Functional Regression: A New Model and Approach for Predicting Market Penetration of New Products

Ashish Sood, Gareth M. James & Gerard J. Tellis

PLEASE DO NOT DISTRIBUTE WITHOUT THE PERMISSION OF THE AUTHORS

Ashish Sood is Assistant Professor Marketing, Goizueta School of Business, Emory University, 1300 Clifton Rd NE, Atlanta, GA 30322; Tel: +1.404-727-4226, fax: +1.404-727-3552, E-mail: ashish_sood@bus.emory.edu.

Gareth James is Associate Professor of Statistics, Marshall School of Business, University of Southern California, P.O. Box 90089-0809, Los Angeles, California, USA; Tel: +1.213.740.9696, fax: +1.213.740.7313 E-mail: gareth@usc.edu

Gerard J. Tellis is Neely Chair of American Enterprise, Director of the Center for Global Innovation and Professor of Marketing at the Marshall School of Business, University of Southern California, P.O. Box 90089-1421, Los Angeles, California, USA. Tel: +1.213.740.5031, fax: +1.213.740.7828, E-mail: tellis@usc.edu

Functional Regression: A New Model and Approach for Predicting Market Penetration of New Products

Abstract

The Bass model has been the standard for predicting the market penetration of new products. Recently a new class of non-parametric techniques, Functional Data Analysis (FDA), has shown impressive results within the statistics community. The authors test the predictive performance of FDA versus variations of the Bass model on 760 new categories across numerous products and countries. Ten fold cross validation shows that Augmented Functional Regression provides distinctly superior predictions by using a) four underlying dimensions of the pattern of penetration across 760 categories, b) deviations of a target category from the mean on these four dimensions, and c) knowledge of the target product.

1

Introduction

Firms are introducing new products at an increasingly rapid rate. At the same time, the

globalization of markets has increased the speed at which new products diffuse across countries,

mature, and die off (Chandrasekaran and Tellis 2007). These two forces have increased the

importance of the accurate prediction of the market penetration of an evolving new product.

While research on modeling sales of new products in marketing has been quite insightful (Peres,

Mueller and Mahajan 2007), it is limited in a few respects. First, most studies rely primarily, if

not exclusively, on the Bass model. Second, prior research, especially those based on the Bass

model, need data past the peak sales or penetration for stable estimates and meaningful

predictions. Third, prior research has not indicated how the wealth of accumulated penetration

histories across countries and categories can be best integrated for good prediction of penetration

of an evolving new product. For example, a vital unanswered question is whether a new

product's penetration can be best predicted from past penetration of a) similar products in the

same country, b) the same product in similar countries, c) the same product itself in the same

country, or d) some combination of these three histories.

The current study attempts to address these limitations. In particular, it makes four

contributions to the literature. First, we illustrate the advantages of using Functional Data

Analysis (FDA) techniques for the analysis of penetration curves (Ramsay and Silverman, 2005).

Second, we demonstrate how information about the historical evolution of new products in other

categories and countries can be integrated to predict the evolution of penetration of a new

product. Third, we compare the predictive performance of the Bass model versus an FDA

approach, and some naïve models. Fourth, we indicate whether information about prior

countries, other categories, the target product itself, or a combination of all three is most important in predicting the penetration of an evolving new product.

One unique aspect of the study is that it uses data about market penetration from most of 21 products across 70 countries, for a total of 760 categories (product x country combinations). The data includes both developed and developing countries from Europe, Asia, Africa, Australasia, and North and South America. In scope, this study vastly exceeds the sample used in prior studies (see Table 1). Yet the approach achieves our goals in a computationally efficient and substantively instructive method.

Another unique aspect of the study is that it uses Functional Data Analysis to analyze these data. Over the last decade FDA has become a very important emerging field in statistics, although it is not well known in the marketing literature. The central paradigm of FDA is to treat each function or curve as the unit of observation. We apply the FDA approach by treating the yearly cumulative penetration data of each category as 760 curves or functions. By taking this approach we can extend several standard statistical methods for use on the curves themselves. For instance, we use functional principal components analysis (PCA) to identify the patterns of shapes in the penetration curves. Doing so enables a meaningful understanding of the variations among the curves. An additional benefit of the principal component analysis is that it provides a parsimonious, finite dimensional representation for each curve. In turn this allows us to perform functional clustering by grouping the curves into clusters with similar patterns of evolution in penetration. The groups that we form show strong clustering among certain products and provide further insights into the patterns of evolution in penetration. Finally, we perform functional regression by treating the functional principal components as the independent variable and future characteristics of the curves, such as future penetration or time to takeoff, as the dependent

variable. We show that this approach to prediction is more accurate than the traditional approach of using information from only one curve. It also provides a deeper understanding of the evolutions of the penetration curves.

The rest of the paper is organized as follows: The next three sections present the method, data and results. The last section discusses the limitations and implications of the research.

Method

We present the method in six sections. The first section describes the spline regression approach to modeling of individual curves. The next three sections outline various applications of functional data analysis. The second section describes functional principal components. The third section illustrates how the functional principal component scores can be used to perform functional cluster analysis and hence identify groupings among curves. The fourth section shows how the functional PCA scores can be used to perform functional regression for predictions. The fifth section describes the alternate models against which we test the predictive performance of the FDA models. The last section details the method used for carrying out predictions.

Modeling of Individual Curves

Functional data analysis is a collection of techniques in statistics for the analysis of curves or functions. Most FDA techniques assume that the curves have been observed at all time points but in practice this is rarely the case. If multiple and frequent observations are available for each curve, a simple smoothing spline can generate a continuous smooth curve from discrete, observations. For example, a smoothing spline is a curve plotting the penetration of CD Players, given ten discrete years of data.

Suppose that a curve, $X(t)$, has been measured at times $t=1, 2,...,T$. Then the smoothing spline estimate is defined as the function, $h(t)$, that minimizes

4

$$\sum_{t=1}^{T} \left(X(t) - h(t) \right)^2 + \lambda \int \{h''(s)\}^2 \, ds$$

(1)

for a given value of λ>0 (Hastie et al., 2001). The first squared error term in Equation (1) forces h(t) to provide an accurate fit to the observed data while the second integrated second derivative term penalizes curvature in h(t). The tuning parameter λ determines the relative importance of the two components in the fitting procedure. Large values of λ force a h(t) to be chosen such that the second derivative is close to zero. Hence as λ gets larger h(t) becomes closer to a straight line, which minimizes the second derivative at zero. Smaller values of λ place more emphasis on h(t)'s that minimize the squared error term and hence produce more flexible estimates. We follow the standard practice of choosing λ as the value that provides the smallest cross-validated residual sum of squared errors (Hastie et al., 2001). Remarkably, even though Equation (1) is minimized over all smooth functions it has been shown that its solution is uniquely given by a finite dimensional natural cubic spline (Green and Silverman, 1994), which allows the smoothing spline to be easily computed. A cubic spline is formed by dividing the time period into L regions where larger values of L generate a more flexible spline. Within the l^{th} region a cubic polynomial of the form

$$h(t) = a_l + b_l t + c_l t^2 + d_l t^3$$

(2)

is fit to the data. Different coefficients, a_l, b_l, c_l and d_l are used for each region subject to the constraints that h(t) must be continuous at the boundary points of the regions and also have continuous first and second derivatives. In a natural cubic spline, the second derivative of each polynomial is also set to zero at the endpoints of the time period. In the more complicated situation where the curves are sparsely observed over time (e.g. due to a different data generating process or data limitations), a number of alternatives have been proposed. For example, James et

al (2000) suggest a random effects approach when computing curves for use in a Functional Principal Components setting.

