efficient one-for-one draw can be implemented in terms of the probability integral transform, as in Chib (1992).

5. Background to the Empirical Application

Despite only recent emergence as a major domestic consumption item, pasta production in Turkey is sizable by world standards. With an annual production capacity of 700,000 metric tonnes, it was the fifth largest producing nation in 1997. Production capacity reached 800,000 tonnes in 1998 and appears to have increased steadily since then. In the last two decades the industry realized unprecedented growth, not only in terms of capacity but also in terms of exports. Pasta producers were quick to take advantage of investment incentives offered by the government with the net result that capacity and actual production have almost doubled in the last decade (table 2).

Table 2. Pasta industry trends (thousand metric tonnes)

year	production	consumption	exports	estimated capacity
1985	217	208	9	340
1990	295	282	13	n.a.
1995	439	328	111	530
1996	430	350	109	n.a.
1997	550	365	134	700

Source: Turkish Pasta Producers' Association.

With domestic demand incapable of absorbing unprecedented production increases, producers turned attentions to foreign markets with the consequence that pasta exports increased more than tenfold in the period 1995-97 (from 9,000 tonnes in 1985 to 134,000 tonnes in 1997). Although Turkey's share in world pasta exports is relatively small compared to the largest exporting nation, Italy, Turkey is actually the second largest world exporter. Whether current trends in growth continue depends in large part on domestic circumstances in its major

destination market, the Russian Federation. Due to an anti-damping tax imposed on Turkish pasta exporters the U.S. market has been effectively closed to Turkish firms.

Including the smallest incumbents, which produce primarily for the local market, there are about 30 firms in the industry. However, production and sales are dominated by the 6 largest firms, each of which has production capacity of more than 100 tonnes per day—large by international standards. Besides skilled labor and capital, the key input for pasta manufacture is semolina, which is processed from Durham wheat. Pasta production from semolina is completely integrated with firms producing both semolina and pasta within the same plant.

Currently, there is significant, growing excess capacity within the industry. Per capita consumption in Turkey was, until 1996, almost stagnant at about 5 kg per annum. Domestic demand has increased slowly during the last decade—far slower than the observed production increases—with the main driving force being, apparently, population (population in Turkey has grown in the order of 1.7 % per annum in the last decade). The loss of the U.S. market and the slow growth of domestic demand has, however, led to large and growing idle capacity and to increasing surplus stocks.

Against these recessionary trends, the Turkish Pasta Producers Association organised during 1996-1998 a series of generic promotion campaigns in order to stimulate domestic demand. The objective of the campaigns was quantified explicitly. The aim was to stimulate a demand expansion by the end of 2000 that would raise per capita production to 8 kg per annum—a 60 % increase over a five-year period.

The monthly data on pasta consumption span the period 1993:5 – 1998:2, including the three months (1996:3, 1997:3 and 1997:10) of the campaigns, which were undertaken in the major city centers—Istanbul, Ankara and Izmir. The first campaign was undertaken in March 1996, the second in March 1997, and the third in October 1997. Expenditures in nominal US dollars for each of the campaigns were $663,338, $126,000 and $559,232, respectively. These investments are small in comparison with mandatory schemes that operate in North America and Australia, but there is one significant difference between those schemes and the present one—the promotion campaigns in Turkey were funded entirely voluntarily, with major contributions from the six largest companies. Another significant difference between the current scheme and others concerns the timing of promotion expenditures. Generally speaking, conventional analyses consider time series within which consecutive expenditures occur; here we have but three, "one-shot" campaigns.[2] It follows that the expenditure amounts and their timing provide an interesting contrast to typical analyses. In addition, each campaign made use of multiple media, with widely varied compositions. In the March '96 campaign 65% of total expenditure was allocated to television commercials, 16% was allocated to newspaper advertisements, 15% was allocated to billboard advertisements, 2% was allocated to journal advertisements, and the remaining 2% of the budget was absorbed by other costs of program administration. In the March '97 campaign the entire promotion budget was allocated to television advertisements. In the October '97 campaign, the proportions of the budget allocated to television, billboard, newspaper and journal advertisements were 35%, 14%, 12% and 0%, respectively. In addition

---

[2] Of course, an interesting question, which we leave for later analysis, is whether it would have been more profitable to spread the expenditures in each month over a series of months or, perhaps, a year.

the October '97 campaign allocated 33% of budget to program sponsorship and 3% of budget was absorbed by other costs. These variations provide an additional point of reference for interpreting the results, comparing the different investment portfolios and identifying future investment strategies that may raise consumption levels in the near future.

