

Teacher Resources

Marketing Science¹

Marketing as defined by the American Marketing Association is “the process of planning and executing the conception, pricing, promotion, and distribution of goods, and services to create exchanges that satisfy individual and organizational objectives.” Since the 1960s, operations researchers have been developing and applying a wide range of OR modeling techniques to assist decision makers in every aspect of the marketing function. One of the first challenges every company faces is to forecast demand for an existing product or one considered for development. This is usually done with sophisticated statistical models. These models forecast overall demand as well as demand within various market segments that are defined by demographic and lifestyle characteristics.

Operations research models help answer a related question: What characteristics should a product have so as to maximize demand or net profit? Companies, however, generally offer more than one product or service and thus the decision maker must plan for a portfolio of products or services targeted at different market segments. Of specific concern is the overlap between products that can lead to cannibalization of one product by another. This problem arises when a company such as General Motors offers an array of vehicles or Hewlett-Packard offers a variety of printers. This problem occurs sequentially when a company such as Intel has to decide when to launch each new, more powerful chip and then decide how to price it as compared to older, less powerful ones.

The biggest change in the marketing function is the availability of vast amounts of data about potential customers starting with ubiquitous data captured by supermarket and retail store scanners. This had led to a specialization known as database marketing. Sophisticated data mining techniques have been developed to cull critical insights from these vast datasets. One conceptually simple concept, scoring, is used to target individuals or segments for different promotions based on past behavior and demographics. Almost all of us interact with these scoring models when we seek credit or a loan since banks routinely score individuals as to their credit worthiness. The Internet has added an entirely new dimension of information overload, as every visit to a website, especially a purchase, is tracked and analyzed to develop a profile of each individual web browser. Not surprisingly, the financial services industry and catalogue based retailers are among the leading users of marketing science.

Operations researchers also model the interaction between the marketing function and other management activities. There is a strong link between demand forecasts and strategies for scheduling production and controlling inventory levels. Decisions on the magnitude, frequency, and duration of special price promotions must be coordinated with these other functions. This coordination becomes even more complex in a global company. The most difficult element to factor into these models is the impact of competition. As information flows ever more rapidly and openly, competitors can quickly react to aggressive marketing actions.

¹ Abstracted from Jehoshua Eliashberg, Gary L. Lilien, and Yoram Wind. Marketing. In Saul Gass and Carl M. Harris, eds., *Encyclopedia of Operations Research and Management Science, Second edition*, (Boston: Kluwer Academic Publishers, 2001)

Definitions

Direct marketing. Any marketing activity in which a business attempts to reach consumers directly, or have consumers reach the business directly, is an example of direct marketing. Specific examples of such direct communication include:

- 1.) A print media ad with a coupon, order form, or telephone number.
- 2.) A direct mail piece, e-mail, pop-up or catalogue sent to households and designed to sell products or services.
- 3.) A television or radio ad featuring a telephone number to call for more information or to place an order.

Database. A list of information organized in some manner.

Marketing database. A database containing information about customers (e.g., name, address, demographics, purchase history) which is used to target customers (or prospective customers) with the goal of increasing customer loyalty and repeat sales.

Case Studies

CDNOW. Launching a company into a new sales environment involves unique challenges as to forecasting sales, markets, and buyer behavior. The music retailer CDNOW set up a business for selling music CDs over the Internet and needed a reliable tool to develop profiles for potential customers, as well as predict CD sales to a cohort of new customers.

Most music stores are located in specific areas and the cross section of customers in that particular locale can be readily assessed. The Internet does not have the boundaries of a brick and block store, so that the customer base must be interpreted through the number of people who visit the site, the number of people who make a purchase, the number of purchases each customer makes, and so forth.

Modeling tools already existed, but they were complex, and required an analyst with experience in managing large data bases. CDNOW was unwilling to invest in the high cost of developing and applying such a complex model. Instead, CDNOW sought a new and simpler model that merely allowed for data entry of total purchases and the number of new customers each week and that could be run in a standard spreadsheet environment such as EXCEL. This would allow management to review the results of the model frequently and make business judgments as the market changed.

The study was conducted in the first quarter of 1997. The model developed was a simple stochastic model of buyer behavior capable of giving managers a medium range tool for forecasting CD sales. Working with the large database, the number of customers each week could be entered, along with the number of units each customer purchased. Then, using an aggregate model, it was found that reasonable forecasts of medium-term sales could be made.

For the cohort of 23,570 people who first purchased from CDNOW in the first quarter of 1997, they extracted records of the total purchases for the 40 weeks beyond the calibration period (that is, April through December 1997). The forecasting performance of the model was evaluated against the actual purchasing numbers. The performance of the model is clearly demonstrated by the fact that the model underpredicts the cohort's first-year purchasing by less than two percent. In contrast, an alternative linear projection model overpredicts the cohort's first-year purchasing by more than 40 percent. This enabled CDNOW to design a reasonable business plan for their company which was not possible to do with the poorly calibrated regression model.

Fader, P. S. & Hardie, B. G. S. (2001). Forecasting repeat sales at CDNOW: A case study. *Interfaces*, 31(3), S94-S107.

Fingerhut. Fingerhut, a mail-order direct marketer, engaged the help of IBM in decreasing 340 million catalog mailings to its 7 million customers annually, because many of the catalogs were redundant and unproductive. Together, Fingerhut and IBM formed a partnership to meet Fingerhut's business need. Fingerhut supplied the statisticians, direct-marketing expertise, and programming resources, while IBM brought the horizontal-marketing concepts and mathematical-optimization expertise.

