European Journal of Operational Research 246 (2015) 232–241

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Accommodating heterogeneity *and* nonlinearity in price effects for predicting brand sales and profits

Stefan Lang[a,1], Winfried J. Steiner[b,*], Anett Weber[b,2], Peter Wechselberger[c,3]

[a] Department of Statistics, Faculty of Economics and Statistics, University of Innsbruck, Universitätsstrasse 15, A-6020 Innsbruck, Austria
[b] Department of Marketing, Clausthal University of Technology, Julius-Albert-Strasse 2, 38678 Clausthal-Zellerfeld, Germany
[c] RSU Rating Service Unit GmbH & Co. KG, Karlstraβe 35, 80333 München, Germany

ARTICLE INFO

Article history:
Received 20 March 2014
Accepted 24 February 2015
Available online 5 March 2015

Keywords:
Forecasting
Sales response modeling
Heterogeneity
Functional flexibility
Expected profits

ABSTRACT

We propose a hierarchical Bayesian semiparametric approach to account simultaneously for heterogeneity *and* functional flexibility in store sales models. To estimate own- and cross-price response flexibly, a Bayesian version of P-splines is used. Heterogeneity across stores is accommodated by embedding the semiparametric model into a hierarchical Bayesian framework that yields store-specific own- and cross-price response curves. More specifically, we propose multiplicative store-specific random effects that scale the nonlinear price curves while their overall shape is preserved. Estimation is fully Bayesian and based on novel MCMC techniques. In an empirical study, we demonstrate a higher predictive performance of our new flexible heterogeneous model over competing models that capture heterogeneity or functional flexibility only (or neither of them) for nearly all brands analyzed. In particular, allowing for heterogeneity in addition to functional flexibility can improve the predictive performance of a store sales model considerably, while incorporating heterogeneity alone only moderately improved or even decreased predictive validity. Taking into account model uncertainty, we show that the proposed model leads to higher expected profits as well as to materially different pricing recommendations.

1. Motivation and literature review

In recent years, two streams of research for estimating sales response models based on store-level data have evolved: on the one hand, researchers have proposed hierarchical Bayesian (HB) store sales models allowing for heterogeneity of marketing effects across stores (e.g., Andrews, Currim, Leeflang, & Lim, 2008; Blattberg & George, 1991; Boatwright, McCulloch, & Rossi, 1999; Hruschka, 2006b; Montgomery, 1997; Montgomery & Rossi, 1999). While some of these studies have shown that considering heterogeneity can improve model fit, the accuracy of sales forecasts, or expected profits (e.g., Hruschka, 2006b; Montgomery, 1997), recent research of Andrews et al. (2008) has demonstrated rather marginal improvements in fit and predictive performance from incorporating store heterogeneity. One possible reason for this latter finding is that the HB models mentioned above assume a strictly parametric functional form thereby limiting the scope for model calibration to an a priori fixed parametrization. Hence, although accounting for heterogeneity, a source of bias remains if the assumed parametric form differs from the true underlying function.

On the other hand, researchers have proposed nonparametric regression models in order to accommodate potential nonlinearities in store sales response (e.g., Brezger & Steiner, 2008; Haupt, Kagerer, & Steiner, 2014; van Heerde, Leeflang, & Wittink, 2001; Kalyanam & Shively, 1998; Steiner, Brezger, & Belitz, 2007). The empirical results of this second stream indicate that own- and cross-price effects may show complex nonlinearities which are difficult or not at all to capture by parametric models. The main potential weakness of this second group of nonparametric approaches, however, is that heterogeneity across stores has not been considered. Consequently, bias due to potential heterogeneity across stores here remains.

There is so far only one approach that has consolidated the two streams: Hruschka (2006a, 2007) proposed a hierarchical Bayesian multilayer perceptron (MLP) that allows for nonlinearity in price effects *and* yields store-specific coefficients. In an empirical study, his flexible heterogenous MLP turned out to be superior in terms of

* Corresponding author. Tel.: +49 5323 72 7650; fax: +49 5323 72 7659.
E-mail addresses: stefan.lang@uibk.ac.at (S. Lang), winfried.steiner@tu-clausthal.de (W.J. Steiner), anett.weber@tu-clausthal.de (A. Weber), peter.wechselberger@rsu-rating.de (P. Wechselberger).
[1] Tel.: +43 512 507 7110; fax: +43 512 507 2851.
[2] Tel.: +49 5323 72 7658.
[3] Tel.: +49 89 4423400 39.

http://dx.doi.org/10.1016/j.ejor.2015.02.047

Table 1
Descriptive statistics for weekly brand prices, market shares, and unit sales.