Functional Principal Components

Suppose we observe n smooth curves, $X_1(t)$, $X_2(t)$, …, $X_n(t)$ (e.g. either observed empirically, or as in our case, approximated from the observed discrete data using the spline regression approach). Then we can always decompose these curves in the form,

$$X_i(t) = \mu(t) + \sum_{j=1}^{\infty} e_{ij}\phi_j(t) \qquad i = 1,..., n$$

…(3)

subject to the constraints

$$\int \phi_j^2(s)ds = 1 \quad \text{and} \quad \int \phi_j(s)\phi_k(s)ds = 0 \quad \text{for } j \neq k.$$

The $\phi_j(t)$'s represent the principal component functions, the e_{ij}'s the principal component scores corresponding to the i^{th} curve and $\mu(t)$ the average curve over the entire population. As with standard principal components, $\phi_1(t)$ represents the direction of greatest variability in the curves about their mean. $\phi_2(t)$ represents the direction with next greatest variability subject to an orthogonality constraint with $\phi_1(t)$. The e_{ij}'s represent the amount that $X_i(t)$ varies in the direction defined by ϕ_j. Hence a score of zero indicates that the shape of $X_i(t)$ is not similar to ϕ_j while a large score suggests that a high fraction of $X_i(t)$'s shape is generated from ϕ_j.

To compute the functional principal components we first estimate the entire path for each $X_i(t)$ using a smoothing spline. Next we divide the time period t=1 to t=T into p equally spaced points and evaluate $X_i(t)$ at each of these time points. Note that the new time points are not restricted to be yearly observations because the smoothing spline estimate can be evaluated at any point in time. Finally, we perform standard PCA on this p dimensional data. The resulting principal component vectors provide accurate approximations to the $\phi_j(t)$'s at each of the p grid

points and likewise the principal component scores represent the e_{ij}'s. One generally chooses p to be a large number such as 50 or 100, to produce a dense grid over which to evaluate the $\varphi_j(t)$'s and hence generate a smooth estimate for $\varphi_j(t)$.

In addition to computing functional principal components on $X_i(t)$ one can also compute the principal components on its derivatives such as $X_i'(t)$ and $X_i''(t)$ which measure the velocity and acceleration of the curves. The velocity of a curve provides information about its rate of change over time i.e. its first derivative. Hence, a high velocity implies a rapid increase or decrease in $X_i(t)$ while a velocity close to zero suggests a stable curve. The acceleration measures the rate of change in the velocity. For example, a straight line has a constant velocity but an acceleration of zero because its first derivative is a constant and the second derivative is zero. The velocity and acceleration curves can be particularly useful because they provide additional information about the penetration curves $X_i(t)$ which are generally curvilinear (Foutz and Jank, 2007).

In theory, n different principal component curves are needed to perfectly represent all n $X_i(t)$'s. However, in practice a small number (D) of components usually explain a substantial proportion of the variability (Ramsay and Silverman, 2005) which indicates that

$$X_i(t) \approx \mu(t) + e_{i1}\varphi_1(t) + e_{i2}\varphi_2(t) + \cdots + e_{iD}\varphi_D(t) \qquad i = 1,\ldots,n \quad \ldots(4)$$

for some positive D << n.

Note that the smooth functions, $X_i(t)$, are infinite dimensional in nature but observed only at a finite number of time points. However, we can use e_{ij} in Equation (4) to reduce the infinite dimensional functional data to a small set of dimensions. This reduction in dimensions is crucial because it allows us to perform functional clustering and functional regression as described in the following two sections.

Functional Clustering

We use functional clustering not for prediction but to better understand the data. In particular, we wish to identify groups of similar curves and relate them to observed characteristics of these curves such as the product and country. We use the principal components described in the previous section to reduce the potentially large number of dimensions of variability and cluster all the curves in the sample.

We apply the standard k-means clustering approach (MacQueen 1967) to the D-dimensional principal component scores, e_i, described in Equation (4) to cluster all the curves in the sample. k-means clustering works by locating D-dimensional cluster centers $c_1,\ldots, c_k$ which minimize the sum of squared distances between each observation and its closest cluster center i.e. find $c_1,\ldots, c_k$ to minimize

$$\gamma_k = \sum_{i=1}^{n} \min_{c_1, c_2, \ldots, c_k} \left\| e_i - c_j \right\|^2$$

(5)

An iterative algorithm is used to minimize γ. First an initial set of candidate centers, $c_1,\ldots, c_k$, are chosen, usually by randomly selecting k of the e_i's, and each curve is assigned to its closest center. Then, for each cluster, a new center is defined by taking the average over all curves currently assigned to that cluster. This algorithm continues until additional iterations do not yield significant changes in the cluster centers.

We use the "jump" approach (Sugar and James 2003) to select the optimal number of clusters, k. In this approach, we compute $\xi_k = \gamma_k^{-Y} - \gamma_{k-1}^{-Y}$ for a range of values of k where γ_k is given by (5) and Y is usually taken to be D/2. Sugar and James (2003) show through the use of information theory and simulations that setting the number of clusters equal to the value corresponding to the largest ξ_k provides an accurate estimate of the true number of clusters in the data.

Once the cluster centers have been computed, each curve is assigned to its closest cluster mean curve. We can then use Equation (4) to project the centers back into the original curve space and hence examine the shape of a typical curve from each cluster.

Functional Regression

We now show how functional regression can be used to predict market penetration of a category by linking each item to be predicted (e.g. marginal penetration level in any year or the year of takeoff) to the functional principal components identified for each curve in the previous section. Let $X_i(t)$ represent the i^{th} curve observed over time such as the first five years of cumulative penetration for a given category. Let Y_i represent a related item to be predicted, such as the penetration in year six.

Functional regression establishes a relationship between predictor, $X_i(t)$, and the item to be predicted, Y_i, as follows:

$$Y_i = f(X_i(t)) + \varepsilon_i \quad i = 1,...,n. \qquad \dots (6)$$

Equation (6) is difficult to work with directly because $X_i(t)$ is infinite dimensional. However, it can be shown that for any function f there exists a corresponding function g such that $f(X(t)) = g(e_1, e_2, ...)$ where e_1, e_2, … are the principal component scores of X(t). This equivalence allows us to use the functional principal component scores to perform functional regression. The simplest choice for g would be a linear function in which case Equation (6) becomes

$$Y_i = \beta_0 + \sum_{j=1}^{D} e_{ij}\beta_j + \varepsilon_j \qquad \dots (7)$$

for some D≥ 1. A somewhat more powerful model is produced by assuming that g is an additive but non-linear, function (Hastie and Tibshirani, 1990). In this case Equation (6) becomes

$$Y_i = \beta_0 + \sum_{j=1}^{D} g_j(e_{ij}) + \varepsilon_j$$

... (8)

where the g_j's are non-linear functions that are estimated as part of the fitting procedure. There are different ways to model the g_j's but one common approach, which we use in this paper, is the smoothing spline discussed in the section on Modeling of Individual Curves. One advantage of using Equations (7) or (8) to implement a functional regression is that once the e_{ij}'s have been computed via the functional PCA, we can then use standard linear or additive regression to relate Y_i to the principal component scores. We can also add covariates that contain information about the curves beyond the principal components, such as product or country characteristics or marketing variables, to the functional regression model in Equation (8).

