



Choice Models and Customer Relationship Management

WAGNER KAMAKURA*

kamakura@duke.edu

CARL F. MELA*

Duke University

ASIM ANSARI

Columbia University

ANAND BODAPATI

University of California, Los Angeles

PETE FADER

University of Pennsylvania

RAGHURAM IYENGAR

Columbia University

PRASAD NAIK

University of California, Davis

SCOTT NESLIN

Dartmouth College

BAOHONG SUN

Carnegie Mellon University

PETER C. VERHOEF

Erasmus University Rotterdam

MICHEL WEDEL

University of Michigan

RON WILCOX

University of Virginia

Abstract

Customer relationship management (CRM) typically involves tracking individual customer behavior over time, and using this knowledge to configure solutions precisely tailored to the customers' and vendors' needs. In the context of choice, this implies designing longitudinal models of choice over the breadth of the firm's products and using them prescriptively to increase the revenues from customers over their lifecycle. Several factors have recently contributed to the rise in the use of CRM in the marketplace:

- A shift in focus in many organizations, towards increasing the share of requirements among their current customers rather than fighting for new customers.

*Co-Chairs.

- An explosion in data acquired about customers, through the integration of internal databases and acquisition of external syndicated data.
- Computing power is increasing exponentially.
- Software and tools are being developed to exploit these data and computers, bringing the analytical tools to the decision maker, rather than restricting their access to analysts.

In spite of this growth in marketing practice, CRM research in academia remains nascent. This paper provides a framework for CRM research and describes recent advances as well as key research opportunities. See <http://faculty.fuqua.duke.edu/~mela> for a more complete version of this paper.

Keywords: customer relationship management, direct marketing

Introduction

What is CRM?

Analytical customer relationship management (CRM) is the process of collecting and analyzing a firm's information regarding customer interactions in order to enhance the customers' values to the firm. Firms exploit such information by designing strategies uniquely targeted to consumer needs. This process enhances loyalty and increases switching costs, as information on consumer preferences affords an enduring competitive advantage. By integrating various data (e.g. across purchases, operations, service logs, etc.), choice researchers can obtain a more complete view of customer behavior. These developments cut across industries, including banking, telephony, Internet, and other areas that have received limited attention in the marketing literature. In addition, each industry likely has unique challenges of its own.

We differentiate between analytical CRM, which is the focus of this paper, and behavioral CRM. Analytical CRM involves using firms' data on its customers to design longitudinal models of choice over the breadth of the firm's products and using them prescriptively to increase the revenues from customers over their lifecycle. In contrast, behavioral CRM uses experiments and surveys to focus upon the psychological underpinnings of the service interaction, or the managerial structures that make CRM effective.

A focus on CRM is warranted given the explosive growth of the analytical CRM applications in industry (Market Research.com forecasts analytical CRM revenue to increase from \$2.4 billion in 2003, to \$43 billion in 2007). Technological enhancement in information technology and increased addressability of customers via new channels has fueled this growth. Thus, it is surprising there are only few papers that seek to assess the state of research in this area, or outline the challenges unique to this area. This paper seeks to address this void.

What Novel Implications of CRM Exist for Choice Modeling?

Choice decisions in the context of CRM include firm choices (whom to target, when and with what) and customer choices (whether, what, when and where to buy). Several aspects make CRM a novel and potentially fruitful domain of inquiry for choice researchers:

- CRM applications typically involve massive amounts of data. These include many observations and many variables.
- CRM applications often involve an inward looking view of the customer, as competitive information is often impossible to obtain.
- Analytical CRM is typically dynamic, as trade-offs in current programs are made against future revenues. In conjunction with large data, this implies that new optimization techniques are needed to cope with such problems.
- Low response rates are often the norm, calling for more flexible response models than the popular logistic regression model.
- Unlike many scanner-panel applications, customers are addressable (Blattberg and Deighton, 1991). As a result, it is common to run large field experiments in these settings and control the sampling approach, affording a greater degree of control over the choice task. Addressability also implies it is easier to target consumers.

A Framework for CRM Research

CRM research can be organized along the *customer lifecycle*, including customer *acquisition*, *development* and *retention* strategies. Customer *acquisition* extends from the channels customers use to first access the firm (Ansari et al., 2004) to the promotions that bring them to a firm. The value of a customer can also be enhanced by the firm through appropriate *development* strategies such as delivering customized products (Ansari and Mela, 2003) and cross-selling (Kamakura et al., 1991, 2003). Finally, early detection and prevention of customer attrition can also enhance the total lifetime of the customer base, if efforts are focused on the *retention* of valuable customers.

