

MKTG 476/776, STAT 476:
Applied Probability Models in Marketing
Spring 2013

Professor Peter S. Fader
faderp@wharton.upenn.edu

Monday 3-6PM (undergrads)
Wednesday 3-6PM (grads)

Teaching assistants:

Grads – Arun Gopalakrishnan (agop@wharton.upenn.edu)

Undergrads – Ramiro del Valle (djesus@wharton.upenn.edu)

Motivations and Objectives

Over the past five decades, statisticians have developed a number of models that have proven to be highly effective in their ability to explain and predict empirical patterns within many areas in business and the social sciences. These models use some basic “building blocks” from probability theory to offer behaviorally plausible perspectives on different types of timing, counting, and choice processes. Researchers in marketing have actively contributed to (and benefited from) these models for a wide variety of applications, such as new product sales forecasting, analyses of media usage, targeted marketing programs, estimation of “customer lifetime value,” and even overall corporate valuation. Other disciplines have seen equally broad utilization of these techniques.

As new forms of information technology provide increasingly rich descriptions of individual-level shopping/purchasing behavior, these models offer great value to practicing managers, particularly those interested in pursuing CRM (“customer relationship management”) activities. Furthermore, as more managers become comfortable with non-linear optimization techniques (using, for example, the “Solver” feature within Microsoft Excel), the specification and interpretation of these models can become a regular part of the sophisticated manager’s toolkit. Taken as a whole, the methodological approaches covered in this course are well-suited to address the types of questions that are being asked with increasing frequency and interest by investors and managers of today’s data-intensive businesses.

The principal objectives of this course are:

- To familiarize students with probability models and their role in marketing, information systems, supply chain management, actuarial science, operations research, and other related areas,
- To provide students with the analytical and empirical skills required to develop probability models and apply them to problems of genuine managerial interest.
- To have students develop good instincts to judge the appropriateness, performance, and value of different kinds of models in a variety of managerial settings.
- To encourage students to think critically about statistical methods and managerial perspectives that are common in certain domains but not always the best ways to approach all data-oriented decision problems.

Prerequisites

This course is open to students at any level (undergraduate, MBA, PhD) who have sufficient mathematical skills to handle the advanced methods that will be introduced and featured here. It is essential that students retain some familiarity with basic integral calculus. Furthermore, a mid-level probability/statistics course would be helpful. But aptitude to learn and fully understand these methods is far more important than mere exposure to them.

Course Organization and Materials

Every session will be lecture-based, with a strong emphasis on real-time problem solving, including mathematical derivations and numerical investigations using Microsoft Excel. Central to the development of the skills associated with probability modeling is hands-on experience. To this end, a set of homework exercises will be assigned for most sessions.

There is no formal textbook for the course (since no suitable book exists), but lecture notes covering most of the material presented in class will be made available immediately after each session. Some (but not all) of the Excel spreadsheets used in class will be made available to the students, and some journal articles will be suggested as illustrations/applications of the techniques discussed. While it is expected that students will read and review all of these materials thoroughly, there will be no pre-class readings assigned for most sessions.

Teaching Approach

The methods covered in this course will be largely unfamiliar to most students. As such, it is important to ensure that the first exposure is impactful and that there are opportunities to work with the materials multiple times and through multiple formats. To address these issues we will utilize a “heads up” learning system in the classroom. The basic elements include:

- Mandatory classroom attendance
- The use of laptops in the classroom is discouraged
- Presentation decks will not be provided until after class
- Each session will be recorded and made available to students (in a unique multimedia format) soon after each class session

These steps are intended to help students keep their “heads up” to focus on the main points in each session. Students are encouraged to ask questions about key conceptual issues, managerial applications, and the overall modeling philosophy; questions about more micro/technical issues can be resolved by reviewing the presentation decks and recordings after class.

Students are expected to create their own complete set of class notes after attending each session and working through the decks/recordings. It is fine for students to collaborate on this task, but it's best for each student to create his/her own notes. Any kind of “divide and conquer” approach will be counterproductive for the student.

Evaluation

Homework Exercises: These exercises will be both analytical and numerical in nature. All of the numerical work can be completed using Excel. One or more homework problems will likely be basis of questions that will be included on the final exam. It is OK for students to communicate about the homework exercises, but each student must write up each problem independently and hand in his/her own work.

Class Participation: Although there are no formal case discussions, students are expected to be actively engaged in the lectures, which will include occasional “cold calls” for students to provide solutions for problems discussed in class.

Final Exam: The final exam will be a structured set of questions to assess students’ conceptual understanding of the course material. It will not require any detailed mathematical derivations or extensive numerical calculations. It will be a closed-book exam but students can bring a one-page “cheat sheet” as a reference.

Paper #1: For the first paper, students will be asked to find a specific type of dataset and analyze it carefully (details will be discussed in class).

Paper #2: The second paper will be more standardized – all students will be given a common dataset to analyze.

Paper #3 (optional): The third paper will be more open-ended with four basic options. Students can: (1) do a detailed analysis of a rich dataset (more complex than those covered in the earlier papers); (2) develop and apply a new probability model; (3) carry out an extensive simulation exercise to explore the properties of one or more models covered in class; or (4) conduct a comprehensive literature review of one application area of probability models in marketing. Generous extensions will be granted to students who wish to complete the paper after the term ends. More details and possible topics will be discussed during the term.