Comparing Alternative Models

To fully understand the advantages of FDA, we compare two implementations or models of FDA with five non-Functional models. We name the two functional models Functional Regression and Augmented Functional Regression and the five non functional models - Estimated Mean, Last Observation Projection, Classic Bass, Meta Bass, and Augmented Meta Bass. Table 2 classifies all the models based on their use of information across curves and nature of the model. We divide the sample into estimation samples for estimating each model and holdout samples for testing the prediction of each model (see Figure 1), using ten fold cross-validation, explained later.

Estimated Mean

The Estimated Mean is a simple model, which fits the mean of the item to be predicted in the estimation sample, as the predicted value of the item in the holdout sample. Specifically, the prediction for the i[th] observation in the holdout sample, $\hat{Y}_i$, is given by

$$\hat{Y}_i = \overline{Y} \qquad \dots (9)$$

where $\overline{Y}$ is the mean on the estimation sample. Note that this is a very simple model which does not use any information in the first T periods of data of the curve in the holdout sample.

Last Observation Projection

The Last Observation Projection is another simple model, which estimates the item to be predicted from only the last observation in each penetration curve. To do so, we first relate the item to be predicted, Y_i, to the final observed penetration level, $X_i(T)$, in the estimation sample. To estimate this relationship, we explore both a standard linear model (Equation (10)) as well as a more flexible non-linear model (Equation (11)),

$$Y_i = \beta_0 + \beta_1 X_i(T) + \varepsilon_i, \qquad \dots (10)$$

$$Y_i = \beta_0 + g(X_i(T)) + \varepsilon_i \qquad \dots (11)$$

We use the non-linear model for our final results. For the prediction, we use the estimated g from Equation (11) and the final observed penetration level ($X_i(T)$) in the holdout sample to get the predicted item in the holdout sample.

Note, this is a slightly superior model to the Estimated Mean, because it uses at least the last observation from each curve to be predicted. However, it still does not use any other prior data from the curve.

Classic Bass

The Classic Bass Model fits the Bass (1969) model for each curve in the sample, thus:

$$s(t) = m\big[F(t) - F(t-1)\big] + \varepsilon(t), \qquad F(t) = \frac{1 - e^{-(p+q)t}}{1 + (q/p)e^{-(p+q)t}} \qquad \dots (12)$$

where t=time period, s(t)=marginal penetration at time t, p=coefficient of innovation and q=coefficient of imitation.

We estimate the model via the genetic algorithm because Venkatesan , Krishnan and Kumar (2004) provide convincing evidence that the genetic algorithm provides the best method for fitting the Bass model relative to all prior estimation methods. For each curve, we use the first T years of data to estimate the three Bass parameters, m, p and q. We then predict the next five years of penetration levels by plugging the estimated parameters back into the Bass model and evaluating at times T+1 through T+5. We predict the time of peak marginal penetration by using t=log(q/p)/(p+q) and the peak marginal penetration using s= m(p+q)2/4q. We do not predict time to takeoff with the Classic Bass Model. Note that the Classic Bass Model does not distinguish between holdout and estimation samples because each curve is fit individually without using information from other curves.

Meta-Bass

The Meta-Bass model extends the Classic Bass Model to use information across curves. To do so, we first estimate m, p and q for each curve using the genetic algorithm, as outlined above. Then, for each item to be predicted, we fit the non-linear additive model,

$$Y_i = \beta_0 + g_1(m_i) + g_2(p_i) + g_3(q_i) + \varepsilon_i \, , \qquad \qquad \dots (13)$$

to the estimation sample where g_1, g_2, and g_3 are smoothing splines as defined previously. We use the estimated parameters from this additive model from the estimation sample and the estimates of m, p, and q for each curve in the holdout sample to compute each item to be predicted for each of the holdout curves.

12

Augmented Meta-Bass

The Augmented Meta-Bass is the same non-linear additive model used for the Meta-Bass except that we add an indicator variable for each of the R products to which each curve belongs, thus:

$$Y_i = \beta_0 + g_1(m_i) + g_2(p_i) + g_3(q_i) + \sum_{r=1}^{R-1} \delta_r I_{ir} + \varepsilon_i \quad \ldots (14)$$

where $I_{ir} = 1$ if the ith curve belongs to product r and 0 otherwise and the δ_r's are regression coefficients that are estimated as part of the model fitting procedure. Note, that the Meta-Bass and Augmented Meta-Bass are extensions of the Classic Bass that make use of all of the information across curves, rather than just utilizing each curve individually. Since, using information across curves is an essential feature of functional regression, doing so puts the Meta Bass and the Augmented Bass on the same platform as the FDA models (see Table 2).

Functional Regression

For the Functional Regression model, we compute the first four principal component scores, two each on the penetration curves, $X_i(t)$, and on the velocity curves $X_i'(t)$. We then use these four scores as the independent variables in an additive regression model, as shown in Equation (8), on the estimation sample. We then use the estimated parameters of this equation and the data from the curves in the holdout sample, to compute the items to be predicted in the holdout sample.

Augmented Functional Regression

Our second functional approach enhances the power of Functional Regression by adding an indicator variable for each of the R products to which each curve belongs, as with the Augmented Meta-Bass model. Hence, the Augmented Functional Regression model involves estimating a non-linear additive model on the estimation sample as follows:

$$Y_i = \beta_0 + \sum_{j=1}^{4} g_j(e_{ij}) + \sum_{l=1}^{L-1} \delta_l I_{il} + \varepsilon_{ij} \qquad \dots (15)$$

where the g_j's are smoothing splines. We then compute the items to be predicted for each curve in the holdout sample from the estimated values of the above parameters and the data in each curve in the holdout sample. This model is directly comparable to Augmented Meta Bass as both models use information across curves and from products.

Method for Prediction

We explain the specific procedure for carrying out the prediction in three parts: items being predicted, computation of errors, and partitioning of sample.

Items Being Predicted

We first truncate each curve at the T^{th} year. We use the penetration in years 1 to T to estimate the model and predict the marginal change in penetration for years T+1 to T+5. For each curve, we also predict the number of years to takeoff, the years to peak marginal penetration, and the level of peak marginal penetration. Thus, we predict a total of eight items for each of seven models, for a total of 56 model-items. We do this whole process once each for T=5 years and T=10 years.

Computation of Errors

For each of these 56 model-items to be predicted, we compute the mean absolute deviation (MAD) over all penetration curves, i.e.

$$MAD = \frac{1}{N} \sum_{i=1}^{N} \left| \hat{Y}_i - Y_i \right| \qquad \dots (16)$$

where Y_i is a particular item for curve i and $\hat{Y}_i$ is the corresponding estimate using a given model.