The customer lifecycle implies that each customer has a value over his or her tenure with a firm. Estimating the lifetime value of a customer by itself requires sophisticated modeling, as it involves predictions of both revenues and retention probabilities. Several approaches exist to measure customer lifetime value (CLV). The relative merits of these different approaches is considered in Jain and Singh (2002) and Venkatesan and Kumar (2003). However, scant evidence exists regarding the accuracy of CLV predictions. There may be considerable room for improvement based on criteria such as out of sample validation, especially at the individual level.

The Pareto/NBD model of lifetime value may offer promise in the forecasting of lifetime value (Reinartz and Kumar, 2003), and recent advances by Fader et al. (2004) mitigate the computational burden of this model. Incorporation of covariates using a proportional hazards model and the addition of discounting to the Pareto/NBD could enrich this literature. Other fruitful areas for inquiry include the role of network effects in lifetime value, the effect of time aggregation and aggregation across a household, the role of predictors other than past purchases on lifetime value (such as inbound contacts or marketing), and a better accounting of costs.

An idea closely related to customer lifetime value is *consumer lifetime value* (Du and Kamakura, 2005). The distinction pertains to the perspective of the decision maker (firm vs. consumer), scope of information, and the approaches used to compute the value of a

customer. Customer lifetime value is typically an inward-looking view of the consumer predicated on firms' internal records for the purpose of determining the value of the customer to the firm. In contrast, consumer lifetime value encompasses all behaviors of a consumer across multiple or competing firms and assumes the perspective of a consumer making inter-temporal choices over categories and time so as to maximize his or her utility. Consumer lifetime value models, which often combine internal and syndicated data, are generally applicable in industries where customer tenures are long, needs change over time and can be linked to life stages, and consumers trade-off future for current utility (such as savings and consumption). Ideally, consumer lifetime value models can be linked with internal firm records to obtain a better sense of which consumers to target. Consumers with a low share of wallet but a high consumer lifetime value may be especially attractive to a firm.

Our discussion of the state of CRM research proceeds using the customer lifecycle framework (acquisition, development and retention), and we shall describe the issues and methodological challenges unique to each stage.

Acquisition

Issues

The objective of acquisition strategies is to obtain more and profitable customers. For example, new home buyers are targeted for home insurance. In spite of its importance, identifying potential customers for acquisition is an area of scant attention. In general, acquisitions are profitable if the expected value of attaining the customer (over the lifetime) exceeds the cost (Blattberg et al., 2001). However, forecasts of likely response are predicated upon past response, and subject to regression to the mean if based on selection from such past response. Deeper analysis of appropriate probabilistic thresholds for mailing could yield significant advances in this area.

Customer acquisition occurs across an array of channels (e.g., direct television, direct mail, Internet, telemarketing, etc.) and researchers have begun to assess the efficacy of channel acquisition strategies and their effect on subsequent behaviors (Bolton et al., 2004; Verhoef and Donkers, 2005; Thomas, 2001). For example, Bolton et al. (2004) argue that customers acquired through channels with a price emphasis tend to be less loyal. A related issue pertains to referral programs, and there has been little analytical research on the efficacy and design of these programs.

Classic behavioral models of consumer adoption (need recognition followed by information search, purchase, and post-purchase service encounters) are useful in comprehending the effects of multi-channel acquisition strategies, as some channels are likely better for information search, while others are better for service or purchase. Thus, acquisition in a multi-channel environment should consider the interaction between these channels (Blattberg et al., 2004).

Approaches

Analytical CRM models of acquisition are plagued by a number of data-related problems. Databases are massive in size (either in number of observations or number of variables), requiring models that can be estimated in large samples, and are easily scalable for large-scale implementations. Extremely low response rates (in the single digits) result in a large proportion of false negatives, calling for new sampling techniques and modeling approaches.