Traditionally, Fingerhut followed the standard marketing approach to managing customer contacts, choosing the best customers for each catalog. This vertical marketing process led to cannibalization or saturation between catalogs. The team developed a customer-selection system called mail-stream optimization (MSO), which focuses on the customer, not the catalog. The MSO system selected the most profitable sequence of catalogs for each customer.

Customers were partitioned into 100 micro classes. This was done based on long-term customer value and length of time as a customer. Statistical prediction models were devised to estimate customer value and advertising productivity over the next 12 months.

The advertising-allocation model used risk/return curves, noting the slope changes to make investment decisions. The curves describe the relationship between the advertising spending (risk) and the revenue (net returns) attained through that spending. The challenge was to invest in advertising up to the point where any continued spending yielded little, if any, profit. These risk/return curves were piecewise linear curves.

Matrices representing the similarity of merchandise, content, and features of various catalogs were used to model saturation of catalogs. A series of regression models used merchandise-similarity measures and promotional- and presentation-similarity measures as independent variables. The dependent variable came from Fingerhut's data warehouse.

The MSO System had several phases: Customer prediction, advertising allocation, customer-profit prediction, catalog saturation and its reduction, and then finally optimization. In the optimization phase, Fingerhut assigns the most profitable set of mail streams to customers while enforcing a variety of budget, management, and technical constraints. Optimization at Fingerhut had four steps:

1. Cluster customers to produce a more easily managed optimization problem.
2. Solve an optimization problem for each cluster, generating potential mail streams, given advertising constraints.
3. Solve a global optimization problem to find the best allocation of mail streams across all customer clusters.
4. Translate the cluster-level solution to the customer level.

The MSO System uses concepts in systems of linear equations, linear programming, linear regression, and organizational matrices. The project paid for itself within the first year, generating an annual profit-gain of \$3.5 million. The MSO system is now being applied to prospective and inactive customers. In the future, it will be translated for use in other aspects of the business.

Campbell, D., Erdahl, R., Johnson, D., Bibelnieks, E., Haydock, M., Bullock, M., & Crowder, H. (2001). Optimizing customer mail streams at Fingerhut. *Interfaces*, 31(01), 77 – 90.

Merrill Lynch. Companies that market their goods or services directly to the public often engage in prospecting. Just as prospectors from the Nineteenth Century mined for gold, these companies “mine” databases containing information on prospective customers for “gold” in the form of new customers.

Merrill Lynch applied various multivariate statistical and analytic techniques, including regression analysis, to develop a database marketing program. A summary of the particular approach used follows:

1. Define and identify the target population.
2. Prepare two random samples: one of known members of the target population and the other of members of the potential prospect universe.
3. Append appropriate demographic and behavioral data to both samples.
4. Conduct a statistical analysis of the data to develop an analytical model that evaluates and prioritizes the prospect universe in terms of its similarity to the target population.
5. Conduct a field test to measure actual performance.

Previous statistical studies at Merrill Lynch had yielded 85-90% accuracy in determining the characteristics of priority customers. These results suggested the potential of the database marketing approach to mining the prospect universe for potential customers. Merrill Lynch conducted a field test with limited usage of a newly selected prospect universe in order to minimize field impact until the benefits of the new database marketing approach could be demonstrated.

A model was developed that generated “scores” for each potential client. The scores were then used to prioritize the list of prospects. A cut-off score for qualifying prospects was established based on the probability that the potential customer would be a priority client.

Merrill Lynch estimates that the impact of the new approach will be between \$3.5 and \$6 billion in additional client assets. The estimates for production credits ranged between \$200,000 and

\$400,000. Production credits are used throughout the brokerage industry to measure commission-related revenue. The success of the database marketing approach at the general business level has led to interest at Merrill Lynch in developing customized models to identify the best prospects for specific investment products. Research efforts are currently underway to develop specialized prospecting models for selected Merrill Lynch products.

Labe, R. P., Jr. (1994). Database marketing increases prospecting effectiveness at Merrill Lynch. *Interfaces*, 24(5), 1-12.

OnStar. A multimethod approach was developed to evaluate strategic alternatives for General Motors' OnStar communications system. OnStar is GM's two-way vehicle communication system that provides a variety of services that enhance safety, security, entertainment, and productivity.

First GM used dynamic modeling to address decisions the company faced in 1997. The dynamic modeling was used to address the choice between incremental and aggressive marketing strategies for OnStar. GM had to decide whether to view OnStar as a car feature or as a service business, and then how to market the car feature. A project team was formed to determine if an evolutionary or a revolutionary strategy should be chosen. At each step, the team interacted with the decision board that was responsible for actually making the decision and committing resources.

The OnStar case was difficult to model because no one had experience or historical data in 1997. An integrated simulation model was used for analyzing the new telematics market. This modeling process would allow integration of various methods and data sources. Concepts and techniques from several management sciences approaches were used. They predicted that no other vehicle manufacturer would launch a competitive system before 2003 and they were correct.

The simulation model was the core tool in the OnStar strategy project. Its six key sectors were (1) customer acquisition and retention, (2) customer choice, (3) alliances and competition, (4) customer service, (5) financial costs and dynamics, and (6) dealer behavior. Over time the sectors interact to generate time series for such important business variables as market share and cash flow.

