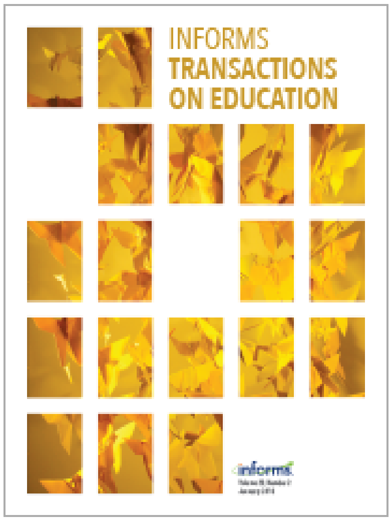

INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Case—Managing Staffing Inefficiencies Using Analytics (B): Business Value in Predictive and Prescriptive Analytics Models

David Kopcso, Dessislava Pachamanova

informs
http://pubsonline.informs.org/journal/ited/

INFORMS TRANSACTIONS ON EDUCATION
Vol. 19, No. 1, September 2018, pp. 43–47
ISSN 1532-0545 (online)

Case

Managing Staffing Inefficiencies Using Analytics (B): Business Value in Predictive and Prescriptive Analytics Models

David Kopcso,[a] Dessislava Pachamanova[a]

[a] Division of Mathematics and Science, Babson College, Wellesley, Massachusetts 02457

Contact: kopcso@babson.edu (DK); dpachamanova@babson.edu (DP)

Received: August 2, 2016
Revised: February 25, 2017; July 20, 2017; July 25, 2017; July 26, 2017
Accepted: July 26, 2017
Published Online in Articles in Advance: October 27, 2017

https://doi.org/10.1287/ited.2017.0186cs

Keywords: business value • models • predictive analytics • prescriptive analytics • classification • ranking • accuracy • confusion matrix • lift chart • ROC curve • profit curve • yield management • simulation

AdviseInvest's Business Problem

AdviseInvest is a venture-backed startup whose mission is to provide ordinary people with access to affordable financial advice. Potential AdviseInvest customers arrive at the AdviseInvest web site and fill out a customer profile with information about their goals and background. They then schedule a call with a sales representative for a particular one-hour window. AdviseInvest has four sales reps on staff during each particular one-hour window. Each sales rep is assigned two customers. The problem is that about 50% of the customers who schedule a call do not answer. Not only does this leave sales reps underutilized, it does so unevenly. For example, it is possible that one sales rep will not get to do any sales calls during that hour, whereas another will be busy the entire hour with two calls. The Director of Sales has approached a small team of analysts asking for a solution to the problem.

One approach to solving the problem posed by the Director of Sales is to use predictive analytics models to identify customers who would be more likely to answer the call scheduled with the sales reps and prescriptive analytics models to decide how to allocate these customers to sales representatives. (See the Appendices for a review of analytics models and applications.) You are asked to analyze the performance and business value of these models, and to recommend a course of action to the Director of Sales.

Appendix A. Business Problems and Analytic Models

While there is an infinite number of business problems, many of them can be framed within a small number of categories of analytical techniques. The process of framing itself is not necessarily straightforward. For example, typical questions such as

1. "How can I offer my customers recommendations for other products/services that they are likely to buy?";
2. "How can I get an idea of my customers' opinions about my company";
3. "Is there a way to use historical inventory data to stock up better in the future?";
4. "Is my new website design better than the old website design?"

should first be translated into questions that specify metrics or point to data and analytical tools:

1. Are two items (customers, products) "similar"? Is a particular customer likely to respond to a particular offer?
2. Can I mine comments customers have made for feedback?
3. Are there patterns or trends over time that I can observe? And if there are, how can I manage my inventory optimally so as to minimize costs?
4. Does my new website design generate significantly more purchases?

Table A.1 summarizes some common questions and analytical techniques that could be used to address them. (See also EMC Education Services 2015 as well as Kokina et al. 2017 for applications to managerial accounting and strategy maps.)

