

MKTG 562
Customer Analytics
Autumn 2022

Professor Simha Mummalaneni
simha@uw.edu
Paccar 487

Course Information

Class times (section A/B):	Monday & Wednesday 3:30pm - 5:20pm
Class times (section C/D):	Wednesday 6:00pm - 9:20pm
Class location:	Paccar 295
Optional lab sessions:	TBA
Additional office hours:	By appointment

Course Objectives

Customer Analytics addresses how to use data analytics to learn about and market to individual customers.

Marketing is evolving from an art to a science. Many firms have extensive information about consumers' choices and how they react to marketing campaigns, but few firms have the expertise to intelligently act on such information. In this course, students will learn the scientific approach to marketing with hands-on use of technologies such as databases, analytics, and computing systems to collect, analyze, and act on customer information. This is the key part of learning how to take advantage of Big Data. While students will employ quantitative methods in the course, the goal is not to produce experts in statistics; rather, students will gain the competency to interact with and manage a marketing analytics team.

The course uses a combination of lectures, cases, and case assignments to learn the material. Students will learn all stages of the analytics process: how to identify marketing problems that can be addressed with analytics, how to decide which analytical method is best suited for a specific marketing problem, how to conduct the analysis using modern-day statistical and machine learning methods in R, how to interpret the results, and how to use the analytical results to make better marketing decisions. This course takes a very hands-on approach with real-world databases and equips students with tools that can be used immediately on the job.

This course focuses specifically on how companies can use individual-level data to target customers, recommend products to them, predict their purchasing patterns, and retain them over the long run. Students who are interested in other marketing analytics topics such as automated product development, algorithmic price optimization, online advertising, platform design, and social media are encouraged to also take the other classes in the Foster marketing analytics specialization: *Analytics for Marketing Decisions*, *Digital Marketing Analytics*, and *Pricing Strategy and Analytics*.

Course Organization

This course is organized around core topics that are of interest to any marketing team:

Customer-centric marketing: This section explains why customer-level data is valuable to companies, and how recent advances in data collection have improved companies' ability to market to individual consumers.

Prospecting and targeting the right customers: In this section, we focus on customer acquisition. We discuss how to target customers, how to decide which ones to target, and how to choose an individualized offer for each customer.

Developing customers: This section focuses on how to improve the profit-per-customer. We discuss how to upsell customers, how to decide which new products to recommend them, and how to use customer ratings to inform these recommendations.

Retaining customers: Holding on to one's existing customers is a crucial task for any marketing team. In this section, we discuss how to predict which customers are most likely to leave, and how to develop an incentive plan that encourages them to stay.

Limitations of customer analytics: Analytics and data enable us to answer important marketing questions with a high level of rigor. However, analytics can also be misinterpreted or misused. In this section, we discuss how to avoid common mistakes in implementing customer analytics.

Analytical Methods

The business problems discussed in this course typically require us to make predictions about how individual customers will behave: predicting customer response to a specific offer, predicting how much a customer will like a specific product, predicting which customers are likely to cancel their service; etc.

We will model these prediction problems using a wide variety of modern analytical methods:

Heuristic models: RFM analysis

Statistical models: Linear regression and logit models

Machine learning & AI models: Neural networks, decision trees, and ensemble methods

Algorithmic models: Market basket analysis and recommendation systems

Although we will be focusing on using these methods in the context of customer analytics, these tools can be applied to many other data contexts as well.

Grading Policy and Assignments

The final grade in this course will be based on the following components:

30%	Individual homework assignments ($4 \times 10\%$ each; lowest gets dropped)
52%	Group homework assignments ($4 \times 13\%$ each)
5%	Individual reflections ($5 \times 1\%$ each)
13%	Class participation and group contribution

In all elective courses at Foster, the distribution of final grades must satisfy school-wide requirements. Grades will be curved accordingly at the end of the quarter.

Assignments:

Much of the learning during the course will happen with the help of individual and group homework assignments. For individual assignments, you are not allowed to work with other students – the write-up should reflect your own work only. For group assignments, you should work in groups of 2 - 4 students and only hand in one write-up

per group. Groups should be formed (self-selected) at the beginning of the quarter and remain constant for all assignments.

Fluency in analytics requires both an ability to analyze data and an ability to explain the findings to a managerial audience. The assignments in this course will typically be a combination of both skills: using R to analyze data, and then explaining what actions a company should take. A combination of programming skills, statistical knowledge, and managerial intuition will be required.

I will ask each of you to submit a confidential “Peer Rating Form” at the end of the quarter. Submitting this form is considered part of your class participation grade. If group members say that someone was exceptional and took a great leadership role in the group, that will help that student’s class participation and group contribution score. Conversely, if group members say that someone was irresponsible and didn’t do enough work, then that will decrease that student’s class participation and group contribution score.

We will discuss each homework assignment in class on the day that it is due. As a result, I cannot accept any late assignments.

Individual reflections:

During the course I will ask you to submit short reflections (max 1/2 page) about what you learned from a lecture and/or an assignment. These reflections are always individual work, even if the assignment you are asked to discuss was a group assignment.

Class participation and group contribution:

Quality contributions which are relevant to the discussion will improve your participation grade. I will call on students at random to open discussions about assignments and in-class examples. Your class participation score will be significantly hurt if you are asked to offer your analysis on a case or assignment question and you are not prepared.

The group contribution component of this grade will be determined in part by the confidential “Peer Rating Form” described above under the “Assignments” section.