In the context of rare events, parametric models such as logit or probit assume specific shapes of the underlying density function, implying a given tail probability expression that remains invariant to observed data characteristics. Further, a customer database may contain information on hundreds of covariates. Thus, space reduction methods (e.g., slice inverse regression (Li, 1991)) need to be combined with non-parametric approaches to predict rare events in these contexts. To redress these considerations, Naik and Tsai (2004) proposed an Isotonic Single-Index Model and developed an efficient algorithm for its estimation. An empirical application in database marketing context with 66 variables and a binary response reveals the standard logistic distribution overestimates the purchase probability substantially.

Many issues remain regarding the modeling of rare events and large data (Balasubramanian et al., 1998). For example, non-parametric approaches offer no restrictions on the response function, such as decreasing returns to scale. Model comparison in the context of very large data is relatively under-researched. It is unclear whether the choice of a link function affects the selection of variables in the response model, and vice versa. Naik and Tsai (2001) address this issue using a single-index model selection criterion to simultaneously select the link function and variables to retain; its extension to discrete response variable is necessary for analyzing customer choices in CRM contexts. In addition, the extraction of meaningful factors with simple structures in the context of binary response is an open issue (see Naik and Tsai, 2005) for a discussion of this issue for non-binary response).

Development

Customer development pertains to the growth of revenues from existing customers. Activities used to develop customer demand include cross-selling, up-selling, and channel management. Cross-selling refers to encouraging customers to buy across categories and up-selling involves increasing demand in existing categories. Cross-selling yields both immediate profit as well as the potential to deepen existing relationships, thereby increasing the switching costs associated with purchasing from another vendor. Channel management refers to migrating customers across channels in order to lower costs or increase demand (via channel-specific promotions or features).

Issues

Several papers in marketing and data-mining that relate to up-selling and cross-selling (Kamakura et al., 1991; Ansari et al., 2000; Ansari and Mela, 2003; Bodapati, 2004;

Iacobucci et al., 2000; Kamakura et al., 2003; Knott et al., 2002; Li et al., 2004; Ying et al., 2004). Most use information on past behaviors, and/or marketing response to predict who would be prone to up-selling and cross-selling and what to recommend. Criterion for targeting of up and cross-selling include the likelihood of purchase as well as the response to the targeting (i.e., if a customer is likely to buy anyway, there is little point in targeting them).

The dynamics of cross-selling are an interesting avenue for further research—banks, for example, by offering multiple services, can make it harder to switch to a competing bank. Cross-selling can be dictated by the consumer life-cycle, with certain needs appearing at certain times. Finally, the role of network effects in cross selling (e.g., viral marketing) remains relatively under-researched.

Channel management has received scant attention. As customers progress through information search, purchase, and service, it is desirable to align channels, stages and customers. Another key issue in customer channel regards whether sales in one channel cannibalize another. Biyalogorsky and Naik (2003) find this not to be the case, but it would be desirable to generalize this result across categories.

Approaches

Much work has been done regarding recommendation systems for ratings data, including correlation approaches, which compute similarity across raters and use this information to make ratings (Breese et al., 1998), and Bayesian model based approaches that consider the effects of the characteristics of goods when making recommendations (Ansari et al., 2000). Yet several issues remain, most notably those relating to what events should be modeled in a recommendations system, scalability to large data, dealing with the large amount of missing ratings data, the role of social contagion in recommendation, the appropriate categorization of products into categories, and the exchangeability of observations (e.g., when the timing of customer behavior matters).

Data Issues Another challenge related to managing customer development pertains to the data used to model response. Companies nowadays record their transactions with each individual customer and store those in customer transaction databases. Rather than information on a sample of customers, companies in question have access to the entire population of customers. But, such databases also lack important information. Observations that are missing due to the design of studies are the rule rather than the exception in CRM applications. These data issues raise unique opportunities and challenges to the customer analyst:

1. Customer transaction records, often consisting of purchase indicators or counts, can be linked to the additional data at the individual customer level based on ZIP-codes. Here, the problem is that the two data sources have different levels of aggregation which needs to be accommodated (Steenburgh et al., 2003).
2. Transaction data are often enriched with supplementary data; geo-demographic, values and life-style data are often used for that purpose. Here, the problem is how to combine data from different sources, a problem known as *data fusion* (Du et al., 2005; Kamakura

and Wedel, 1997, 2000; Rassler, 2002; Gilula et al., 2004). Data on two different sets of variables are obtained from two independent samples, while a number of variables (usually demographics) are measured in both samples.