Analytical models are employed to various degrees in organizations, and the degree to which they are utilized in an organization is associated with the organization's analytics maturity. Figure A.1 summarizes the stages of analytics maturity from analytics maturity models suggested in the literature. The first two stages involve data collection as well as using data for describing and observing events. The third stage uses data and basic models for setting alerts

Table A.1. Correspondence Between Common Business Questions and Analytics Techniques

Business problem	Analytical technique
Observe trends and relationships; formulate hypotheses	Data visualization
Decide if a new situation is "better" than the current situation	Hypothesis testing
Determine the impact of various factors on an output variable of interest	Linear regression, logistic regression, classification and regression trees
Identify groups of observations that are similar in terms of specified characteristics	Cluster analysis
Identify transactions that tend to occur together	Association rules (Market basket analysis)
Assign a score (likelihood, ranking) to an observation (customer, loan, transaction, etc.)	Logistic regression, classification and regression trees
Classify an observation (customer, loan, transaction, etc.) into a category	Logistic regression, classification and regression trees, Naïve Bayes, *k* Nearest Neighbors, Support Vector Machines (SVM), Neural networks
Analyze data over time; forecast	Time series analysis
Analyze text data (customer call transcripts, social media posts, etc.); perform sentiment analysis	Text analytics, Natural Language Processing (NLP)
Find an optimal mix of products given limited resources; find an optimal route/schedule/price	Optimization

and making observations about factors driving events. The fourth stage is predictive—predictive models are employed to forecast and identify opportunities. The fifth stage is prescriptive—models of increasing sophistication are used to determine the best course of action for a particular situation.

A sign of analytics maturity for an organization is the ability to use the output of predictive models as input to prescriptive analytics models. To realize the business value from analytics models, analytics professionals at the company also need to estimate the value of using these models for various stakeholders, such as the company's customers and employees, and to adjust those models to achieve the desired value. There are many challenges with integrating predictive and prescriptive analytics effectively.

Appendix B. Evaluating Model Performance

Because of the binary nature of the variable to be predicted in the AdviseInvest problem (Answer (1)/Not Answer (0)), classification and regression trees or logistic regression are possible choices for predictive modeling techniques (see Table A.1 in Appendix A). A decision tree is created with the `rplot` package in *R* for the customer data. The tree is built using 90% of the data as a training set, and 10% of all observations are set aside as a test data set. The output from classification models such as decision trees is a score, interpreted as a probability, which is then used to classify the observations based on a cutoff (threshold) value. AdviseInvest customers in the test data set are scored with a score between 0 and 1 based on the decision tree output. A base case cutoff (threshold) level of 0.50 is used for deciding whether to assign an observation to the "Answer" (if the score is ≥ 0.50) or "Not Answer" (if the score is < 0.50) category.

B.1. Classification Accuracy and Value

Predictive model classification accuracy can be evaluated through a variety of metrics and charts. These include total accuracy, percentage of false positives and false negatives, and Receiver Operating Characteristics (ROC) curves. Please refer to Evgeniou (2016), Shmueli et al. (2016), or Provost and Fawcett (2013) for more detail.

Figure A.1. Stages of Analytics Maturity and Relationship of Analytics Maturity and Business Value

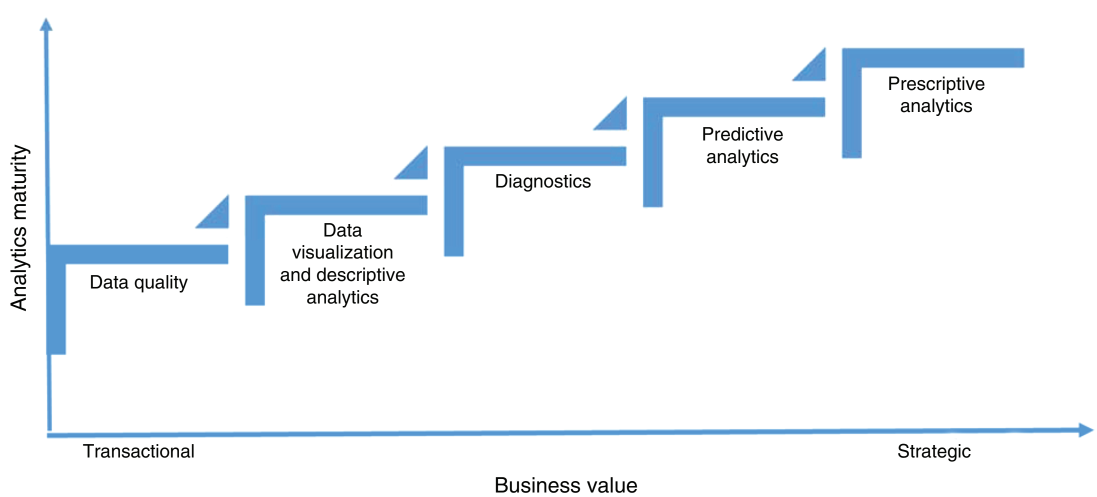

Note. Compilation based on multiple sources: INFORMS, Deloitte Consulting, SAS.