GM's modeling effort had important financial, organizational, and societal results. The OnStar business in 2002 had two million subscribers, an 80 percent share of the emerging telematics market, and was valued at between \$4 and \$10 billion. The OnStar project set the stage for a broader GM initiative in service businesses that ultimately could yield billions in incremental earnings. Finally, and most importantly, the OnStar project has saved lives.

Barabba, V., Huber, C., Cooke, F., Pudar, N., & Paich, M. (2002). A multimethod approach for creating new business models: The General Motors OnStar Project, *Interfaces*, 32(1), 20-34.

SilverScreener. Every Monday morning, the management of Pathé Theaters in the Netherlands decides which movies in its cinemas to retain and which to replace. It must choose replacement movies from those available at that time. The SilverScreener model, a mathematical

programming system, was implemented to help Pathé managers make those decisions for one six-screen theater.

SilverScreener is a marketing management support system developed in 1999 by Swami, Eliashberg, and Weinberg. It is designed to provide exhibitors of movies an effective attendance forecasting system. ¶ A theater owner trying to manage screens well faces complexity from various sources:

1. There is a large number of movies to compete for the few screens available.
2. As distributors release new movies each week, they pressure the owners to provide screens and play time for them.
3. Theater owners often base a number of theaters in the same geographical area and must manage the interdependency among several facilities.
4. The distributor/exhibitor contract imposes complicating constraints, such as obligating the exhibitor to play the film for a certain period of time even when audience demand is weak.

In addition, the contractual agreements between movie distributors and exhibitors are unique to the industry. They include a sliding scale system so that over the number of weeks that a movie plays at a theater, the initial advantage in the distribution of revenue is allocated to the distributor, but as the weeks continue, the advantage gradually shifts to the exhibitor. Keeping the bottom line in mind, carefully selecting the movies to be shown is essential for a theater to be profitable. Left to their own devices, most managers make such decisions using their intuition and do not rely on any data sources.

An attendance forecasting system was developed using Pathé's historical data, managerial judgment, and theater specific factors. The performance of the test theater was measured against the performance of two unaided similar multi-screen cinemas. The revenues at the theater using SilverScreener recommendations were higher than those at the two comparison theaters. Managerial attitudes towards the modeling system improved after implementation of SilverScreener.

Although there were other models in place, they had proved very complex from an analysis perspective. The average movie theater manager was not statistically sophisticated enough to either input the data or to interpret the results for accurate decision-making. The results from the experimental implementation of the SilverScreener model were comparable to the results from the much more complex models already in place. The key personnel were comfortable using the SilverScreener model and found it to be a useful tool.

Eliashberg, J., Swami, S., Weinberg, C. B., & Wierenga, B. (2001). Implementing and evaluating SilverScreener: A marketing management support system for movie exhibitors. *Interfaces*, 31(3), S108–S127.

Objectives of the Module

This module is intended for students in an Advanced Placement Statistics course or any high school course that addresses inferential statistics. In this module, students will explore the concepts of direct marketing, a marketing database, and a sales promotion. They will use statistical techniques to make decisions about the promotion and predict the amount of net profit it is likely to generate. In particular, students will:

1. Analyze information drawn from a customer database to assess the feasibility of a direct sales promotion.
2. Construct confidence intervals for a population proportion and for the difference between two such proportions.
3. Construct hypotheses and test them using:
 - a.) a large sample test for the difference of two population proportions.
 - b.) a large sample test for the difference of two means.
 - c.) a Chi-Square test for the homogeneity of population proportions.
4. Make a distinction between statistical significance and practical significance.

Initiating Activity

Open a dialogue with students aimed at getting them to think about marketing (especially via mail, telephone or the Internet). Some questions you might ask include the following:

- Have you or any family member ever been solicited to purchase a product by mail, telephone or e-mail?
- How do you think you were selected? How did they get your address, telephone number or e-mail address?
- Have you or a member of your family ever purchased an item by mail order, telephone order or over the Internet?
- If you or a family member ever did make such a purchase, what sort of information were you required to provide to the company?
- Did any of the information seem unrelated to the purchase? If so, why do you think they asked for it?
- Have you ever completed a product registration form? If so, what personal information did they ask for?
- How do companies keep a record of all of this information?
- What is a database?
- What do you think a marketing database is?

Teaching Notes – Part I

2. Some students may realize that this question requires an assumption about the number of games each customer purchases. In that case, you might suggest assuming that each customer purchases exactly one game, or you might use the profit value per customer data immediately following question 5 to calculate the average number of games purchased per customer: $\$33.10 \text{ profit per customer} / \$30 \text{ profit per game} \approx 1.1 \text{ games/customer}$. (You may want to ask students why this average need not be an integer.) When choosing the latter option, you could skip immediately to question 6.

The optional direction between questions 9 and 10 sets up the use of a spreadsheet for the later computations, if you prefer.

Whenever you repeatedly determine confidence intervals for different elements of the same variable, the individual value of p used does not reflect the overall level of significance of all your calculations combined. In addition, you have introduced a subtle bias by selecting the two extreme values for a comparison of significance. If you are searching for statistically significant variations among the six profit values, a more advanced statistical tool called one-way, analysis of variance (ANOVA) is more appropriate.