All relevant University of Pennsylvania policies regarding academic integrity must be followed. Students may not submit work that has been prepared by (or in conjunction with) someone else, without explicit instructor permission. Any student who in any way misrepresents somebody else's work as their own will face severe disciplinary consequences.

Grading

Students will follow one of two options, depending on whether they choose to do Paper #3:

	<i>Option 1</i>	<i>Option 2</i>
Homework	10%	8%
Class Participation	10	8
Final Exam	30	24
Paper #1	25	20
Paper #2	25	20
Paper #3	--	20

More details about each assignment will be discussed in class.

Class Scheduling Issues – PLEASE READ CAREFULLY!

There are two sections of the course: one for undergraduate students on Mondays and one for graduate students on Wednesdays. The same basic material will be covered in each one. Both sessions will be recorded each week, and students are welcome (but certainly not expected) to examine both videos if they wish to gain additional perspectives on any session.

This semester's schedule is awkward due to the unusual timing of the Martin Luther King, Jr. Day observance. Because we want to keep both sections operating in parallel, we will use the following scheduling plan:

- Week 1 (W 1/9 only): The initial grad session.
- Week “1A”: On Monday 1/14, we will repeat the “Week 1” material for the undergrads. (And grads who can't make it to the 1/9 session are also welcome to attend.)
- On Wednesday 1/16, we will hold a “bonus” session for the grads, which will cover some interesting extensions to the “Week 1” material. Undergrads are encouraged to attend (or watch the recorded version) as well.
- Week 2 (W 1/23 only): Due to the holiday we will only meet on Wednesday, but this material is required for the undergrads as well. Again, they can come to the class if they can, but must watch the recorded version if they can't.
- Weeks 3-13: Normal schedule
- Week 14 (M 4/22 only): A different “bonus” session for the undergrads; grad students are invited to attend as well.

Course Schedule (M=undergrad section, W=grad section)

Week 1 (W 1/19, M 1/14): Introduction to probability models

Motivating problem: Forecasting customer retention. Comparisons to traditional regression-based models: “curve-fitting” vs. “model-building.” Careful derivation of a parametric mixture model (the beta-geometric). Coverage of maximum likelihood estimation and the Microsoft Excel Solver tool. Discussion about the philosophy and objectives of probability modeling.

Week “1A” (W 1/16): Bonus session: “The Paradox of Increasing Loyalty”

We further extend the model covered in Week 1 to demonstrate some unusual patterns about customer retention. All students from both sections are welcome to attend.

Week 2 (W 1/23 only): Models for count data

Introduction to the Poisson process and its extension to the negative binomial distribution. Evaluating goodness-of-fit. Generalizing the model to allow for “spikes” at 0 or 1.

Week 3 (M 1/28, W 1/30): More count models and repeated choice processes

Alternative estimation approaches for count models (“Means and zeroes” and “method of moments”). Dealing with problems of limited/missing data: truncated and shifted NBD models. Choice vs. counting. The binomial distribution. The beta distribution as a mixture model. Parameter estimation and inference.

Week 4 (M 2/4, W 2/6): Empirical Bayes methods

Conditional distributions and expectations. Combining population information (“priors”) with observed data for individuals. Regression-to-the-mean.

Week 5 (M 2/11, W 2/13): Timing models

Motivating problem: forecasting new product adoption. Implementing and evaluating different timing models, particularly the exponential-gamma. Dealing with grouped data and right censoring. Introducing hazard functions. Derivation and discussion of other timing models (e.g., Weibull), and the linkages among them. Exploring the interplay between timing and counting processes.

Week 6 (M 2/18, W 2/20): Customer-base analysis

Project #1 (counting model) due

Combining the basic building blocks to create integrated models to estimate customer lifetime value and related concepts.

Week 7 (M 2/25, W 2/27): Customer-base analysis (cont.)

More CLV-oriented applications.

Spring Break – no class (M 3/4, W 3/6)

Week 8 (M 3/11, W 3/13): Introducing covariates

Poisson regression and NBD regression for counting models. Beta-logistic (and alternative approaches) for choice models. Proportional hazard methods and covariate effects for timing models. General discussion about the different role of covariates from the perspective of an econometrician vis-à-vis a probability modeler. Applications.

Week 9 (M 3/18, W 3/20): Finite mixture and latent class methods

Looking at non-parametric (discrete) approaches to capturing heterogeneity. Interpreting support points versus cluster characteristics. Estimation issues. Overview of selection criteria for non-nested models.

Week 10 (M 3/25, W 3/27): Multi-item choice models

The multinomial choice process and the Dirichlet mixing distribution. Interplay between the beta and Dirichlet distributions.

Week 11 (M 4/1, W 4/3): Fun with Dirichlet!

Project #2 (covariate model) due

Further examination of the Dirichlet-multinomial choice model and its astonishing patterns. Discussion of Ehrenberg's "empirical laws."

Week 12 (M 4/8, W 4/10): Integrated models

Combined models of counting, timing, and/or choice. Particular focus on the BB/NBD as a working example.

Week 13 (M 4/15, W 4/17): Nonstationary processes

Overview and comparison of techniques such as renewal processes, learning models, hidden Markov methods, and other approaches to capture dynamics over time

Week 14 (M 4/22 only): Applications and wrap-up

More applications/extensions/speculations, etc.