Partitioning of Sample

We use ten-fold cross-validation by randomly partitioning the curves into ten equal groups. We hold out one group, estimate each of the models on the remaining nine groups and then form predictions on the held out group. We repeat this process ten times, for each of the ten held-out groups of data. For each of the 56 model-items, we compute the mean absolute deviation (MAD) as an average of these ten iterations. Note, that k-fold cross-validation is superior to simple splitting of data into one holdout and training group, because all of the data are used (randomly) as a holdout once.

Data

This section details our sample, sources, and procedure for data collection.

Sample

Most of the prior studies are limited in scope in terms of both product type and geographical breadth (see Table 1). We collect a large database of products (see Table 3a) introduced across a large set of countries (see Table 3b). The sample includes a broad sample from three categories - household white goods, computers and communication, and entertainment and lifestyle.

Sources

The information required for this study is penetration rates of different products introduced in different markets from the year of introduction to the year of first slowdown. The primary source of our data is the Global Market Information Database of Euromonitor International. Euromonitor International's Global Market Information Database is an integrated online information system that provides business intelligence on countries, consumers and

lifestyles. We also use press releases, industry reports and archived records to identify the year of introduction from databases like Factiva and Productscan.

Procedure

We follow the general rules for data collection for the historical method (Golder 2000). We explain specific problems we encounter and the rules we use to resolve them. We examine the data for each category for three things. First, we identify all curves that have penetration rates of more than 1% in the first year. For these categories, we check the year of introduction from historical reports or press releases. We exclude all categories where data is not available from the first year of introduction. Second, we exclude from our analysis any categories that do not contain at least T+5 years of observations or have not reached peak marginal penetration. Third, the data from this source is only available from 1977. Hence, we exclude all categories where the product had been introduced or taken off earlier than 1977.

Results

We present the results on functional principal components, functional clustering, and functional regression.

Functional Principal Components

Figures 2a and 2b provide plots of $\varphi_1(t)$ and $\varphi_2(t)$ computed from the first ten years of observations on the 760 penetration curves. The first principal component represents the amount by which a curve's penetration is above or below average by year ten. Categories with a positive score on the first component end up with above average last period penetration levels while those with negative scores have below average last period penetration. Alternatively, the second principal component represents the way that the penetration levels evolve. Categories with a positive score on the second component grow most rapidly in the early years but slow down by

year ten while those with a negative score are associated with slow initial growth and a rapid increase towards year ten.

An alternative way of visualizing these curves is presented in Figures 2c and 2d. Here the black line corresponds to $\mu(t)$, the average penetration level over all 760 curves. The red lines represent $\mu(t) \pm \varepsilon_j \varphi_j(t)$ where ε_j is a constant proportional to the standard deviation of e_{ij}. Figure 2c shows that categories with a positive value for e_{i1} will have above average last period penetration levels at year ten while ones with a negative e_{i1} will remain stagnant over time and last period penetration levels below the overall average. Alternatively, Figure 2d shows that categories with a positive value for e_{i2} will grow somewhat faster than average to begin with but then fall below average after 10 years while curves with a negative e_{i2} will have the opposite pattern.

Remarkably, $\varphi_1(t)$ and $\varphi_2(t)$ together explain over 99% of the variability in the smoothed penetration curves which indicates that e_{i1} and e_{i2} provide a highly accurate two-dimensional representation of $X_i(t)$. However, it should be noted that the smoothing spline approach removes some of the variability in the data so $\varphi_1(t)$ and $\varphi_2(t)$ explain somewhat less than 99% of the variation in the observed penetration data. As mentioned previously one can also compute principal components for the velocity of the penetration levels. When we perform this decomposition on the penetration curves the principal components of $X'_i(t)$ have a very similar structure to those for $X_i(t)$.

Functional Clustering

Figure 3 provides several approaches to viewing the results from the functional clustering using k-means on e_{i1} and e_{i2}. The "jump" approach of Sugar and James (2003) suggests between six and nine clusters. We opt for six to provide the most parsimonious representation (see Table

3). Figure 3a plots the centers of the six clusters on the original time domain. The figure illustrates the pattern of growth of a typical curve in each cluster. Alternatively, Figure 3b plots all 760 curves in the reduced two dimensional space, using the same colors to represent each cluster as for Figure 3a. The six cluster centers are represented as solid black circles.

Each cluster differs from the other clusters in the pattern of penetration over time. Broadly speaking, Clusters 1 through 3 represent high growth products while the last three correspond to lower growth rates. Cluster 1 takes on large values in both the first and second principal component dimensions. Recall that a positive value in the first dimension corresponds to overall high last period penetration while a positive value in the second dimension represents a fast growth at the beginning but a slow down by year 10. The black curve in Figure 3a shows this pattern with the fastest overall growth but a slight slowdown by year 10. Cluster 2 is close to zero for second dimension indicating no overall slowdown as we can see from the pink curve. Clusters 3 and 4 provide an interesting contrast. Cluster 3 has a negative value in the second dimension while cluster 4 is positive. This suggests a slow start for Cluster 3 but with increasing momentum by year 10 and the opposite pattern for Cluster 4. Figure 3a shows precisely this pattern with Cluster 4 starting ahead of Cluster 3 but then falling rapidly behind. Cluster 5 represents a moderate rate of growth while Cluster 6, which contains the largest number of products, corresponds to a much slower improvement in penetration.

We also examine whether the penetration patterns differ across products. Figure 3c illustrates the growth patterns for the twenty-one different products in the sample. We plot all 760 curves in our two-dimensional space using a different plotting symbol for each product. There are very clear patterns within the same product. For example, the green stars correspond to internet-compatible personal computers and have almost uniformly large values on the first

dimension indicating rapid increases in penetration levels. Notice that one product may have both positive and negative values for the second dimension, suggesting more rapid takeoff in some markets over others. Alternatively, the yellow squares represent DVD players and have a very tight clustering with almost uniformly moderate scores on the first principal component and negative scores on the second principal component. These results suggest a slow initial growth with much more rapid expansion towards year 10. The tighter clustering suggests that the takeoff for these products is largely similar across different markets in the sample. Finally, the blue solid dots, representing Video Tape Recorders, show the opposite pattern with large positive scores on the second dimension suggesting fast initial growth but then a slow down in later years.

Curves for each product are from a variety of countries. Table 4 provides the fraction of curves of each product that fall within each of the six clusters. The functional clustering suggests three groups - fast growth electronics, slower growth electronics, and household goods. The first three clusters consist entirely of the six fast growth electronics with Cluster 1 primarily internet personal computers, Cluster 2 a mixture and Cluster 3 mainly DVD players. Video game consoles, satellite TV, and CD players make up the bulk of Cluster 4. Cluster 5 contains many products but seems to principally concentrate on countries with slower growth for CD and DVD players, Satellite TV, and Video game consoles. Finally, Cluster 6, the slowest growth cluster, contains the vast bulk of household appliances.