11. In questions 11-14, we want students to begin to think about the distinction between statistical significance and practical significance. Many times, a small difference may be statistically significant due to the large size of the samples involved. However, it remains a case-specific matter of opinion as to whether a small difference is of any practical significance.

Part II

3. Of course, eyeballing the data is no substitute for a formal test of significance. The question of statistically significant differences in response rates across segments will be addressed later in questions 8-12.
5. You may want to ask why the Z-test was performed at the 95% level of confidence, while in question 4, a 90% confidence interval was constructed.
8. This item should clearly demonstrate the difference between statistical and practical significance.
10. The Chi-square test of association is used here, because the expected proportions (response rates) must be estimated from the data under analysis. In order to properly apply the Chi-square goodness-of-fit test, it is necessary that the expected proportions be calculated *a priori* on the basis of theory.
12. & 13. One way to interpret degrees of freedom is as the number of independent pieces of information under analysis. For the 2x6 tables under analysis here, the information contained

in each of the 12 cells is not independent information, because the row and column sums are known. If we know the values for the right five cells, all of the remaining cell values may be computed. (Students who have trouble seeing this might be encouraged to try it out for themselves.) Thus, the minimum number of cell values needed in order to complete the table is 10, and that determines the number of degrees of freedom.

18. You might want to ask students why it makes sense to omit the two test samples from this final analysis.

Project Ideas

1. Talk to a manager at your local supermarket. and ask the following:
 - a.) Do they send out direct mail advertisements?
 - b.) If so, how do they select the recipients?
 - c.) How do they use customer purchase data?
 - d.) Do they give out store coupons at check out?
 - e.) If so, how is the particular coupon determined?

2.
 - a.) Draw a random sample consisting of 1% of the seniors in your high school.
 - b.) Draw a random sample of 1% of all of the students in your high school.
 - c.) If a random sample of 1% of the residents of your hometown were drawn, how many people would that be? Is drawing such a sample feasible?
 - d.) Explain why a large direct marketer such as Lands' End would probably not select a random sample consisting of 1% of its customer base.
 - e.) If Lands' End drew a random sample of 1,000 customers from its customer data base, estimate the percentage of the customer base sampled.

3. Contact your local Republican or Democratic party and ask how they make up their calling lists for Election Day.

Extension 1: Random Samples

In statistical studies such as the test samples of the promotion with and without an incentive that were carried out by the Marketing department at Gamz, Inc., it is important that the samples used be representative of the populations from which they are drawn. However, the question of representativeness leads to a dilemma. How can we be certain that any sample is representative of a population without knowing the characteristics of the population that the sample is intended to estimate?

We can overcome this quandary once we realize that, in the long run, samples drawn randomly from the population will be representative of the population. But what makes any sample a *random* sample? A random sample is any sample drawn from a population in such a way that every member of the population has the same likelihood of being included in the sample. This process of randomization removes any possible sampling bias and yields a sample that can be considered representative of the population.

1. Describe one way in which a random sample of 1,000 customers could have been drawn from the Gamz, Inc. database of 50,000 customers.
2. Obtain the Gamz, Inc. database as an Excel spreadsheet from your teacher and draw a random sample of 1,000 customers using the following method.
 - a. Add a column to the database. Insert the RAND function in the first customer's cell. This will generate a random number between 0 and 1. Copy this cell to every customer in the database.
 - b. Copy the column to itself by using the Copy command followed by the "Paste Special" command and select the "Values" button. This is done to so that the random numbers do not change every time a key is pressed.
 - c. Sort the entire database including the new column by the random number column.
 - d. The first 1000 records is your random sample.
3. Record the random number assigned to the last customer in your sample. Is this number close to .02? Explain.
4. Use the sort features of Excel to separate the random sample into the six market segments used in the test of the promotion.
5. Calculate the proportion of customers in each segment of the random sample. How do the proportions in the sample compare with the corresponding proportions in the entire database?
6. Calculate the profit value per customer for each segment of the sample and compare with the profit value per customer by segment in the entire database.
7. Redo steps 5 and 6 for the next group of 1000 customers.

Extension 2: The Chi-Square Test

Recall that in Part II of the Student Activity, the Chi-square test indicated that there is no statistically significant difference in the response rates to the promotion across the six segments in the sample that did not receive the incentive. Using the same means and standard deviations as in Table 3 of the Student Activity, manipulate the number of purchasers in each of the six segments so as to produce a statistically significant Chi-square test. Keep the total number of purchasers at 109, but move some of them from one segment to another. Challenge: What is the minimum movement that results in a statistically significant difference in the response rates?

Homework

Classic Music, a direct marketer of music cassettes and CDs, has recently purchased the exclusive rights to market a collection of four CDs containing the greatest hits of the “King of Rock and Roll,” Elvis Presley. Classic Music plans to distribute the four CD set by means of a direct mail promotion. They have obtained a list of 100,000 music lovers from a list broker and plan to mail the offer to the entire list.

First, Classic Music decided to conduct a test on a sample of 5000 potential customers drawn at random from the list. Each potential customer in the test was mailed a color brochure with a cover letter and an invitation to order the set of CDs. 240 customers responded positively by purchasing the set.