Attendance and punctuality:

Learning to articulate your analysis and to evaluate and respond to the analysis of others is an important part of what you will learn in this class. If you miss class, you will miss this, and there is not a way to “make it up.” As a result, you should make every effort not to miss class.

If you miss class or are late, it will lower your class participation grade. If you must miss class, you should do the readings, prepare and turn in the assignments on time (late assignments will not be accepted), and arrange to get notes from a friend about what you missed in class.

Course Materials

You should prepare for each class session by watching the lecture videos and doing the reading. These materials are specified on Canvas.

Readings in this course consist of a mixture of cases and articles – no textbook is required.

Textbook: None required

Course pack: None required

Lecture videos: Distributed on Canvas

Lecture notes: Distributed on Canvas

Articles and additional course material: Distributed on Canvas

The syllabus, lecture videos, lecture notes, and additional course material are posted on Canvas. All material is organized by the class number it belongs to. Lecture notes will be posted on Canvas on the day of the lecture.

Software

This course uses the statistics program “R” during the course. R is a very widely used program, and learning how to conduct analysis in R is a major benefit from this course. R is free to use, and has a strong community of users who provide online help and support.

Learning a new programming language can be a challenging task. If you want additional support, I can help with troubleshooting or basic R advice. If you want to learn more about R or want a textbook to use as a reference, I will be happy to provide some suggestions.

Honor Code

Students are expected to respect the Foster code of conduct at all times.

Written assignments will be either individual assignments or group assignments, as specified on Canvas. For individual assignments, you are not allowed to work with other students – the write-up should reflect your own work alone.

It will be a violation of academic integrity if you base any of your assignments on discussions or solutions that you find online or that are shared by people who have taken similar courses in the past. I reserve the right to fail you for the course if I become aware of such a violation. To ensure that future classes benefit from a similar learning experience, do not share any case or assignment materials with students outside the class.

Religious Accommodations

Washington state law requires that UW develop a policy for accommodation of student absences or significant hardship due to reasons of faith or conscience, or for organized religious activities. The UW's policy, including more information about how to request an accommodation, is available at [Religious Accommodations Policy](https://registrar.washington.edu/staffandfaculty/religious-accommodations-policy/) (<https://registrar.washington.edu/staffandfaculty/religious-accommodations-policy/>).

Accommodations must be requested within the first two weeks of this course using the [Religious Accommodations Request form](https://registrar.washington.edu/students/religious-accommodations-request/) (<https://registrar.washington.edu/students/religious-accommodations-request/>).

Course Schedule and Assignment Due Dates

Section A/B

Date	Class #	Class Topic	Assignment Due
Sep 28	1	Course overview Customer analytics and centrality	
Oct 3	2	Quantifying customer value Using R for customer analytics	R first steps
Oct 5	3	How to distinguish good from bad analytics	
Oct 10	4	Case analysis: Col Solare part 1 Testing and experimentation	Col Solare part 1
Oct 12	5	Case analysis: Goop Predicting response with RFM analysis	Goop
Oct 17	6	Statistics review	
Oct 19	7	Case analysis: Lyft Lift and Gains	Lyft
Oct 24	8	Predicting response with logistic regression	<i>Reflection</i> on RFM
Oct 26	9	Interpreting interaction effects Predicting response with neural networks	
Oct 31	10	Case analysis: Col Solare part 2	Col Solare part 2
Nov 2	11	Evaluating model performance	<i>Reflection</i> on logistic regression
Nov 7	12	Predicting response with decision trees	
Nov 9	13	Case analysis: Coca-Cola	Coca-Cola
Nov 14	14	Next-product-to-buy models	<i>Reflection</i> on Coca-Cola case
Nov 16	15	Recommendation systems	
Nov 21	16	Case analysis: Neiman Marcus	Neiman Marcus
Nov 23		NO CLASS — Day before Thanksgiving	
Nov 28	17	Predicting and managing attrition	<i>Reflection</i> on NPTB models
Nov 30	18	When customer analytics fails	
Dec 5	19	Case analysis: Xbox Live From prediction to prescription	Xbox Live
Dec 7	20	Course wrap-up	<i>Reflection</i> on course

Course Schedule and Assignment Due Dates

Section C/D

Date	Class #	Class Topic	Assignment Due
Sep 28	1	Course overview Customer analytics and centrality How to distinguish good from bad analytics	
Oct 5	2	Quantifying customer value Using R for customer analytics Testing and experimentation	R first steps
Oct 12	3	Case analysis: Goop Predicting response with RFM analysis Statistics review	Goop Col Solare part 1
Oct 19	4	Case analysis: Lyft Lift and Gains Predicting response with logistic regression	Lyft
Oct 26	5	Interpreting interaction effects Predicting response with neural networks Case analysis: Col Solare part 2	<i>Reflection</i> on RFM Col Solare part 2
Nov 2	6	Evaluating model performance Predicting response with decision trees	<i>Reflection</i> on logistic regression
Nov 9	7	Case analysis: Coca-Cola Next-product-to-buy models	Coca-Cola
Nov 16	8	Recommendation systems Case analysis: Neiman Marcus	<i>Reflection</i> on Coca-Cola case Neiman Marcus
Nov 23		NO CLASS — Day before Thanksgiving	
Nov 30	9	Predicting and managing attrition When customer analytics fails	<i>Reflection</i> on NPTB models
Dec 7	10	Case analysis: Xbox Live From prediction to prescription Course wrap-up	Xbox Live <i>Reflection</i> on course