Similarly, we also examine whether the penetration patterns differ across countries. We categorize the data into seven geographic regions (see Table 5). For each region, Table 5 shows the fraction of curves that fall in each of the six clusters. For example, for countries from Africa and developing Asia 86% of curves fall into the slowest growth Cluster 6. In contrast, North

American, Western Europe and Australasia have curves that are more spread out over the six clusters, with only 30% in the slowest growth, Cluster 6.

Functional Regression

This section presents the performance of the seven models on the eight items to be predicted. Tables 6a and 6b present the cross-validated mean absolute deviation scores for each model using cutoffs of T=5 and T=10 years of training data respectively.

In order to assess the superiority of functional data analysis for predicting items of penetration curves, we compare the Functional Regression model to the five non-functional models. Functional Regression is superior to Estimation Mean, Last Observation Projection, and Classic Bass at predicting all eight items at both cutoff times (see Table 6). The reason is that the Estimation Mean and the Last observation Projection use minimal information from prior time periods while Classic Bass uses no information across curves. Functional Regression is also better than Meta Bass on all items for both cutoff times except for time to peak marginal penetration at cutoff time = 10 years.

The performance of Functional Regression is mixed when compared with Augmented Meta Bass. At the Cutoff of T=5 years, Functional Regression is superior to Augmented Meta-Bass for the T+1, T+2, and T+3 years, similar for T+4 years but inferior for the other four items (see Table 6a and Table 7a). At the Cutoff of T=10 years, Functional Regression outperforms Augmented Meta-Bass for T+1 through T+5 years as well as time to takeoff but not for time to peak marginal penetration and peak marginal penetration (see Table 6b and Table 7b). The reason is that the Augmented Meta Bass uses information about product while the Functional Regression does not.

On the other hand, with the sole exception of the Classical Bass predicting year T+1 with for cutoff of T=10, the Augmented Functional Regression model is superior to all non functional models including Augmented Meta Bass, for every item to be predicted and for both cutoff times. The Augmented Functional Regression is also superior to Functional Regression, except in six instances where it is equal or slightly inferior (for T+1, T+2 years at Cutoff of T=5 years and T+1 to T+4 years at Cutoff of T=10 years). The superiority over Functional Regression is most noticeable in the time to takeoff and time to peak marginal penetration.

We also compute the fraction of curves for which Augmented Functional Regression outperforms each of the other methods (see Table 7). Augmented Functional Regression is superior to all other models on almost all of the items for more than 50% of the curves. It appears that the Functional Regression model is slightly superior for predicting T+1 but the augmented version is preferable for any longer range predictions. Most of the differences in Tables 6 and 7 are highly statistically significant. We also tested out the Augmented Functional Regression model with the addition of a predictor for geographic region as defined in the clustering section but found that the performance deteriorated slightly. In summary, the Augmented Functional Regression model outperforms other models in over 96 % of the comparisons with six alternate models to predict seven items across two cutoff times.

Discussion

Predicting the market penetration of new products is currently growing in importance due to increasing globalization, rapid introduction of new products, and rapid obsolescence of newly introduced products. Moreover, good record keeping has generated a wealth of new product penetration histories. The Bass model has been the standard model for analyzing such histories. However, the literature has not shown how exactly researchers should integrate the rich record of

penetration histories across categories with the penetration of an evolving new product to predict future characteristics of its penetration. Functional data analysis, which has gained significant importance in statistics, is well suited for this task. Our goal is to demonstrate and assess the merit of functional data analysis for predicting the market penetration of new products and compare it with the Bass model.

We compare the predictive performance of Functional Regression and Augmented Functional Regression with five other models: two simple or naïve models, the Classic Bass model, the Meta Bass model and the Augmented Meta-Bass Model. We compare the performance of these models on 8 items to be predicted. Ten fold cross-validation leads to three important results:

- The essential logic of integrating information across categories, which is the foundation of functional data analysis, provides superior prediction for an evolving new product.

- For a vast variety of items that need to be predicted, Augmented Functional Regression is distinctly superior to a variety of models, including simple or naïve models classic and enhanced Bass models and Functional Regression.

- Augmented Functional Regression provides superior predictions by using a) four underlying dimensions of the pattern of penetration across 760 categories b) deviations of a target category from the mean on these four dimensions and c) knowledge of the target product.

First, why don't simple models work well for prediction, as some researchers assert they do (Fader and Hardie 2005; Armstrong 1984; 1978; Armstrong and Lusk 1983) Our analysis makes it clear that there are two dimensions of information that are not captured by simple models. First, there is valuable information in the prior history of the new product, which as the

Bass model suggests, probably arises from consumers' innovative and imitative tendencies. Second, there is intrinsic information across products and countries, which may be used effectively to predict the penetration of an evolving new product. Despite their intuitive appeal, simple models that do not capture these sources of information will fail to predict well.

Second, why does the Classic Bass model not work as well for prediction? We suspect that it does not fully capture the two dimensions of information. First, the classic Bass model ignores other categories. This fact is borne out by the superiority of the Meta Bass and the Augmented Meta Bass in predicting items further out into the future. Both of these latter models capture information from other categories. Second, the classic Bass model is relatively flexible but nevertheless parametric. So it is limited in the range of shapes that it can take on. In particular, it is constrained to symmetric shapes for certain values of p and q. On the other hand, FDA provides higher flexibility by using a non-parametric approach. So it can capture a variety of flexible patterns, with the help of the principal components as explained earlier. The main disadvantage of a non-parametric method is that the increased flexibility can produce variability in the estimates. However, functional regression builds strength across the 760 curves to mitigate the problem of variability while generating more flexible estimates than those produced by the Classic Bass model.

Third, why does the Augmented Functional Regression outperform Functional Regression, especially for items further out into the future? The probable reason is that a particular product has a distinct pattern of penetration over time. Adding knowledge of that product further stabilizes the variability of predictions around their true value. This pattern can be seems in both the improvement of Augmented Meta-Bass over Meta-Bass and Augmented

Functional Regression over Functional Regression. Note also, that the improvement is greatest in peak marginal penetration, an item which arguably is most closely associated with a product.

Fourth, why is product seemingly more relevant for predicting market penetration than is country? The probable reason is that the evolution of market penetration seems to follow more distinct patterns by the nature of the product than by the country. For example, electronic products with universal appeal diffuse rapidly across countries both large and small and developed and developing. On the other hand, culturally sensitive products such as food appliances diffuse slowly overall and very differently across countries. Moreover, our data is only after 1977. Due to increasing industrialization of developing countries and flattening of the world economy, inter-country differences are much smaller after 1977, than before it.

Fifth, is the exclusion of marketing variables a limitation of Functional Regression? We posit that it not. Indeed, we show the superiority of Augmented Functional Regression, which includes a covariate for the product to which the curve belongs. In like manner, this model could also include covariates for marketing variables, such as price, quality, or advertising.

 To illustrate some of the above points, Figure 4 demonstrates 6 plots of the predictive performance of the Classic Bass model (red) and the Functional Regression model (green) relative to actual (black). These six plots are drawn from among those where Functional Regression does the best. For each plot, the first 10 periods are fitted on the estimation sample, while the last 5 periods are predictions on the hold out sample. Both models do well in the estimation periods. However, performance varies dramatically in the hold out periods.