Refrigerated orange juice category (64 oz)[a]								
Brand	Retail price			Market share			Unit sales	
	Range ($)	Mean ($)	SD ($)	Range (percent)	Mean (percent)	SD (percent)	Minimum	Maximum
Premium brands								
Tropicana Pure[b]	[1.60; 3.55]	2.95	.53	[3;73]	15	15	6388	100,712
Florida Natural	[1.57; 3.16]	2.86	.33	[1;53]	5	7	1138	56,037
National brands								
Citrus Hill	[1.09; 2.82]	2.31	.31	[1;78]	8	12	2006	151,570
Minute Maid	[1.29; 2.92]	2.23	.40	[3;87]	21	22	4805	243,711
Tropicana	[1.49; 2.75]	2.20	.35	[2;75]	21	23	3041	102,629
Florida Gold	[.99; 2.83]	2.17	.39	[1;63]	4	8	325	150,945
Tree Fresh	[1.07; 2.48]	2.15	.27	[1;42]	4	6	916	39,401
Store brand								
Dominick's	[.99; 2.47]	1.75	.4	[1;83]	22	22	2170	189,462

[a] The unit sales of all eight brands amount to 96.25 percent of the total sales volume in the refrigerated orange juice category (64 oz) during the time span considered.
[b] Reading example: For Tropicana Pure the lowest observed price across all stores and weeks was 1.60 $, its lowest market share (unit sales) in a week pooled across stores was 3 percent (6388 units), and its mean price level averaged across all stores and weeks was 2.95 $.

posterior model probability (Hruschka, 2006a) and further with respect to predictive validity (Hruschka, 2007) compared to a heterogeneous parametric multiplicative model, respectively. In addition, Hruschka (2007) analyzed profit implications for the flexible heterogeneous MLP from a retailer's point of view. Specifically, he shows that taking menu costs into account a moderately risk averse retailer may prefer a clusterwise pricing strategy to a store-specific pricing strategy. If menu costs are ignored, expected profits however increase with the number of clusters and reach their maximum for a store-specific pricing strategy.[3] Our approach proposed here differs from that of Hruschka (2006a, 2007) in several ways: First, we will show that accounting for store heterogeneity alone might not be advantageous per se (as is assumed by Hruschka), and that accounting for functional flexibility is the primary driver for model improvement (at least for our data). In particular, we find that allowing for heterogeneity in addition to functional flexibility can improve the predictive performance of a store sales model considerably, while incorporating heterogeneity alone only moderately improves or even decreases predictive validity. Second, we illustrate why accommodating store heterogeneity may pay off only once nonlinearity in price response is modeled appropriately. And third, we compare our flexible heterogeneous model in terms of expected profits to competing models that capture heterogeneity or functional flexibility only (or neither of them). This way, we investigate how much loss in expected profits management incurs by not using the model with the highest predictive performance.

In the following, we develop a store sales model which accommodates *both* functional flexibility *and* heterogeneity within a unified regression framework. We propose a structured additive approach where own- and cross-price response is estimated flexibly using P-splines, while heterogeneity in price response across stores is simultaneously accommodated by multiplicative store-specific random effects that scale the nonlinear price curves while their overall shape is preserved.

The rest of the paper is organized as follows. In Section 2, we introduce our hierarchical Bayesian semiparametric model. Using store-level data from Dominick's Finer Foods we analyze in Section 3 whether, how much, and why the performance of a store sales model can be improved from considering either heterogeneity, functional flexibility, or both features. Our results indicate that the most complex model accommodating both heterogeneity and functional flexibility outperforms competing models in predictive validity for most brands, and provides substantial increases in expected profits for those brands. We further show that the proposed model leads to materially different pricing implications. We conclude in Section 4 with an outlook on future research perspectives.

2. Data and model framework

2.1. Data

To model store sales response, we use weekly store-level scanner data for eight brands of orange juice offered by Dominick's Finer Foods (DFF), a major supermarket chain in the Chicago metropolitan area. The data were provided by the James M. Kilts Center, GSB, University of Chicago, and include unit sales (Q_{st}), price ($price_{st}$) and a deal code indicating the use of a display ($display_{st}$) for each of the eight brands in each of $s = 1, \ldots, 81$ stores of the chain over a time horizon of $t = 1, \ldots, 89$ weeks (resulting in about 7,000 data points per brand). The price data further reveal whether a 9- or 99-ending price ($end9_{st}$, $end99_{st}$) has been set. The brands can be grouped into three price-quality tiers: the premium brands which are made from freshly squeezed oranges, the national brands which are reconstituted from frozen orange juice concentrate, and the retailer's own store brand. Table 1 provides summary statistics pooled across stores for weekly prices, market shares and unit sales of the individual brands.

Since cross-item price effects are usually much weaker than own-item price effects (e.g., see Hanssens, Parsons, & Schultz, 2003), we capture them in our demand equation in a more parsimonious way at the tier level. Following Brezger and Steiner (2008), we define $price\_national_{st}$ ($price\_premium_{st}$) as the lowest price of a competing national (premium) brand in store s and week t. For example, if a store sales model is estimated for the national brand Citrus Hill, $price\_national_{st}$ captures the lowest price level of either of the *other* national brands in store s and week t.[4] $price\_dominicks_{st}$ denotes the observed price for Dominick's, the only private label brand in our data.

We further use 11 characteristics (collected in the vector $\boldsymbol{v}_s$) of each store's trading area to explain possible store-level variation in price response due to sociodemographic and competitive effects. A detailed description of these background covariates (among others relating to age, education, family size, income, distances to nearest

[3] The advantages of using nonparametric or seminonparametric techniques for estimating response functions were demonstrated previously by Hruschka in the context of market share modeling (Hruschka, 2002), brand choice modeling (Hruschka, Fettes, & Probst, 2004) and catalog allocation modeling (Baumgartner & Hruschka, 2005), too.

[4] Note that for the computation of $price\_national_{st}$, we consider Sunny Delight as another national brand in the refrigerated orange juice category. However, we did not estimate store sales models for Sunny Delight due to the lack of (substantial) price variation of this brand.