1. Give a two-sided 99% confidence interval for the response rate.
2. Give a one-sided 90% confidence interval, with a lower bound, for the response rate.
3. The cost of the mailing, which includes the costs of the brochure, the cover letter, acquisition of the list, and postage, is \$2 per customer mailed. The net profit per sale, excluding the cost of the mailing, is \$55. Was the test mailing profitable?
4. Classic Music has decided to go ahead with the promotion and mail the offer to the remaining 95,000 potential customers in the list. The Director of Marketing is very excited about the promotion and claims there is a probability of 90% that the company will earn in excess of \$40,000 from it. Can you support this claim?
5. The Director of Marketing wants to test whether the type of residential area has any effect on the response rate. Analyzing the responses to the test mailing reveals the following distribution:

	Large Metro Areas	Small to Mid-size Metro Areas	Rural Areas	Totals
Responders	135	55	50	240
Non-responders	2265	1545	950	4760
Total Mailed	2400	1600	1000	5000

Test the hypothesis that the response rate and the type of residential area are independent.


6. Using the data from question 5, test the hypothesis that the following pairs of response rates are similar:
 - a.) Large Metropolitan Areas and Small to Mid-sized Metropolitan Areas.
 - b.) Large Metropolitan Areas and Rural Areas.
 - c.) Small to Mid-sized Metropolitan Areas and Rural Areas.

Solutions

Gamz, Inc. Targets the Market

Part I: Analyzing Information in a Marketing Database

GAMZ Inc. markets computer games online. It has a database of fifty thousand customers, aged 16 to 80. The database contains the recent purchase history of each customer, as well as three basic demographic data points: age, gender, and community type they live in – urban, suburban or rural. These customers have all purchased at least one game from GAMZ within the last 12 months, and no customer has purchased more than three games in the last 12 months.



GAMZ is planning on developing a high quality catalogue that describes a wide range of products targeted at different types of customers, called market segments. The fixed cost of developing the marketing campaign material is estimated to be \$25,000. The variable cost of printing and mailing each catalogue is \$2.50 per potential customer who receives a mailing. The computer games are priced at \$90 each and the gross profit is \$30 per game. The other \$60 covers production costs and royalty payments to the developers but not mailing and catalogue costs.

1. If GAMZ carried out a mass mailing to the entire list, how many games would they have to sell to break even; that is, to recover the costs of developing and mailing the catalogue? *5000 games*


One statistic direct marketers use to decide which segments of a database to target is the response rate: the proportion of customers who bought the product.

2. What would the response rate have to be to break even? *0.1 or 10%*
3. What assumption did you make to determine this number? *That each customer who makes a purchase purchases exactly one game.*
4. Is this assumption realistic? *Probably not.*

Upper management has said they are not really interested in any mass marketing effort unless it will generate at least \$10,000 in profit.

5. How many games would they have to sell to make \$10,000 profit? *5334*

GAMZ also uses a statistic called “profit value per customer” to account for the fact that some customers in the database have purchased more than one game in the past 12 months. Suppose the profit value per customer is \$33.10 for the entire database.



6. On average, how many games did each customer in the database purchase in the past year? *1.1*
7. How does this information change the break even response rate? *It lowers it to 0.091 or 9.1%.*
8. Using the purchase value information, calculate the average response rate necessary to generate \$10,000 in profit. *0.097 or 9.7%.*

In order to maximize the response to a marketing promotion, marketers sometimes separate a customer database into various segments. A segment is a partition of the database that uses one or more factors such as gender, age, or location to characterize it. When a database has been partitioned into segments, every record in the database will belong to exactly one segment. GAMZ has decided to use gender and age to partition its customer database. GAMZ uses three categories of age: under 20, 20-29, and 30 and over. Thus, one of the segments it uses is “females under 20,” and there are six segments altogether. Table 1 contains data on the sizes, response rates, and profit for the last 12 months for the various segments of the GAMZ database.



Segment	Size	Number of Buyers who Bought "N" games			Profit
		N = 1	N = 2	N = 3	
Females under 20	9,029	8391	539	99	\$292,980
Males under 20	9,407	8100	847	460	\$335,220
Females, 20-29	8,202	7740	420	42	\$261,180
Males, 20-29	9,221	8018	802	401	\$324,750
Females 30 & over	7,789	7441	251	97	\$247,020
Males 30 & over	6,352	6256	75	21	\$194,070

Table 1: Response data for GAMZ customers

9. For “Males under 20,” how was the profit of \$335,220 derived? That is, write an equation for that profit that uses the numbers in that row of the table. $\$335220 = \$30(1(8100) + 2(847) + 3(460))$

Optional: Enter the data from Table 1 into a spreadsheet, and translate your equation into a spreadsheet equation to calculate the profit for each segment. *If \$335220 occurs in cell F4, then $F4 = 30(C4 + 2(D4) + 3(E4))$*

10. Looking at the data in Table 1, the 9,029 females under 20 generated \$292,980 in profit. The profit-value per customer of that segment of the market is therefore $\$292,980 / 9,029 = \32.45 . The standard deviations for the profit value per customer of each segment are provided. Compute the profit value per customer for each of the remaining market segments and enter those values in Table 2 below:

Segment	Size	Profit	Profit-Value	Standard Deviation
Females under 20	9,029	\$292,980	\$32.45	\$9.34
Males under 20	9,407	\$335,220	\$35.64	\$15.01
Females, 20-29	8,202	\$261,180	\$31.84	\$7.82
Males, 20-29	9,221	\$324,750	\$35.22	\$14.41
Females 30 & over	7,789	\$247,020	\$31.71	\$8.42
Males 30 & over	6,352	\$194,070	\$30.55	\$4.71