Note how for curves, 676, 557, and 126, a generally flat curve with a late takeoff in the last two period, tricks the classic Bass into over predicting penetration for the holdout period. However, Functional Regression, which draws strength from other categories, is not so

influenced by the last two periods. Also, note how for curves 582, 572, and 121, the parameterization of the Bass model leads it to predict symmetric curves which are quite far from the actual.

This study has the following limitations. First, while the data is from a single source, the source itself does not record data before 1977. Indeed, we drop categories in some countries where we consider the year of introduction precedes 1977. We also drop categories in countries where penetration is not high enough till 2006. So the data are not balanced by country. Thus, substantive estimates about time to takeoff or about penetration by countries must be made with caution. However, that fact should not affect the comparison of the models, because all models have access to the same data. Second, depending on the release patterns of a particular product the product predictor used in Augmented Meta-Bass and Augmented Functional Regression may or may not be available. Third, our data do not include any marketing variables. However, the strength of the Augmented Functional Regression is that it can include such marketing variables. Thus, future research could address how better improvements can be obtained by using such information when available.

Table 1:

Scope of Prior Studies[*]

Authors	Categories	Countries
Gatignon, Eliashberg and Robertson (1989)	6 consumer durables	14 European countries
Mahajan, Muller and Bass (1990)	Numerous studies	
Sultan, Farley and Lehmann (1990)	213 applications	US, European countries
Helsen, Jedidi and DeSarbo (1993)	3 consumer durables	11 European countries and US
Ganesh and Kumar (1996)	1 industrial product	10 European countries, US, Japan
Ganesh, Kumar, Subramaniam (1997)	4 consumer durables	16 European countries
Golder and Tellis (1997)	31 consumer durables	USA
Putsis et al (1997)	4 consumer durables	10 European countries
Dekimpe, Parker and Sarvary (1998)	1 service	74 countries
Kumar, Ganesh and Echambadi (1998)	5 consumer durables	14 European countries
Golder and Tellis (1998)	10 consumer durables	USA
Kohli, Lehmann and Pae (1999)	32 appliances, house wares and electronics	USA
Dekimpe, Parker and Sarvary (2000)	1 innovation	More than 160 countries
Mahajan, Muller and Wind (2000)	Numerous studies	
Van den Bulte (2000)	31 consumer durables	USA
Talukdar, Sudhir, Ainslie (2002)	6 consumer durables	31 countries
Agarwal and Bayus (2002)	30 innovations	USA
Goldenberg, Libai and Muller (2002)	32 innovations	USA
Tellis, Stremersch and Yin (2003)	10 consumer durables	16 European countries
Golder and Tellis (2004)	30 consumer durables	USA
Stremersch and Tellis (2004)	10 consumer durables	16 European countries
Van den Bulte and Stremersch (2004)	293 applications	28 countries
Chandrasekaran and Tellis (2007)	16 products and services	40 countries

[*] Adapted from Chandrasekaran and Tellis (2007)

Table 2:

Classification of Models

		Analysis of Curves	
		Non Functional Analysis	**Functional Analysis**
Uses Information Across Curves	**No**	Classical Bass	
	Yes	Estimation Mean Last Period Projection Meta-Bass Augmented Meta-Bass	Functional Regression Augmented Functional Regression

Table 3:

Sample

Table 3a: Product Scope

Entertainment and Lifestyle	Household White Goods	Computers/ Communication
Cable TV	Air conditioner	Internet PC
Camera	Dishwasher	Personal Computer
CD Player	Freezer	Fax
Color TV	Microwave Owen	Satellite TV
DVD Player	Tumble Drier	Telephone
HiFi Stereo	Vacuum Cleaner	
Video Camera	Washing Machine	
Video Tape Recorder		
Videogame Console		

Table 3b: Geographical Scope

North America	Western Europe	Australasia	Eastern Europe	East Asia	West Asia	South America	Africa	Developing Asia
Canada	Austria	Australia	Belarus	China	Israel	Argentina	Algeria	Azerbaijan
USA	Belgium	N_Zealand	Bulgaria	HongKong	Jordan	Bolivia	Egypt	India
	Denmark		Croatia	Japan	Kuwait	Brazil	Morocco	Indonesia
	Finland		CzechRep	S_Korea	S_Arabia	Chile	Nigeria	Kazakhstan
	France		Estonia	Singapore	UAE	Colombia	Tunisia	Malaysia
	Germany		Hungary	Taiwan		Ecuador		Pakistan
	Greece		Latvia			Mexico		Philippines
	Ireland		Lithuania			Peru		Thailand
	Italy		Poland			Venezuela		Turkmenistan
	Netherlands		Romania					Vietnam
	Norway		Russia					
	Portugal		Slovakia					
	Spain		Slovenia					
	Sweden		Ukraine					
	Switzerland							
	Turkey							
	UK							

Table 4:

Proportions of Each Type of Product within Each Cluster

Product	Clusters					
	1	**2**	**3**	**4**	**5**	**6**
Fast growth electronics						
Cable TV	0.167	0.167	0.105	0.077	0.089	0.042
CD Player	0.083	0.167	0.184	0.154	0.103	0.020
DVD Player	0.083	0.167	0.447	0.051	0.205	0.027
Internet PC	0.583	0.361	0.105	0.077	0.082	0.056
Satellite TV	0.000	0.028	0.105	0.167	0.151	0.053
Videotape Recorder	0.083	0.111	0.053	0.077	0.021	0.002
Slower growth electronics						
Camera	0.000	0.000	0.000	0.000	0.007	0.018
Color TV	0.000	0.000	0.000	0.026	0.007	0.024
Fax	0.000	0.000	0.000	0.038	0.021	0.000
HiFi Stereo	0.000	0.000	0.000	0.026	0.014	0.076
PC	0.000	0.000	0.000	0.026	0.048	0.089
Telephone	0.000	0.000	0.000	0.000	0.014	0.027
Video camera	0.000	0.000	0.000	0.026	0.048	0.018
Videogame consol	0.000	0.000	0.000	0.192	0.110	0.022
Household goods						
Air conditioner	0.000	0.000	0.000	0.000	0.007	0.129
Freezer	0.000	0.000	0.000	0.000	0.000	0.071
Microwave oven	0.000	0.000	0.000	0.038	0.041	0.111
Tumble Drier	0.000	0.000	0.000	0.013	0.014	0.036
Vacuum cleaner	0.000	0.000	0.000	0.000	0.000	0.051
Washing Machine	0.000	0.000	0.000	0.013	0.007	0.020

Table 5:

Distribution of Each Geographic Region over Clusters

Geographic Region	Clusters					
	1	**2**	**3**	**4**	**5**	**6**
N. America, W. Europe and Australasia	5.4%	11.4%	14.5%	16.3%	22.9%	29.5%
Eastern Europe	0.0%	3.4%	3.4%	8.0%	24.7%	60.3%
East Asia	5.1%	8.5%	3.4%	20.3%	13.6%	49.2%
West Asia	0.0%	6.7%	11.1%	13.3%	22.2%	46.7%
South America	0.0%	0.0%	0.8%	10.0%	24.2%	65.0%
Africa	0.0%	0.0%	0.0%	3.1%	10.9%	85.9%
Developing Asia	0.0%	2.3%	0.0%	3.8%	8.3%	85.6%