Casual empirical assessments, and the perception among producers, particularly those in the larger companies, suggest that the campaigns have been 'a success.' Nevertheless, precise measures of response are unavailable, leaving considerable scope for empirical inquiry. In order to shed light on this important question we focus attentions on carryover rates across the three campaigns, their longevity, and their magnitude.

6. Empirical Evidence

In modelling pasta consumption over the sample period, we focus on a single reduced-form equation containing each of the essential factors affecting consumption in Turkey. Pasta is something of a 'new' product for many Turkish consumers, many of whom remain unresponsive to price adjustments, furnishing their basic subsistence with traditional rice and bread products. Nevertheless, there is a slight upward trend in consumption patterns over the sample period, with mean percent changes in consumption (1993:5-1998:2) around one half of one percent. Of interest for our analysis is that this trend was somewhat stronger in the pre-promotion period (1993:5-1996:2) than the post-promotion time series (1996:3-1998:2), with growth rates of 0.67 and 0.51 of one percent, respectively. It is whether this decline in trend in the post-promotion period would have been higher in the absence of the program that we wish, in essence, to investigate.

Because Turkish pasta consumers' consumption habits are highly seasonal, conjectured to be income inelastic, but possibly responsive to own and competing product prices, we adopt a fairly parsimonious specification consisting of twelve monthly dummy variables and the real prices of pasta, rice and bread as covariates in (3). The dummy variables are constructed in the usual way with ones corresponding to the month in question and zeros elsewhere. The pasta, rice and bread prices are national averages[3] at retail (Turkish lira per kilogram) normalized by the consumer price index (1993:5 = 1.0). The dependent variable in the regression is retail disappearance (in thousands of metric tons) and the expenditure data for the three campaigns is in millions of nominal US dollars. Given these denominations, the lagged response rates that we seek to estimate measure returns in metric tons of additional consumption per thousand dollars expended.

Although the generalization to Turkish consumers may be wholly inaccurate, the authors usually consume bread with their pasta and sometimes choose, instead, a rice dish when a mood for pasta arises, leading to *a priori* expectations that pasta and rice will be substitutes and that pasta and bread will be complements. One issue arising in the pasta, rice and bread series is an unfortunate, highly collinear relationship between pasta- and rice-price movements. The sample correlation coefficient between nominal rice and pasta prices is 0.99. This collinearity poses problems for estimating precisely the individual responsiveness of the consumption data to own

---

[3] National averages for Turkey, as in this case, typically exclude the southern eastern region of the country where political unrest and considerable military activity prohibit data collection.

price movements and to movements in the price of rice. Such imprecision could also bias estimates of the lagged responses to the promotion campaigns. For this reason, we investigated the lagged responses for three variations of the model, namely a specification that includes the twelve quarterly dummy variables and the real prices of pasta, rice and bread (Model I); another that includes the twelve dummies, the real price of pasta and the real price of bread (Model II); and a third specification that includes the twelve dummy variables and the real prices of rice and bread (Model III). Although slight numerical differences in the lagged responses exist, each of the qualitative impacts is the same and, in the interests of focusing on one 'preferred specification,' we confine attentions to Model I.[4]

Table 3 reports estimates of posterior means of the parameters in the regressions with the missing data unrestricted in sign (space limitations prevent the reports of the full posterior densities). The first column reports parameters; the second column presents definitions; and the third column reports estimates of the parameters for the model that includes the prices of pasta, rice and bread as covariates (Model I). The numbers in parentheses are the reported means divided by their standard deviations in the Gibbs sample.

Table 3. Unrestricted estimates of posterior means of lagged responses in pasta consumption (1993:5 – 1998:2) to promotion campaigns (1996:3, 1997:3 and 1997:10).