B.1.1. Total Accuracy. The total accuracy (hit ratio) is the percentage of all observations that are correctly classified by the algorithm. Suppose that in this example it is 75%. This means that 75% of all observations in the test data set were either customers who did not answer and were correctly classified as "Not Answer" or customers who answered and were correctly classified as "Answer."

The simplest model that can be used as a benchmark is the so-called Maximum Chance Criterion or Naïve Rule. This is the accuracy that would be obtained if, instead of building a predictive model, we follow the rule of the dominant class. Suppose that 47% of all customers in the test data set did not answer and 53% answered. If we had simply assumed that everybody will answer, we would have achieved a prediction accuracy of 53%.[1] The decision tree model would have therefore improved prediction accuracy overall.

B.1.2. Confusion Matrix. The problem with using total accuracy as a measure of algorithm performance is that it weighs all misclassifications equally. In other words, observations that are 1 (Answer) and are classified as 0 (Not Answer) are counted in the same way as observations that are 0 (Not Answer) and are classified as 1 (Answer). The confusion matrix (Table B.1) helps separate the cases and calculate not only the overall model accuracy but also the model accuracy for each particular class.

For example, 67.92% (= 0.36/0.53) of all customers who answered were classified correctly, and 82.98% (= 0.39/0.47) of all customers who did not answer were classified correctly. The former is referred to as true positive rate or sensitivity, whereas the latter is referred to as true negative rate or specificity.

Observations that were 1 but were classified as 0 are referred to as false negatives. Observations that were 0 but were classified as 1 are referred to as false positives. Correctly classified 1s are referred to as true positives, whereas correctly classified 0s are referred to as true negatives.

B.1.3. ROC Curve. The default cutoff value for classification is 0.5—if the predictive model assigns an observation a score of 0.5 or more, the observation would be classified as a 1 (Answer); if the score is less than 0.5, the observation would be classified as a 0 (Not Answer). Varying the cutoff value changes the classification and results in different confusion matrices and different values for the errors.

Changing the cutoff also changes the percentage of observations that are correctly identified (true positives and true negatives). A receiver operating characteristics (ROC) curve represents this tradeoff by plotting the true positive rate versus the false positive rate (which is 1 minus the true negative rate) for various levels of the cutoff. The false positive rate is used to calculate how bad the algorithm is at identifying false positives. The ROC curve for the AdviseInvest problem is shown in Figure B.1.

Table B.1. Confusion Matrix for the AdviseInvest Problem Using a Cutoff Value of 0.5 for Classification

	Predicted		
Actual	Not answer (0)	Answer (1)	Total
Not answer (0)	0.39	0.08	**0.47**
Answer (1)	0.17	0.36	**0.53**
Total	**0.56**	**0.44**	**1.00**

Figure B.1. ROC Curve

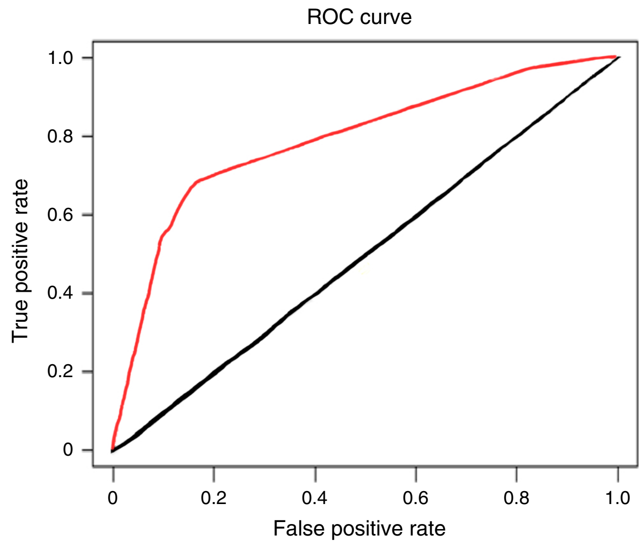

The higher and the farther away the ROC curve is from the 45-degree diagonal line (i.e., the larger the area under the curve, which in itself is a metric of performance, usually denoted AUC), the better the model is considered in terms of effectiveness.