Table 6:

Mean Absolute Deviation by Model and Item

Table 6a: Using five years (T=5) of training data

Item to be Predicted	Method						
	Estimation Mean	Last Period Projection	Classical Bass	Meta-Bass	Augmented Meta-Bass	Functional Regression	Augmented Functional Regression
T +1	0.009	0.004	0.003	0.008	0.008	0.002	0.002
T +2	0.012	0.008	0.007	0.011	0.010	0.005	0.005
T +3	0.015	0.010	0.012	0.013	0.011	0.008	0.007
T +4	0.017	0.014	0.017	0.016	0.012	0.012	0.008
T +5	0.020	0.017	0.020	0.020	0.014	0.016	0.010
Takeoff	3.358	2.828	NA	2.693	2.420	2.663	2.337
Peak Time	5.817	5.082	9.554	4.663	3.332	4.623	3.177
Peak Marginal Penetration	0.034	0.032	0.140	0.035	0.024	0.030	0.021

Table 6a: Using ten years (T=10) of training data

Item to be Predicted	Method						
	Estimation Mean	Last Period Projection	Classical Bass	Meta-Bass	Augmented Meta-Bass	Functional Regression	Augmented Functional Regression
T+1	0.011	0.005	0.004	0.008	0.008	0.004	0.005
T +2	0.011	0.007	0.006	0.009	0.008	0.005	0.006
T +3	0.012	0.007	0.008	0.010	0.010	0.006	0.006
T +4	0.012	0.009	0.011	0.009	0.009	0.008	0.008
T +5	0.012	0.010	0.012	0.011	0.011	0.009	0.009
Takeoff	3.604	2.918	NA	2.884	2.729	2.683	2.490
Peak Time	4.681	3.952	7.536	3.442	2.956	3.668	2.913
Peak Marginal Penetration	0.027	0.024	0.043	0.025	0.021	0.023	0.018

Table 7:

Superiority of Augmented Functional Regression Over Other Models

Table 7a: Fraction of curves for which Augmented Functional Regression outperforms other models using five years of training data

Item to be Predicted	Method					
	Estimation Mean	**Last Period Projection**	**Classical Bass**	**Meta-Bass**	**Augmented Meta-Bass**	**Functional Regression**
T+1	0.89	0.66	0.51	0.84	0.85	0.51
T+2	0.80	0.70	0.50	0.76	0.77	0.57
T+3	0.77	0.73	0.52	0.71	0.75	0.69
T+4	0.78	0.77	0.6	0.72	0.70	0.75
T+5	0.78	0.80	0.64	0.76	0.73	0.78
Takeoff	0.64	0.63	NA	0.6	0.52	0.59
Peak Time	0.75	0.7	0.77	0.66	0.55	0.68
Peak Marginal Penetration	0.70	0.69	0.66	0.73	0.59	0.69

Table 7b: Fraction of curves for which Augmented Functional Regression outperforms other models using ten years of training data

Item to be Predicted	Method					
	Estimation Mean	**Last Period Projection**	**Classical Bass**	**Meta-Bass**	**Augmented Meta-Bass**	**Functional Regression**
T+1	0.86	0.52	0.33	0.71	0.72	0.37
T+2	0.81	0.63	0.49	0.73	0.72	0.49
T+3	0.79	0.62	0.56	0.70	0.71	0.57
T+4	0.74	0.61	0.63	0.63	0.62	0.59
T+5	0.66	0.62	0.61	0.62	0.64	0.59
Takeoff	0.66	0.59	NA	0.54	0.53	0.53
Peak Time	0.70	0.64	0.77	0.54	0.49	0.57
Peak Marginal Penetration	0.71	0.67	0.62	0.66	0.57	0.68

Figure 1:

Analytic Framework

Penetration

8 Items for Prediction
- T+1 to 5 periods,
- time to takeoff,
- time to peak penetration,
- peak penetration

Penetration curve over time

7 Prediction Models
- Estimated Mean
- Last Observation Projection
- Classic Bass,
- Meta Bass,
- Augmented Meta Bass,
- Functional Regression,
- Functional-Bass

Estimation Sample (90%)

Holdout Sample (10%)

Time

Cross-validated   Training Period   Testing Period

Figure 2:

Illustration of First Two Functional Principal Component Curves
(Based On 10 Years Of Training Data)
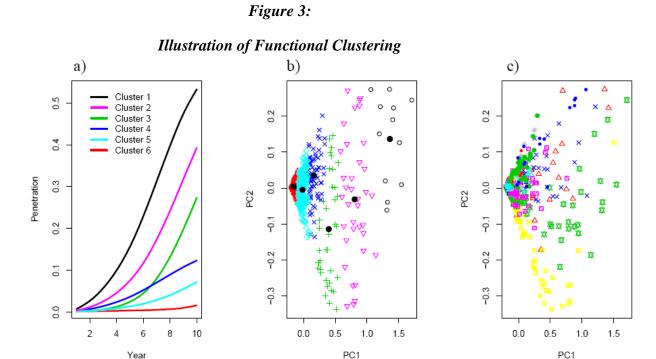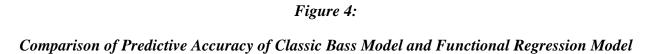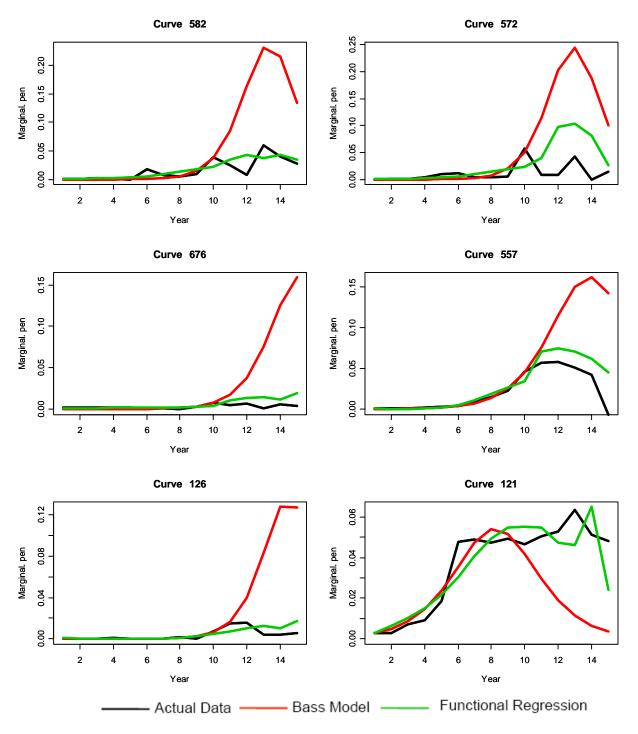
Illustration of Functional Clustering

Figure 2 a) The shapes of the average penetration curves within each of the six clusters.

b) The first two principal component scores for all 760 curves. A different color and plotting

symbol has been used for each cluster with a black solid circle for the cluster centers. c)Same as

b) but with different symbols for each product.