Table 2: Profit-Value per Customer by Segment

11. Do any of the segments of the GAMZ database appear to be more or less promising as target segments for their marketing campaign? *Answers vary, but students may notice the sizeable difference between the largest and smallest values.* In your opinion, how great a difference in the profit value per customer would be of *practical significance*? *Answers vary; at this point, we just want students to begin to think about the possible difference between statistical significance and practical significance.*

Using a one-sided, two-sample Z-test, is the difference between the *largest* and *smallest* profit values per customer from Table 2 statistically significant? *Yes*. What is the value of p , the calculated level of statistical significance? *Very close to zero*. See TI-83 screen captures, which follow.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

12. Using a one-sided, two-sample Z-test, is the difference between the *largest* and *next largest* profit values per customer from Table 2 statistically significant? *Yes*. What is the value of p ? *.026*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

13. If statistically significant, is either of these differences practically significant? *The difference between the largest and smallest is probably of practical significance. The difference between the largest and next largest is almost certainly not practically significant.*

Repeat the two tests of significance from numbers 11 and 12 using the same means and standard deviations, but assume a sample size of 400 for each segment.

14. Is either difference statistically significant assuming the new sample sizes? Explain. *The difference between the largest and smallest remains statistically significant ($z \approx 6.47$, $p \approx 4.89 \times 10^{-11}$), but the difference between the largest and next largest is no longer statistically significant ($z \approx .404$, $p \approx .343$). Hopefully, students will attribute the different results in this case to the dramatically smaller sample sizes.*

Part II: Using Statistical Analysis to Make a Marketing Decision



In order to increase the response to their promotion, the Marketing Department at GAMZ, Inc. is considering using a \$10-off incentive on the initial offer. They decide to test the response to the promotion with or without the incentive. To do so, they draw two random samples of 1,000 customers each from their 50,000-customer database. One of these test samples receives the promotion catalogue with the \$10-off incentive. The other test sample receives the promotion catalogue only, with no additional incentive. Recall from part I that the profit per game is \$30, and the cost of printing and mailing is \$2.50 per potential customer.

1. If a \$10 incentive is used, what is the net profit per game sold? $\$30 - \$10 = \$20$
2. Using the net profit per game sold for the test sample that receives the incentive, calculate the response rates needed to break even and to return at least a \$10,000 profit, as you did in Part I, questions 7 and 8.
 $\$150,000/\$20 = 7500 \text{ games}; 7500 \text{ games}/1.1 = 6819 \text{ customers}; 6819/50,000 = 0.136$
 $\$160,000/\$20 = 8000 \text{ games}; 8000 \text{ games}/1.1 = 7273 \text{ customers}; 7273/50,000 = 0.146$

Table 3 contains a summary of the test data for the two samples.

With Incentive:

Segment	Total sample size	Number of responses	Number of responders who bought:			Total number of games purchased	Response rate	Net Profit per Segment
			1	2	3			
Females under 20	181	27	23	3	1	32	0.149	$32 * (30-10) - 181 * 2.5 = 187.50$
Males under 20	191	32	25	5	2	41	0.168	342.50
Females 20-29	158	16	14	2	0	18	0.101	-35.00
Males 20-29	193	25	18	4	3	35	0.130	217.50
Females 30 & over	132	15	14	0	1	17	0.114	10.00
Males 30 & over	145	18	17	1	0	19	0.124	17.50
TOTAL	1000	133	111	15	7	162	0.133	\$740.00

Without Incentive:

Segment	Total sample size	Number of responses	Number of responders who bought:			Total number of games purchased	Response rate	Net Profit per Segment
			1	2	3			
Females under 20	175	16	13	2	1	20	0.091	$20 * (30) - 175 * 2.5 = 162.50$
Males under 20	184	19	14	2	3	27	0.103	350.00
Females 20-29	165	23	21	2	0	25	0.139	337.50
Males 20-29	180	24	19	4	1	30	0.133	450.00
Females 30 & over	150	16	14	2	0	18	0.107	165.00
Males 30 & over	146	11	11	0	0	11	0.075	-35.00
TOTAL	1000	109	92	12	5	131	0.109	\$1430.00

Table 3: Sample Test Data With and Without an Incentive

- Calculate the net profit per segment for both test samples. Do any of the segments of the GAMZ database appear to be more or less promising as target segments for their marketing campaign? *See answers in Tables above. Males 30 over and Females 30 & over seem to be weak segments in both tests. Some segments seem unprofitable with incentives because of the added costs and random differences in demand.*
- Construct a 90% confidence interval for the difference between the response rates for the two test groups. (.000026, .048) *The interval does not include 0. This demonstrates with 90% confidence that the incentive will produce a statistically significant different response rate.*

5. Use the data in Table 3 to state null and alternative hypotheses regarding the difference between the response rate of the test sample that received the incentive and the response rate of the test sample that did not receive the incentive.

H_0 : There is no difference in the response rates of the two test samples.

H_A : The response rate of the test sample that received the incentive is greater than the response rate of the test sample that did not receive the incentive.