B.1.4. Incorporating Business Value. Suppose a sales rep is paid $20 per hour and can be hired for flexible time periods. Given the current scheduling method of two customers per sales rep per hour, one can approximate the cost of a sales rep sitting idle when a customer does not answer the phone at $10. Suppose the expected revenue per customer who answers the phone is $110. Hence, the expected profit from a customer who answers the phone is $110 − $10 = $100. Customers who are predicted to not answer the call are not assigned a sales rep. The payoff structure is represented in Table B.2.

Using the confusion matrix in Table B.1 and the payoff structure in Table B.2, one can calculate useful information. For example, one can calculate the profit that would be realized if the predictive model with the current cutoff value is employed to decide which customers to schedule.

B.2. Ranking and Value

In practice, one rarely has sufficient information to predict scores accurately as a proxy for probabilities. Scores from

Table B.2. Payoff Matrix

	Predicted ($)	
Actual	Not answer (0)	Answer (1)
Not answer (0)	0	−10
Answer (1)	0	100

Figure B.2. Decile-Wise Lift Chart: Proportion of Answers in Each Decile Relative to Proportion of Answers in the Entire Data Set

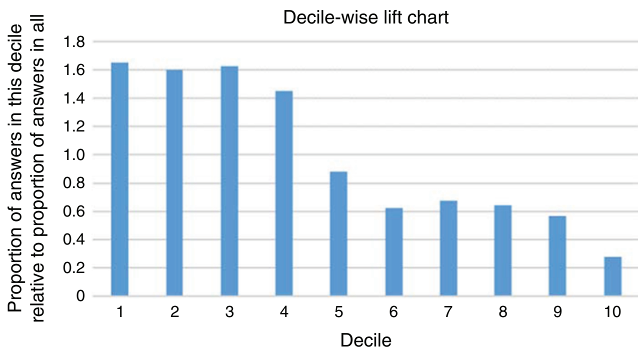

predictive models, however, are often useful for ranking all observations on a relative basis. This is consistent with many applications of such predictive models in practice, where one usually is assigned a budget or has a constraint on the number of potential customers to reach. So, ranking all customers and reaching out to a subset of them who are the highest ranked within an allowable budget (rather than reaching out only to those with scores higher than a threshold) makes sense in a business context.

A score-ranking-based graph that is often used in practice is a decile-wise lift chart. Its idea is to show how much more likely it is to identify positive observations (in this case, customers who would answer the call) in the top deciles by score than it would be by random chance. The decile-wise lift chart is determined by first calculating the percentages of "Answers" that are in each decile and then dividing those percentages by the percentage of "Answers" in the overall data. The resulting chart is given in Figure B.2.

Another useful graph is a cumulative gains chart. A cumulative gains chart plots the percentage of true positives (i.e., positives correctly classified as such) in a proportion of the (score-ranked) test set, ranging from 0% to 100% of all observations in the data set. Figure B.3 shows the cumulative gains chart for the AdviseInvest example. The "higher" the cumulative response curve is over the 45-degree line, the better the model does at identifying true positives. (The 45-degree line shows the number of true positives that would be identified if observations are selected at random. For example, if we pick 20% of the test observations at random, we would expect to find 20% of the true positives in that subset.)