Figure 4:

Comparison of Predictive Accuracy of Classic Bass Model and Functional Regression Model

36

References

Agarwal, Rajshree and Barry L. Bayus (2002) "Market Evolution and Sales Takeoff of Product Innovations" Management Science, 48 (8), 1024–1041.

Armstrong, J. Scott (1978), "Forecasting with econometric methods: Folklore versus Fact," Journal of Business, 51, 4, 549-564.

____ (1984), "Forecasting by Extrapolation: Conclusions from 25 Years of Research," Interfaces, 14, Nov-Dec, 52-66.

____ and Edward J. Lusk (1983), "The accuracy of alternative extrapolation models: Analysis of forecasting competition through open peer review," Journal of Forecasting, 2, 3, 259-262.

Bass, F. M. (1969), "A new product growth for model consumer durables," Management Science 15, 215–227

Chandrasekaran, Deepa and Gerard J. Tellis (2007), "Global Takeoff of New Products," forthcoming, Marketing Science.

Dekimpe, Marnik, Philip Parker, and Miklos Sarvary. (1998) "Staged Estimation of International Diffusion Models: An Application to Global Cellular Telephone Adoption." Technological Forecasting and Social Change, 57, 105–132.

Fader, P. S. and Hardie, B. G. 2005, "The value of simple models in new product forecasting and customer-base analysis: Research Articles," Applied Stochastic Modeling Bus. Ind. 21, 4-5 (Jul. 2005), 461-473.

Foutz, N. and Jank, W. (2007), "The Wisdom of Crowds: Pre-release Forecasting via Functional Shape Analysis of the Online Virtual Stock Market," Presentation, Marketing Science Conference, Singapore.

Ganesh, Jaishankar, and V. Kumar. (1996) "Capturing the Cross-National Learning Effect: An Analysis of an Industrial Technology Diffusion." Journal of the Academy of Marketing Science, 24 (4), 328–337

____, ____, and V. Subramaniam. (1997) "Learning Effect in Multinational Diffusion of Consumer Durables: An Exploratory Investigation." Journal of the Academy of Marketing Science, 25 (3), 214–228.

Gatignon, Hubert, Jehoshua Eliashberg, and Thomas S. Robertson, (1989) "Modeling Multinational Diffusion Patterns: An Efficient Methodology." Marketing Science, 8 (3), 231–247.

Goldenberg, Jacob, Barak Libai, and Eitan Muller (2001) "Using Complex Systems Analysis to Advance Marketing Theory Development: Modeling Heterogeneity Effects on New Product Growth through Stochastic Cellular Automata." Academy of Marketing Science Review

Golder, Peter N., Gerard J. Tellis. (1997) "Will It Ever Fly? Modeling the Takeoff of Really New Consumer Durables" Marketing Science 16 (3), 256–270.

____ and ____ (1998) "Beyond Diffusion: An Affordability Model of the Growth of New Consumer Durables." Journal of Forecasting, 17 (3/4), 259–280.

____ and ____. (2004) "Going, Going, Gone: Cascades, Diffusion, and Turning Points of the Product Life Cycle." Marketing Science, 23 (2), 207–218.

____ (2000), "Historical Method in Marketing Research with New Evidence on Long-term Market Share Stability" Journal of Marketing Research 37 (May) 156-72.

Green, P. J. and Silverman, B. W. (1994). Nonparametric Regression and Generalized Linear Models. Chapman and Hall

Hastie, T. J. and Tibshirani, R. J. (1990). Generalized Additive Models. Chapman and Hall.

____., ____ and Friedman, J. (2001) The Elements of Statistical Learning: Data Mining,

> Inference and Prediction. Springer.

Helsen, Kristaan, Kamel Jedidi, and Wayne DeSarbo. (1993) "A New Approach to Country

> Segmentation Utilizing Multinational Diffusion Patterns" Journal of Marketing, 57 (4),

> 60–71.

James, G. M., Hastie, T., and Sugar, C. A. (2000), "Principal Component Models for Sparse

> Functional Data", Biometrika 87, 587-602

Kohli, Rajeev, Donald Lehmann, and Jae Pae, (1999) "Extent and Impact of Incubation Time in

> New Product Diffusion" Journal of Product Innovation Management, 16, 134–144.

Kumar, V., Jaishankar Ganesh, and Raj Echambadi (1998) "Cross National Diffusion Research:

> What Do We Know and How Certain Are We?" Journal of Product Innovation

> Management, 15, 255–268.

MacQueen, J. B. (1967). Some Methods for classification and Analysis of Multivariate

> observations, Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and

> Probability, Berkeley, University of California Press, 1:281-297

Mahajan, Vijay, Eitan Muller, and Frank M. Bass (1995) "Diffusion of New Products: Empirical

> Generalizations and Managerial Uses." Marketing Science, 14 (3), Part 2 of 2, G79–G88.

____, Eitan Muller, and Yoram Wind. (2000) "New Product Diffusion Models: From Theory to

> Practice." In New Product Diffusion Models, ed. Mahajan, V., Eitan Muller, Yoram

> Wind, pp. 3-24, Boston: Kluwer Academic.

Peres, Renana, Eitan Mueller and Vijay Mahajan (2007), "Review of Research on Diffusion,"

> forthcoming, Marketing Science.

Putsis, William P. Jr., Sridhar Balasubramanian, Edward Kaplan, and Subrata Sen. (1997) "Mixing Behavior in Cross-Country Diffusion" Marketing Science, 16 (4), 354–369.

Ramsay, J.O. and Silverman, B.W. (2005), "Functional Data Analysis", Springer (2nd Edition)

Stremersch, Stefan, and Gerard J. Tellis. (2004), "Managing International Growth of New Products" International Journal of Research in Marketing

Sugar, C. A. and James, G. M. (2003), "Finding the number of clusters in a data set: An information theoretic approach," Journal of the American Statistical Association 98, 750–763.

Sultan, Fareena, John U. Farley, and Donald R. Lehmann (1990), "A Meta-Analysis of Diffusion Models" Journal of Marketing Research, 27, 70–77.

Talukdar, Debabrata, K. Sudhir, and Andrew Ainslie (2002) "Investigating New Product Diffusion Across Products and Countries" Marketing Science, 21 (1), 97–114.

Tellis, Gerard J., Stefan Stremersch, and Eden Yin (2003) "The International Takeoff of New Products: The Role of Economics, Culture and Country Innovativeness" Marketing Science, 22 (2), 188–208.

Van Den Bulte, Christophe, and Stefan Stremersch (2004) "Social Contagion and Income Heterogeneity In New Product Diffusion: A Meta-Analytic Test" Marketing Science, 23 (4), 530–544.

____. (2000) "New Product Diffusion Acceleration: Measurement and Analysis" Marketing Science, 19 (4), 366–380.

Venkatesan, Rajkumar, Trichy V. Krishnan and V. Kumar (2004), "Evolutionary estimation of macro-level diffusion models using genetic algorithms: An alternative to nonlinear least squares," Marketing Science 23, 451–464. 12.