3. Data on the use of products and services from competitors, and “soft data” such as customer satisfaction, often provide important additional insights to the company, but are lacking in the transaction database and need to be collected in separate surveys. Due to the survey costs, such data are usually only collected from a sample of customers in the database. The information from these surveys is needed for all customers and the database needs to be augmented with the auxiliary data (Kamakura and Wedel, 2003).
4. Large amounts of supplementary ‘soft’ data such as product-specific life styles or customer satisfaction are often collected. Here, a questionnaire is split in two or more parts and those different parts are administered to different groups in a sample (Adiguzel and Wedel, 2004). Because marketing management demands more and more information, while the responsiveness to lengthy questionnaires has decreased, split questionnaires are regularly employed in customer satisfaction surveys and pop-up questionnaires on the internet.
5. In a panel, items are often administered in a rotating fashion to different parts of the sample in different episodes. In each of these conditions, data are missing intentionally, i.e., they are specified missing in the design of the study to reduce respondent burden, increase the response rate or reduce costs of data collection.

Several imputation models have been used, including location-scale models, mixture models and factor models (Kamakura and Wedel, 1997, 2000). These models capture the covariance of the variables in the two databases, and use those to impute the missing values. Several issues are important in data-fusion, sub-sampling and split questionnaire design. The most important one is the search for powerful fusion variables, since the resulting complete dataset is as good as these fusion variables allow for. Second, there is a need for the development of fast and efficient—for example sequential- procedures that are scalable to the size of current transaction databases and that can be applied on line. Third, current methods are strongly dependent on distributional assumptions, and nonparametric approaches are called for. Gilula et al. (2004) provide important steps towards addressing the latter issues.

Dynamic Structural Modeling and Strategic Behavior in CRM The concept of using multiple services and products to raise consumer switching costs over time suggests that dynamic programming (DP) might be a useful foundation for structural models describing consumer behavior in the context of multi-category buying, especially when these needs are a natural result of changes in the consumer lifestage or future rewards (e.g., loyalty programs). This could be coupled with DP on the part of the firm in its marketing to determine the optimal mix of offers at the optimal time. The structural nature of this approach indemnifies it from the Lucas critique, making it especially useful for forecasting and policy experiments (Franses, 2005). Past work in marketing and elsewhere has used DP for CRM. Shi and Gonul (1998) developed a framework for prescribing an optimal mailing policy for a direct marketer. The model allowed for forward-looking behavior for both firm and

customers. The forward looking models for customer and firm were simultaneously solved and an optimal mailing policy for the firm was specified. In a similar vein, Bitran and Mondschein (1996) and Simester et al. (2004) allowed for a forward looking behavior on the part of firms and specified an optimal mailing policy. Lewis (2003) presents a dynamic model wherein firms offer introductory discounts and then raise prices to extract profits after consumer switching costs increase with usage. Another area of interest regarding DP in CRM is the optimal allocation of services, such as the outsourcing of inbound service calls (Li et al., 2004). To guarantee both long-term cost efficiency and overall customer satisfaction, a firm might sacrifice short-term efficiency by strategically assigning a few callers to the various call centers to learn about their service quality. A better knowledge of service centers' efficiencies can reduce future service time and improve overall customer satisfaction. Several issues arise when applying DP techniques in the context of CRM:

1. The estimation of dynamic programs can use either historical data or data from randomized field experiments (Anderson and Simester, 2004). Historical data is easily available, but is subject to endogeneity between the intervention and customer purchase history (since firms typically use purchase patterns of their customers to decide upon the time of intervention).
2. As CRM data typically contains many variables, the states and controls "explode," leading to the appropriately called "curse of dimensionality," which hinders estimation.
3. Optimal contact strategy typically involves making real time decisions. Dynamic programming techniques in environments characterized by rapid updates of customer information are computationally intensive and time-consuming. In such environments, heuristics that closely resemble the optimal rules from a dynamic program but are faster to recalculate are perhaps a more practical solution.
4. The efficacy of applying DP to customize recommendations remains an open issue. Zhang and Wedel (2004) use an optimization approach predicated upon a response model calibrated on individual level data. They find an aggregate optimization performs close to segment level and individual-level models, with some more efficacy of the individual-level models in online venue. Clearly, this is an issue that warrants further study.