Perform a one-sided two-proportion Z-test of your hypothesis at the 95% confidence level, and state the results of the test. *With 95% confidence, the null hypothesis that the response rates are the same may be rejected in favor of the alternative hypothesis that the response rate with the incentive exceeds the response rate without the incentive.*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

6. Ignoring the cost of producing and mailing the catalogue, compute the total net profit for each test sample.
With the incentive: $161(\$30 - \$10) = \$3220$.
Without the incentive: $131(\$30) = \3930 .
7. Looking back at the response rates that you calculated in questions 7 and 8 in Part I, perform a one-sided one-proportion Z-test at the 90% level of confidence of the following hypotheses if the incentive is *not* used:

- a. H_0 : The response rate will not be sufficient to break even.
 H_a : The response rate will be sufficient to break even.

If the incentive is *not* used, is the promotion likely to break even? *Yes*. What level of confidence do you have in your response? *97.6%*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

- b. H_0 : The response rate will not be sufficient to earn at least \$10,000 in profit.
 H_a : The response rate will be sufficient to earn at least \$10,000 in profit.

If the incentive is *not* used, is the promotion likely to generate \$10,000 in profit? *Yes*. What level of confidence do you have in your response? *90%*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

8. Looking back at the response rates that you calculated in question 2 in Part II, perform a one-sided one-proportion Z-test at the 90% level of confidence of the following hypotheses if the incentive *is* used:
- H_0 : The response rate will not be sufficient to break even.
 H_a : The response rate will be sufficient to break even.

If the incentive *is* used, is the promotion likely to break even? *No. The breakeven response rate is 0.136 and the observed value was only 0.133. What level of confidence do you have in your response? If the null hypothesis is true, there is a 60.9% probability of obtaining a response rate as low or lower than was observed. Thus, there is no evidence for rejecting the null hypothesis, and it must be maintained.*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

- H_0 : The response rate will not be sufficient to earn at least \$10,000 in profit.
 H_a : The response rate will be sufficient to earn at least \$10,000 in profit.

If the incentive *is* used, is the promotion likely to generate \$10,000 in profit? *No. The breakeven response rate is 0.146 and the observed value was only 0.133. What level of confidence do you have in your response? If the null hypothesis is true, there is an 87.8% probability of obtaining a response rate as low or lower than was observed. Thus, there is no evidence for rejecting the null hypothesis, and it must be maintained.*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

The Marketing Department at GAMZ, Inc. has decided to go forward with the promotion, but not to offer the \$10-off incentive.

- Explain why this decision does or does not make sense, based on your tests of the available data. *Despite the fact that we can be 95% confident that the incentive results in a greater response rate, the higher response rate did not offset the cost of the incentive in the test sample. Moreover, even without the incentive, there is better than a 95% chance that the promotion will break even, and a 90% chance that it will generate at least \$10,000 in profit.*

Next, the Marketing Department decides to investigate the test sample that did not receive the incentive to learn whether the response rates for any segments of the database were higher than would be expected. If so, it would make sense to target such segments in the promotion. Table 4 contains the observed numbers of responders and non-responders for each segment of the test sample that did not receive the incentive.

Observed	Females under 20	Males under 20	Females 20-29	Males 20-29	Females 30 & over	Males 30 & over	Totals
Responders	16	19	23	24	16	11	109
Non-Responders	159	165	142	156	134	135	891
Totals	175	184	165	180	150	146	1000

Table 4: Observed Numbers of Responders and Non-Responders by Segment

Notice that of the 1000 customers in the sample who did not receive the incentive, there were 175 females under 20. This is 0.175 of the sample. If there is no segment of the database that is more likely to respond than any other, the expected value of the number of responders for females under 20 is 0.175 of the 109 total responses; that is, $(0.175)(109)=19.075$. Similarly, the expected value of the number of non-responders is $175-19.075=155.925$, as shown in Table 5.

10. Compute the expected value of the number of responders and non-responders from each of the other five segments of the database and enter those values in Table 5.



Expected	Females under 20	Males under 20	Females 20-29	Males 20-29	Females 30 & over	Males 30 & over	Totals
Responders	19.075	20.056	17.985	19.62	16.35	15.914	109
Non-Responders	155.925	163.944	147.015	160.38	133.65	130.086	891
Totals	175	184	165	180	150	146	1000

Table 5: Expected Numbers of Responders and Non-Responders by Segment

Now, a Chi-square test of association may be performed to test the null hypothesis that the observed response rates across all of the segments are the same against the alternate hypothesis that these response rates differ significantly. To calculate the value of Chi-square, determine for each segment the difference between the observed number of responders in Table 4 and the expected value of that number in Table 5. Similarly, the difference between the observed number of non-responders and its expected value must be computed. Those differences are then squared, and the squared differences are divided by the expected values. Finally, these twelve terms are summed to get the value of Chi-square.

The mathematical formula is: $\chi^2 = \sum \frac{(o-e)^2}{e}$, where

o and e represent the observed numbers and their expected values, respectively. Obviously, computing the value of Chi-square in this case involves a good deal of tedious computation.



Fortunately, statistics software on computers or calculators can be utilized to accomplish this task. The set-up for computing Chi-square using a TI-83 Plus calculator is described below. To use a different calculator or computer software, refer to the accompanying documentation.