Parameter	Description	Model I
$\hat{\beta}_1$	*January dummy*	11.78 (6.98)
$\hat{\beta}_2$	*February dummy*	11.94 (7.05)
$\hat{\beta}_3$	*March dummy*	12.11 (6.99)
$\hat{\beta}_4$	*April dummy*	11.20 (6.11)
$\hat{\beta}_5$	*May dummy*	10.22 (5.80)
$\hat{\beta}_6$	*June dummy*	10.09 (6.25)

---

[4] In the second model (rice price excluded) the response to own price change is negative (though insignificant); and in the third model (pasta price excluded) the response to change in the real price of rice is negative and significant. Across each of the three specifications the response to bread-price is insignificant, but appears to be stable to the exclusion (inclusion) of the other price series. Across the three regressions, point estimates of the standard error are similar (0.50, 0.56, 0.49, respectively across the three models) and estimates of the carryover effects are fairly similar, implying that the carryover effects are invariant to the alternative price specifications.

Table 3 (continued)

$\hat{\beta}_7$	*July dummy*	10.28
		(6.47)
$\hat{\beta}_8$	*August dummy*	10.02
		(6.19)
$\hat{\beta}_9$	*September dummy*	10.98
		(6.61)
$\hat{\beta}_{10}$	*October dummy*	11.18
		(7.05)
$\hat{\beta}_{11}$	*November dummy*	11.42
		(7.43)
$\hat{\beta}_{12}$	*December dummy*	11.37
		(6.96)
$\hat{\beta}_{13}$	*Retail Pasta Price*	0.0001
		(0.7870)
$\hat{\beta}_{14}$	*Retail Rice Price*	-0.0002
		(-1.7518)
$\hat{\beta}_{15}$	*Retail Bread Price*	0.0001
		(0.3156)
$\hat{\sigma}$	*Standard Error*	0.50
		(2.87)
α_{35}	*Lag 0*	-1.79
		(-1.23)
α_{36}	*Lag 1*	0.53
		(0.35)
α_{37}	*Lag 2*	2.18
		(1.60)
α_{38}	*Lag 3*	1.20
		(0.95)
α_{39}	*Lag 4*	0.41
		(0.32)
α_{40}	*Lag 5*	1.85
		(1.50)
α_{41}	*Lag 6*	0.81
		(0.63)
α_{42}	*Lag 7*	1.43
		(1.15)
α_{43}	*Lag 8*	-0.09
		(-0.07)
α_{44}	*Lag 9*	-0.36
		(0.29)
α_{45}	*Lag 10*	-0.69
		(-0.54)
α_{46}	*Lag 11*	-0.49
		(0.38)
	March 97 Campaign	
α_{47}	*Lag 0*	-3.32
		(-0.46)
α_{48}	*Lag 1*	2.69
		(0.34)
α_{49}	*Lag 2*	12.17
		(1.64)
α_{50}	*Lag 3*	12.30
		(1.80)
α_{51}	*Lag 4*	12.30
		(1.80)

Table 3 (Continued)

α_{52}	*Lag 5*	12.30
		(1.80)
α_{53}	*Lag 6*	12.30
		(1.80)
	October '97 Campaign	
α_{54}	*Lag 7*	12.30
		(1.80)
α_{55}	*Lag 8*	12.30
		(1.80)
α_{56}	*Lag 9*	12.30
		(1.80)
α_{57}	*Lag 10*	12.30
		(1.80)
α_{58}	*Lag 8*	12.30
		(1.80)
R^2	*Coefficient of Determination*	0.23
κ_{35}	*March '96 Expenditure ($US million)*	0.66
κ_{47}	*March '97 Expenditure ($US million)*	0.13
κ_{54}	*October '97 Expenditure ($US million)*	0.56
Condition($\mathbf{x'x}$)	*Condition number*	995.06
n	*Sample size*	58
S^1	*Burn-In sample size*	5000
S^2	*Gibbs-sample size*	10000
Execution Time	*Execution time (minutes)*	7.22

Note: Numbers in parentheses are parameter estimates divided by their standard deviations in the Gibbs sample.

Monthly dummy variables are significant but very similar, suggesting lack of evidence of significant seasonal effects in pasta consumption. Consumption responds positively to the real price of bread and responds positively (though insignificantly) to own price and negatively (though insignificantly) to the price of rice. Although the apparent insignificance of the reports for rice and bread prices belies firm conclusions, the results indicate (somewhat inexplicably and contrary to prior belief) that rice is a complement and bread a substitute for pasta. Notwithstanding the latter comment, the remaining patterns of response are the ones that we would expect given the high collinearity between rice and pasta prices. While they prevent firm conclusions about the relative responsiveness of pasta demand to pasta- and rice-price changes, they do not appear to affect markedly the stability of the other parameters in the model.