Customer Retention

Customer retention has a significant impact on firm profitability. Gupta et al. (2004) find that a 1% improvement in retention can increase firm value by 5%. Churn refers to the tendency for customers to defect or cease business with a company. Marketers interested in maximizing lifetime value realize that customer retention is a key to increasing long-run firm profitability. A focus on customer retention implies that firms need to understand the determinants of customer defection (churn) and are able to predict those customers who are at risk of defection at a particular point in time. An understanding of the drivers for customer defection can help companies in designing CRM strategies and interventions aimed toward increasing customer loyalty and prolonging the lifetime of customers.

Issues

Given the importance of customer retention, companies use a variety of mechanisms for reducing churn. These efforts can be grouped into three main areas: improving service quality, targeting interventions to prevent churn, and loyalty programs.

Firms' investment in improving service quality and customer satisfaction is based on the assumption that they improve customer retention. While some studies have found a link between satisfaction and retention (Rust and Zahorik, 1993), others have questioned this link. For example, Mittal and Kamakura (2001) find the link between customer satisfaction and retention to be moderated by customer characteristics.

Recent research finds that retention rates are affected by the channel(s) utilized by the customer. Ansari et al. (2004) find that e-mails tend to drive persons to the Internet, and that purchases on the Internet lessen inertia in buying and loyalty. They conjecture that this arises from lower service levels and lower switching costs. Zhang and Wedel (2004) find the opposite effect in the context of grocery purchases, perhaps due to the use of e-shopping lists, which might actually raise switching costs. In light of these conflicting findings it would be desirable to better ascertain the role of optimal channel mix in retention.

Since the introduction of frequent flier program by American Airlines in the 1980s, loyalty programs have become ubiquitous in almost every industry. The interest in loyalty programs has increased over time as more and more companies use them for developing relationships, stimulating product or service usage, and retaining customers. However, three important questions remain: (a) Do loyalty programs work, (b) which customers do they work for the best, and (c) how best to design a loyalty program?

In spite of the pervasiveness of loyalty programs, their effectiveness is far from clear. Some studies find that loyalty programs increase customer retention and customer development (Bolton et al., 2000; Leenheer et al., 2004; Verhoef, 2003); others find no impact on retention but improvement in share of wallet (Sharp and Sharp, 1997); and yet others find almost no difference in the behavior of loyalty program members and non-members (Dowling and Uncle, 1997). Kopalle and Neslin (2003) investigate the economic viability of frequency reward programs in a competitive environment, and find brands benefit from reward programs when customers value future benefits, reward programs expand the market and if the brand has a higher preference. They also caution that if the gains from a reward program come mainly from competitors, a classic prisoner's dilemma situation evolves and profitability is eroded. Accommodating endogeneity becomes crucial in evaluating the success of different program interventions (Leenheer et al., 2004) because of customers' self-selection in the loyalty programs.

Optimal targeting of loyalty programs is also an open issue. Conventional wisdom suggests that loyalty programs should be designed to reward a firm's best customers. However, Lal and Bell (2003) found that, in the context of grocery stores, loyalty programs do not affect the behavior of best customers. Instead, these programs have the biggest impact on a store's worst (or low spending) customers.

Several questions pertain to loyalty program design, including whether rewards should use cash or merchandise, offer luxury or necessity goods, be probabilistic or deterministic, or

whether to use the firm's own products. Recent behavioral research provides some guidelines on these important issues (Kivetz, 2003; Kivetz and Simonson, 2002; Kivetz et al., 2004), and these findings have implications for modeling loyalty program design.

Approaches

Many aspects of churn have been modeled in the literature. First, whether churn is hidden or observable influence the overall approach to modeling. In some industries, customer defection is not directly observed, as customers do not explicitly terminate a relationship, but can become inactive. In other industries, however, the defection decision is observable as customers cease their relationship via actively terminating their contract with the firm (Schmittlein et al., 1987; Fader et al., 2004).

The modeling approach could also depend critically on the relative importance placed on explanation/interpretation *vis a vis* prediction. Models that are better at explanation may not necessarily be better at prediction. The empirical literature in marketing has traditionally favored parametric models (such as logistic or probit regression or parametric hazard specifications and zero-inflated poisson models) that are easy to interpret. Similar to the previous discussion on acquisition, churn is a rare event that may require new approaches from data mining, machine learning and non-parametric statistics that emphasize predictive ability (Hastie et al., 2001). These include projection-pursuit models, jump diffusion models, neural network models, tree structured models, spline-based models such as Generalized Additive Models (GAM), and Multivariate Adaptive Regression Splines (MARS), and more recently approaches such as support vector machines and boosting (Lemmens and Croux, 2003). Further work is needed to understand the relative merits and disadvantages of these different approaches for CRM applications. Neslin et al. (2004) provide some insights on the relative performance of different approaches to predictive modeling for churn based on the Teradata tournament for modeling defections.