To begin the set-up, the information in Table 4 above must be stored in a 2x6 matrix, A , on the calculator. To do so, press “2nd” and “MATRIX” to obtain the menu for working with matrices, highlight “1:” and “EDIT”, and press “ENTER”. You may now edit “MATRIX [A]”. Notice that the default dimensions are 1x1, so change that to 2x6 and begin entering the values from Table 4. Matrix A is too wide to be viewed on one screen, so the last two screens in the set-up show what the completed matrix A should look like.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

Figure 1: TI-83 Matrix Set-Up

Now, to calculate the value of Chi-square, press the “STAT” key to access the statistics menu, highlight “TESTS” and “C:” (“ χ^2 -TEST ...”), and press “ENTER”. Then highlight “CALCULATE” and press “ENTER” again.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

Figure 2: Calculating the Value of Chi-square Using a TI-83

11. Compute the value of χ^2 .

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

12. For these data, what is the probability of obtaining a value of Chi-square at least as large as the computed value? *0.416* What does this probability tell you about the null hypothesis? *The null hypothesis that the response rates across all six segments are the same cannot be rejected, except at a very low level of confidence.*
13. From the output of the computation of the value of Chi-square, how many degrees of freedom did this computation have? *5*

One way to interpret the number of degrees of freedom in a data set is as the number of independent pieces of information contained in the data set. If any piece of information in the table may be computed from other pieces of information in the table, then it is not an independent piece of information.

13. In light of the above interpretation of degrees of freedom, explain why your answer to question 9 does or does not make sense. *It makes sense, because the row and column totals in Table 4 are known. Therefore, the minimum number of entries in the table that must be known in order to calculate all of the other entries is five. Thus, there are only five independent pieces of information in the table.*
14. Would you recommend to the Marketing Department at GAMZ, Inc. that they target specific segments of their database in the promotion? *No* Why or why not? *The Chi-square test indicates that there is no statistically significant difference in the response rates across segments.*


Finally, the Marketing Department at GAMZ, Inc. must forecast the total expected net profits from the promotion in order to estimate whether the promotion is viable. Recall that management requires a net profit of at least \$10,000 to go forward with any promotion. If the promotion is mailed to the remaining 48,000 customers in the database, use the data from the test sample that did not receive the incentive to answer the following questions.

15. What is the expected value of the number of purchasers? *.109(48,000) = 5232*

16. On average, what is the forecasted total number of games that will be purchased? $(131/109)(5232) = 6288$.
Alternatively, 131 games/1000 customers were purchased in the test sample, so $.131(48,000) = 6288$.
17. Compute the expected value of the gross profit on the sale of the games. $6288(\$30) = \$188,640$

Recall from Part I that there is a fixed cost of \$25,000 associated with the promotion, as well as a cost of \$2.50 for each customer to whom the promotion is mailed.

18. Compute the expected value of the total net profit on the entire promotion, excluding the two test samples.
 $\$188,640 - \$25,000 - (\$2.50)48,000 = \$43,640$
Does it appear that the promotion will meet management's requirement of net profits of at least \$10,000?
Yes, based on the result of the forecast. More conservatively, even if the lower tail estimate of a 0.091 response rate had been used, the result would still be almost \$16,000 in net profit.



Extension 1: Results vary.

Extension 2: One way to do so follows. Cells where changes were made in the observed responses are bolded.

Observed	Females under 20	Males under 20	Females 20-29	Males 20-29	Females 30 & over	Males 30 & over	Totals
Responders	16	19	26	24	17	7	109
Non-Responders	159	165	139	156	133	139	891
Totals	175	184	165	180	150	146	1000

QuickTime™ and a TIFF (Uncompressed) decompresso are needed to see this picture.

Homework:

1. (.040, .055)

QuickTime™ and a TIFF (Uncompressed) decompresso are needed to see this picture. QuickTime™ and a TIFF (Uncompressed) decompresso are needed to see this picture.

2. (.044, .048)

QuickTime™ and a TIFF (Uncompressed) decompresso are needed to see this picture. QuickTime™ and a TIFF (Uncompressed) decompresso are needed to see this picture.

3. Yes; $(\$55/\text{order})(240 \text{ orders}) - (\$2/\text{mailing})(5,000 \text{ mailings}) = \3200 .
4. Choosing the lower bound of .044 from the one-sided 90% confidence interval in question 2, the forecasted number of responses is 4180. Then the profit can be forecast as computed in question 3: $(\$55)(4180) - (\$2)(95,000) = \$39,900$. Thus, the forecast by the Director of Marketing was only slightly optimistic.
- 5.

	Large Metro Areas	Small to Mid-size Metro Areas	Rural Areas	Totals
Responders	135	55	50	240
Non-responders	2265	1545	950	4760
Total Mailed	2400	1600	1000	5000

QuickTime™ and a TIFF (Uncompressed) decompresso are needed to see this picture. QuickTime™ and a TIFF (Uncompressed) decompresso are needed to see this picture.

Since the probability that the computed value of χ^2 is 10.16 or larger purely by chance is only 0.006, we reject the hypothesis that there is no difference in the response rates across type of residential area.

6. a.) *There is a statistically significant difference in the response rates of large metropolitan areas and small to mid-sized metropolitan areas in favor of large metropolitan areas ($z=3.19$, $p=.00072$).*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

Alternatively, the 90% confidence interval on the difference of the two proportions does not include 0, so the null hypothesis that there is no difference in the proportions can be rejected in favor of the alternate hypothesis that there is a difference.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

- b.) *There is not a statistically significant difference in the response rates of large metropolitan areas and rural areas ($z=0.73$, $p=.23$).*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

- c.) *There is a statistically significant difference in the response rates of small to mid-sized metropolitan and rural areas in favor of rural areas ($z=-1.97$, $p=.024$).*

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompresso
are needed to see this picture.