Turning to the lag effects, each of the negative point estimates (occurring at observations 35, and 43-47) is insignificant at conventional significance levels. The pattern of appearance of negative impacts is noteworthy for one other reason, namely the timing of the three campaign expenditures. The latter, it should be noted, occurred at observation numbers 35, 47 and 54. With this knowledge the results seem to suggest that promotion may have had important impacts beginning immediately after disbursement of the first two campaigns. These are reasonable conclusions to draw from the exercise and are ones that probably would be drawn under more conventional estimation methods.

Turning to comparisons of the three campaigns, the second expenditure (March '97) stands out as the most influential, at least in terms of maximizing response rates[5]. Rates increased significantly during the second month in the post-promotion period, and peaked at a rate of about +21 in the fifth post-promotion period. Response rates in the other periods were mostly insignificant.

Figure 1 shows, in different terms, the pattern of response rates. It serves also to highlight the proximity of the rates derived from the separate specifications of the model and illustrates, again, that the carryover estimates are not affected markedly by the alternative specifications. More importantly, the graphic suggests in fairly clear terms that post-promotion responses in the Turkish pasta sector were generally, significant; that the first campaign was moderately successful; that the second campaign had a significant impact that subsided quite rapidly in the sixth and seventh months that followed it; and that the third campaign was successful in stemming the rapidity of this decline. Conclusions beyond these ones are prohibited due to data limitations.

Although Benson, Breidt and Schroeter (2002) call into question the practice of restricting carry-overs to be positive, we reinvestigate the regression in table 3 under the assumption that the draws for the latent responses in Step 3 in the algorithm in (5) are truncated normal and are positive. Reports of the estimates are suppressed in the interests of space but can be summarized concisely in comparison with the reports in table 3 and figure 1. In fact, the parameter estimates are little altered from those in table 3 and the figure. That is, the carryover response rates exemplify the same pattern across the time series and are virtually indistinguishable across the three models. Collectively, the results suggest that the sign restriction on the carryover impacts is a relatively weak imposition on the data.

Finally, it must be mentioned, the campaigns, while they appear to have had a short-term impact on consumption, do not appear to have met their explicit target of stimulating per capita demand to 8 kilograms per annum—current projections estimate per capita consumption at around 5.8 kilograms per annum.

---

[5] A (more conventional) money-metric rate of return using producer surplus calculations can be obtained from these results by making assumptions about the retail supply elasticity for Turkish pasta. In the absence of reliable assumptions, we present the rate of return measured in units of consumption quantity per unit expenditure, which is the fundamental input into the latter calculation.

**Figure 1. Promotion carryover—unrestricted parameter estimates
(disbursement periods 1996:3=35, 1997:3=47, 1997:10=54).**

7. Conclusions

    Subject to some limitations which we outline below, we conclude that the promotion programs conducted in the three months March '96, March '97 and October '97—although they appear to have failed to meet their explicit target—were successful in raising short-term consumption levels and stemming declines in downward trends. The results suggest that responses were greatest to the second campaign, which allocated its entire budget to television media; that its impact peaked in the sixth month following expenditure; and that the rate of return (measured in metric tons additional consumption per thousand dollars expended) was around a factor of 20. Whether these responses have led to significant increases in the short-run profitability of the industry or altered general trends among competing products we leave unanswered. Future analysis should focus attentions on these issues—issues for which our missing-data approach to promotion carryover provides some important, basic information. In this regard, our principal objective has been to propose an alternative perspective to the conventional dummy-variables formulation for estimating flexible promotion carryover. We argue that the interpretation is, perhaps, most useful when expenditures do not occur consecutively within a time series, there is reasonably sound prior information that an impact occurred, and interest centers on the existence of lagged responses to once-and-for-all investments. But the approach, of course, is subject to some important limitations.

There are two principal limitations.   The first and most obvious limitation concerns miss-specification errors arising as a result of wrongly attributing the carryover responses to promotion when other missing data or omitted variables are responsible; although this source of bias is a familiar one that is inherent in more conventional methods, its effects are likely to be more dramatic in the present setting. Extensions of the theory and procedures should focus attention on minimizing the likelihood of this type of error.