Research opportunities also exist in modifying flexible semi-parametric models to handle unique facets of marketing data. For example, one can extend semi-parametric models to handle sources of unobserved heterogeneity in longitudinal-data situations involving multiple records for customers. Finally, more research is needed for the automatic handling of a large number of explanatory variables and for modifying different models to handle rare-event data (Donkers et al., 2003; King and Zheng, 2001).

Summary

Given the rise of analytical CRM, we argue that the topic will assume an increasingly central role in research in marketing. As such, we sought to assess the state of the field in this area, and outline challenges unique to choice researchers in CRM. As we have argued, there are a plethora of issues and analytical challenges that remain unresolved. We hope that this article inspires new answers and new approaches to resolve them.

References

- Adiguzel, Feray and Michel Wedel. (2004). "The Design of Split Questionnaires," Working Paper, University of Michigan.
- Anderson, Eric T. and Duncan Simester. (2004). "Long-Run Effects of Promotion Depth on New Versus Established Customers: Three Field Studies," *Marketing Science* 23(1), 4–20.
- Ansari, A., S. Essegai, and R. Kohli. (2000). "Internet Recommendation Systems," *Journal of Marketing Research*, 37, 363–375.
- Ansari, Asim and Carl F. Mela. (2003). "E-Customization," *Journal of Marketing Research* 40(2), 131–145.
- Ansari, Asim, Carl Mela, and Scott A. Neslin. (2004). "Customer Channel Migration," Working Paper, Columbia University School of Business, NY.
- Balasubramanian, S., S. Gupta, W. Kamakura, and M. Wedel. (1998). "Modeling Large Data Sets in Marketing," *Statistica Neerlandica* 52, 303–323.
- Bitran, Gabriel R. and Susana V. Mondschein. (1996). "Mailing Decisions in the Catalog Industry," *Management Science* 42(9), 1364–1381.
- Biyalogorsky, Eyal and Prasad Naik. (2003). "Clicks and Mortar: The Effect of Online Activities on Offline Sales," *Marketing Letters* 14(1), 21–32.
- Blattberg, Robert C. and John Deighton. (1991). "Interactive Marketing: Exploiting the Age of Addressability," *Sloan Management Review* 33(Fall), 5–14.
- Blattberg, Robert C., Gary Getz, and Jacquelyn S. Thomas. (2001). *Customer Equity: Building and Managing Relationships as Valuable Assets*. Boston, MA: Harvard Business School Press.
- Blattberg, Robert C., Kim and Scott A. Neslin. (2004). *Database Marketing* (forthcoming).
- Bodapati, A. V. (2004). "Recommendation Systems," Working Paper, Anderson School of Management, UCLA.
- Bolton, Ruth, P. K. Kannan, and Matthew Bramlett. (2000). "Implications of Loyalty Program Membership and Service Experiences for Customer Retention and Value," *Journal of the Academy of Marketing Science* 28(1), 95–108.
- Bolton, Ruth N., Katherine N. Lemon, and Peter C. Verhoef. (2004). "The Theoretical Underpinnings of Customer Asset Management: A Framework and Propositions for Future Research," *Journal of the Academy of Marketing Science* 32(3), 271–292.
- Breese J. S., D. Heckerman, and C. Kadie. (1998). "Empirical Analysis of Predictive Algorithms for Collaborative Filtering," In G. F. Cooper and S. Moral (eds.), *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pp. 43–52, Morgan Kaufmann.
- Bas Donkers, Philip Hans Franses, and Peter C. Verhoef. (2003). "Selective Sampling for Binary Choice Models," *Journal of Marketing Research* 40(4), 492–497.
- Dowling, Grahame R. and Mark Uncles. (1997). "Do Customer Loyalty Programs Really Work?" *Sloan Management Review* 38(4), 71–82.
- Du, Rex and Wagner A. Kamakura. (2005). "Household Lifecycles and Life Styles in America," *Journal of Marketing Research*, (forthcoming).
- Du, Rex, Wagner A. Kamakura, and Carl F. Mela. (2005). "Customers' Share of Category Requirements," Working Paper, Duke University Marketing Department.
- Fader, P. S., B. G. S. Hardie, and K. L. Lee. (2004). "Counting Your Customers' the Easy Way: An Alternative to the Pareto/NBD Model," Working Paper, Wharton Marketing Department.
- Franses, Philip, H. (2005). "On the Use of Econometric Models for Policy Simulation in Marketing," *Journal of Marketing Research* (42), 4–14.
- Gilula, Zvi, Robert E. McCulloch, and Peter E. Rossi. (2004). "A Direct Approach to Data Fusion," Working Paper, University of Chicago.
- Gupta, Sunil, Donald R. Lehmann, and Jennifer Ames Stuart. (2004). "Valuing Customers," *Journal of Marketing Research* 41(1), 7–18.
- Hastie, T., R. Tibshirani, and J. Friedman. (2001). *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. New York: Springer-Verlag.
- Iacobucci, D., P. Arabie, and A. V. Bodapati. (2000). "Recommendation Agents on the Internet," *Journal of Interactive Marketing* 14(3).