Figure B.3. Cumulative Gains Chart for AdviseInvest Example

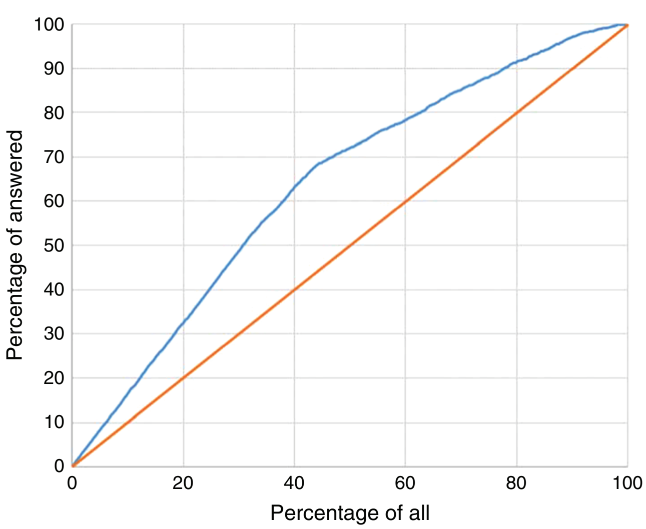

Appendix C. Beyond Predictive Models: Prescriptive Analytics and Value

The predictive model output from Appendix B provides information about the customers scheduled for a phone call during each one-hour interval. There is now information about their score and their relative ranking, which helps when thinking about a scheduling method that is better than the status quo. As one prescribes a "better" system, however, one needs to clarify what "better" means. Three common management objectives are (see, for example, Weatherford and Bodily 1992):

1. Maximize profit/contribution. Profit is revenues minus variable costs minus fixed costs minus taxes. Contribution towards fixed costs is revenue minus variable costs. Profit and contribution are often used synonymously, and typically if one is maximized, so is the other. We use profit in the sense of contribution here.

2. Maximize capacity utilization. In the context of the AdviseInvest problem, this would mean making sure that the eight customer slots for each hour are always full.

3. Maximize average revenue per customer. This goal will explicitly reward scheduling higher-value customers. Although important, this goal can sometimes lead to results opposite to goal 1. In the most extreme case, it might lead to a situation in which only a single customer with the highest expected value is scheduled for each time slot. So, it should be used with caution.

Determining one or more correct objectives should happen after consultation with the stakeholders in the project; although more often than not the most useful information one can extract from such consultations is a prioritization of the goals. It would be a good idea to keep track of how new solutions do according to multiple criteria of interest to the business sponsor of the analytics project.

The current system is to allocate customers to sales reps in the order in which they have requested the time slot. Table C.1 shows how the system would work on an example of eight customers extracted from the test set. As customers are signing up, they get distributed to sales reps. Customers C1 and C5 are assigned to sales rep S1, customers C2 and C6 to sales rep S2, and so on. Customers C2, C3, C5 and C6 do not answer the call (Answer = 0 for those rows). Sales rep S4 services two customers but sales reps S1 and S3 service one, and sales rep S2 does not service any. The average utilization of the sales reps is $(50 + 0 + 50 + 100)/4 = 50\%$, and it is unevenly distributed among the sales reps.

Before the outcome of the call is known, each customer who answers is expected to realize $110 in average (expected) revenue (see column "Expected Revenue" in Table C.1).

Table C.1. System A (Status Quo)

Sales rep	Customer	Answer	Score	Expected revenue ($)
S1	C1	1	0.36	110.00
S2	C2	0	0.09	0.00
S3	C3	0	0.36	0.00
S4	C4	1	0.87	110.00
S1	C5	0	0.82	0.00
S2	C6	0	0.09	0.00
S3	C7	1	0.88	110.00
S4	C8	1	0.36	110.00

Endnotes

[1] Note that this number can be exaggerated in situations in which one class is much less likely to occur than the other. Customer response to a direct order marketing campaign, for example, is in the range 1.5%–2%. If, instead of using a predictive model to score customers on their likelihood of response to the campaign, a Naïve Rule is used in which no customers are contacted, the Naïve Rule would achieve an accuracy of 98%–99.5%.

References

EMC Education Services (2015) *Data Science and Big Data Analytics: Discovering, Analyzing, Visualizing and Presenting Data* (John Wiley & Sons, Hoboken, NJ).

Evgeniou T (2016) Classification methods. Retrieved July 14, 2016 from http://inseaddataanalytics.github.io/INSEADAnalytics/Report_s67.html.

Kokina J, Pachamanova D, Corbett A (2017) The role of data visualization and analytics in performance management: Guiding entrepreneurial growth decision. *J. Accounting Ed.* 38:50–62.

Provost F, Fawcett T (2013) *Data Science for Business: What You Need to Know About Data Mining and Data-Analytic Thinking* (O'Reilly Media, Sebastopol, CA).

Shmueli G, Bruce P, Patel N (2016) *Data Mining for Business Analytics: Concepts, Techniques, and Applications with XLMiner*, 3rd ed. (John Wiley & Sons, Hoboken, NJ).

Weatherford L, Bodily S (1992) A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Oper. Res*. 40(5):831–844.