The second limitation arises from an identification problem that precludes the estimation of multiple response rates when the lagged responses overlap within the time series. The problem can be observed in terms of equation (2). Because the parameters rather than the data vary in the post-promotion sample, only a single latent variable is identified from the complete (data $\times$ parameter) responses. For example, it appears (from table 3 and figure 2) that carryover from the second promotion campaign overlaps with the third. For this reason we do not advocate the use of the approach to more conventional data—and this point needs to be stressed. Whether the framework can be extended directly or, indirectly (perhaps by prompting more use of sampling-based approaches to commodity promotion) remains to be seen.

For one-shot campaigns, or when promotion data are intermittent in the time series, the present work suggests an alternative interpretation and approach to estimation that is intuitive; is flexible; and has the potential to be useful in situations where the length of the time series prohibits the estimation of a traditional distributed-lag—the present example is one such situation. At the least the theory and procedures proposed in this paper offer an alternative, perspective for estimating carryover rates and provide a useful starting point from which to consider additional, institutional detail.

Holloway is Reader in Agricultural Economics and Statistics, Department of Agricultural and Food Economics, School of Agriculture, Policy and Development, University of Reading, Earley Gate, Whiteknights House, Berkshire, RG6 6AR, United Kingdom, phone: (44)-(118)-931-6775; fax: (44)-(118)-975-64-67; email: garth.holloway@reading.ac.uk.

Aydogus is Professor and Chair, Ege University, 35040 Bornova, Izmir, Turkey. Email: aydogus@bornova.ege.edu.tr.

References

Benson, J., F. Breidt and J. Schroeter. (2002). "Television Advertising and Beef Demand: Bayesian Inference in a Random Effects Tobit Model." *Canadian Journal of Agricultural Economics* 50, 201-219.

Casella, G and E. George. (1992). "Explaining the Gibbs Sampler." *American Statistician* 46, 167-74.

Chib, S. (1992). "Bayes Inference in the Tobit Censored Regression Model." *Journal of Econometrics* 52, 79-99.

Chang, C. and H. M. Kaiser. (1998). "Determinants of Temporal Variations in Advertising Effectiveness." *National Institute for Commodity Promotion Research and Evaluation Discussion Paper* 98-02, Cornell University, Ithaca, New York.

Dorfman, J. (1996). "Modeling Multiple Adoption Decisions In A Joint Framework." *American Journal of Agricultural Economics* 78, 547-57.

———. (1997). *Bayesian Economics Through Numerical Methods: A Guide to Econometrics and Decision-Making With Prior Information*. New York: Springer.

Gelfand, A. and A. Smith. (1990). "Sampling-Based Approaches to Calculating Marginal Densities." *Journal of the American Statistical Association* 85, 972-985.

Liu, D. J. and O. D. Forker. (1988). "Generic Fluid Milk Advertising, Demand Expansion and Supply Response: The Case of New York City." *American Journal of Agricultural Economics* 70, 229-36.

———. (1990). "Optimal Control of Generic Fluid Milk Advertising Expenditures." *American Journal of Agricultural Economics* 72, 1047-55.

Kinnucan, H. W., H.-S. Chang and M. Venkateswaran. (1993). "Generic Advertising Wearout." *Review of Marketing and Agricultural Economics* 61, 401-15.

Reberte, J. C., H. M. Kaiser, J. E. Lenz and O. D. Forker. (1996). "Generic Advertising Wearout: The Case of the New York City Fluid Milk Campaign." *Journal of Agricultural and Resource Economics* 21, 199-209.

Pritchett, J. G., D. J. Liu and H. M. Kaiser. (1998). "Optimal Choice of Generic Milk Advertising Expenditures by Media Outlet." *Journal of Agricultural and Resource Economics* 23, 155-69.

Tanner, M. H. and W. H. Wong. (1987). "The Calculation of Posterior Distributions By Data Augmentation." *Journal of the American Statistical Association* 82, 528-50.

Vande Kamp, P. R. and H. M. Kaiser. (1999). "Irreversibility In Advertising-Demand Response Functions: An Application to Milk." *American Journal of Agricultural Economics* 81, 385-96.

Ward, R. and L. Myers. (1979). "Advertising Effectiveness and Coefficient Variation Over Time." *Agricultural Economics Research* 31, 1-11.

Zellner, A. (1971). *An Introduction to Bayesian Inference in Econometrics.* New York: Wiley, Wiley Classics Library Edition.