- Jain, D. and S. S. Singh. (2002). "Customer Lifetime Value Research in Marketing: A Review and Future Directions," *Journal of Interactive Marketing* 16(Spring), 34–46.
- Kamakura, Wagner A., S. Ramaswami, and R. Srivastava. (1991). "Applying Latent Trait Analysis in the Evaluation of Prospects for Cross-Selling of Financial Services," *International Journal of Research in Marketing* 8, 329–349.
- Kamakura, Wagner A. and Michel Wedel. (1997). "Statistical Datafusion for Cross-Tabulation," *Journal of Marketing Research* 34(November), 485–498.
- Kamakura, Wagner A. and Michel Wedel. (2000). "Factor Analysis and Missing Data," *Journal of Marketing Research* 37, 490–498.
- Kamakura, Wagner A. and Michel Wedel. (2003). "List Augmentation with Model Based Multiple Imputation: A Case Study Using a Mixed-Outcome Factor Model," *Statistica Neerlandica* 57(1), 46–57.
- Kamakura, Wagner A, Michel Wedel, Fernando de Rosa, and Jose A. Mazzon. (2003). "Cross-Selling Through Database Marketing: A Mixed Data Factor Analyzer for Data Augmentation and Prediction," *International Journal of Research in Marketing* 20, 45–65.
- King, Gary and Langche Zeng. (2001). "Logistic Regression in Rare Events Data," *Political Analysis* 9(2), 137–163.
- Kivetz, Ran. (2003). "The Effects of Effort and Intrinsic Motivation on Risky Choice," *Marketing Science* 22, 477–502.
- Kivetz, R. and I. Simonson. (2002). "Earning the Right to Indulge: Effort as a Determinant of Customer Preferences Toward Frequency Program Rewards," *Journal of Marketing Research* 39(2), 155–170.
- Kivetz, Ran, Oleg Urminsky, and Yuhuang Zheng. (2004). "Goal-Motivated Purchase Acceleration: Evidence and Consequences in Reward Programs," Working Paper Columbia University.
- Knott, Aaron, Andrew Hayes, and Scott A. Neslin. (2002). "Nest-Product-to-Buy Models for Cross-Selling Applications," *Journal of Interactive Marketing* 16(3), 59–75.
- Kopalle and Neslin. (2003). "The Economic Viability of Frequency Reward Programs in a Strategic Competitive Environment," *Review of Marketing Science* Vol. 1, Article 1.
- Lal, Rajiv and David E. Bell. (2003). "The Impact of Frequent Shopper Programs in Grocery Retailing," *Quantitative Marketing and Economics* 1(2), 179–202.
- Leenheer, Jorna, Tammo H. A. Bijmolt, Harald J. van Heerde, and Ale Smidts. (2004). "Do Loyalty Programs Enhance Behavioral Loyalty? A Market-Wide Analysis Accounting for Endogeneity," Working Paper, Tilburg University.
- Lemmens, Aurélie and Christophe Croux. (2003). "Bagging and Boosting Classification Trees to Predict Churn," Working Paper, Teradata center.
- Lewis, Michael. (2003). "Customer Lifetime Value as a Function of Acquisition Discount: An Analysis of Promotionally and Regularly Acquired Customers," Working Paper, University of Florida.
- Li, K. C. (1991). "Sliced Inverse Regression for Dimension Reduction, With Discussions," *Journal of the American Statistical Association* 86, 316–342.
- Li, Shibo, Baohong Sun, and Rong Zhou. (2004). "Learning and Optimal Matching of Customers and Service Providers," Working Paper, Carnegie Mellon University.
- Li, Shibo, Baohong Sun, and Ron Wilcox. (2004). "Cross-selling Naturally Ordered Products: An Application to Consumer Banking Services," Working Paper, U. Virginia.
- Mittal, Vikas and Wagner A. Kamakura. (2001). "Satisfaction, Repurchase Intent and Repurchase Behavior: Investigating the Moderating Effect of Customer Characteristics," *Journal of Marketing Research* 38(February), 131–142.
- Naik, Prasad A. and Chih-Ling Tsai. (2001). "Single-Index Model Selections," *Biometrika* 88(3), 821–832.
- Naik, Prasad A. and Chih-Ling Tsai. (2004). "Residual Information Criterion for Single-Index Model Selections," *Journal of Nonparametric Statistics* 16(1–2), 187–197.
- Naik, Prasad A. and Chih-Ling Tsai. (2004). "Isotonic Single-Index Model for Database Marketing," *Computational Statistics and Data Analysis*, 47, 775–790.
- Naik, Prasad A. and Chih-Ling Tsai. (2005). "Constrained Inverse Regression for Incorporating Prior Information," *Journal of the American Statistical Association* 100(469), 204–211.

- Neslin, S. S. Gupta, W. Kamakura, J. Lu, and C. Mason. (2004). "Defection Detection: Improving Predictive Accuracy of Customer Churn Models," Working Paper, Teradata Center at Duke University.
- Rassler, Suzanne. (2002). *Statistical Matching*. New York: Springer.
- Reinartz, W. J. and V. Kumar. (2003). "Customer Lifetime Duration: An Empirical Framework for Measurement and Explanation," *Journal of Marketing* 67(January), 77–99.
- Rust, Roland and Anthony Zahorik. (1993). "Customer Satisfaction, Customer Retention, and Market Share," *Journal of Retailing* 69(2).
- Schmittlein, David, Donald Morrison, and Richard Colombo. (1987). "Counting Your Customers: Who Are They and What Will They Do Next?" *Management Science* 33(1).
- Sharp Byron and Anne Sharp. (1997). "Loyalty Programs and Their Impact on Repeat-Purchase Loyalty patterns," *International Journal of Research in Marketing* 14, 473–86.
- Shi, Mengze and Fusun Gonul. (1998). "Optimal Mailing of Catalogs: A New Methodology Using Estimable Structural Dynamic Programming Models," *Management Science* 44(9), 1249–1262.
- Simester, Duncan, Peng Sun, and John N. Tsitsiklis. (2004). "Dynamic Catalog Mailing Policies," Working Paper, MIT.
- Steenburg, Thomas J., Andrew Ainsle, and Peder Hans Engbretson. (2003). "Massively Categorical Variables, Revealing the Information in ZIP-Codes," *Marketing Science* 22(1), 40–57.
- Thomas, J. S. (2001). "A Methodology for Linking Customer Acquisition to Customer Retention," *Journal of Marketing Research* 38, 262–268.
- Verhoef, Peter C. (2003). "Understanding the Effect of CRM Efforts on Customer Retention and Customer Share Development," *Journal of Marketing* 67(4), 30–45.
- Venkatesan, R. and V. Kumar. (2003). "Using Customer Lifetime Value in Customer Selection and Resource Allocation," MSI Report 03–112.
- Verhoef, Peter C. and Bas Donkers. (2005). "The Effect of Acquisition Channels on Customer Retention and Cross-Buying," *Journal of Interactive Marketing* (forthcoming).
- Ying, Yuanping, Fred Feinberg, and Michel Wedel. (2004). "Improving Online Product Recommendations by Including Nonrated Items," Working Paper, University of Michigan Business School.
- Zhang, Jie and Michel Wedel. (2004). "The Effectiveness of Customized Promotions in Online and Offline Stores," Working Paper, University of Michigan Business School.