

# Partnerships in Training

**Jason R. W. Merrick, Jill R. Hardin**

Department of Statistical Sciences & Operations Research,  
Virginia Commonwealth University, PO Box 843083, Richmond, Virginia 23284,  
jrmerric@vcu.edu, jrhardin@vcu.edu

**Russell Walker**

Strategic Technology, Capital One Financial Corporation, 4801 Cox Road,  
Glen Allen, Virginia 23060, russell.walker@capitalone.com

Capital One Financial Corporation is a major supplier of credit products in the US market. A key part of their growth and success over the last decade has been an *Information Based Strategy* collecting large quantities of data and performing advanced analysis to leverage the information obtained to make good business decisions. The ability to continue this successful strategy is contingent on training their analysts in key analytical techniques. Through a partnership with the Department of Statistical Sciences and Operations Research at Virginia Commonwealth University, three training courses were developed in forecasting, optimization, and simulation. These courses have not only been well received by the analysts, but have led to considerable early return on investment with one project forecasted to reduce costs at their Richmond (VA) mail center by more than \$2 million annually. We describe this successful partnership as an example of the kind of partnerships from which universities and private corporations can jointly benefit and highlight what we believe to be the key to the training's success: the use of internal and external trainers who work as a team to provide sound, comprehensive training that is relevant to the problems faced by the analyst.

## **1. Introduction**

Credit card companies are in a ideal position to exploit Operations Research and Management Science (OR/MS) methodologies like forecasting, optimization, and simulation (Board et al. 2003). At the core of most credit card companies are their credit risk models that predict customers' credit worthiness based on certain criteria, allowing decisions to be made about awarding credit cards and the credit limit and interest rates that should be offered (Trench et al. 2003). Credit card companies collect large quantities of data on transactions, customers, and their internal processes and operations. They use this data to make forecasts of diverse quantities, such as their customers' non-payment rates, incoming mail and phone call volumes, and response rates to mail solicitations. They use the forecasts in making complex decisions from marketing strategies to delinquency cure to call center staffing, each of which requires sophisticated analysis. Many of their internal systems, such as mail and call centers, are dynamic in nature and must be designed to be robust to the variability experienced in their incoming volumes (Saltzman and Mehrotra 2003).

Capital One® is a leader in the direct marketing of MasterCard® and Visa® credit cards, auto loans, and other consumer financial services. With more than 47 million accounts, it has one of the world's largest financial services franchises. The use and management of data is the key to their success, making research on how to remain the best in their class critical to maintaining continued growth. One of their current research programs endeavors to evaluate new data sources and analysis techniques, with the aim of accelerating and enabling the introduction of new methods into their business decision processes. As a result, they must cultivate their talent pool to improve the

knowledge of key decision techniques, including forecasting, optimization, and simulation, to facilitate the improvement of current and future models used to manage their portfolio of accounts.

Capital One commonly compares its hiring process for analysts to recruiting by the National Basketball Association (NBA). The NBA often drafts high-school athletes, focusing more on talent and athleticism than on making sure their draft choices have experienced four years of college play. For Capital One, this strategy involves hiring MBAs from top business schools who have had one or two quantitative courses and then training them to perform the high-level analysis necessary in their roles in the company. While most of these MBA graduates have had exposure to forecasting, optimization, and simulation (Albritton et al. 2003), they do not consider themselves experts in these methods. They often feel as if they are reinventing the wheel rather than applying well-developed analytical methodologies. Although Capital One does tend to hire core groups of graduates with PhDs and master's degrees in statistics, they hire fewer graduates from OR/MS programs.

Rather than hiring advanced OR/MS graduates to act as internal consultants, Capital One's strategic objective is to train all of their analysts to use applicable analytical methods. However, they recognize that there is a gap in the current level of knowledge in the key analytical methods that they are seeking to leverage. Many training courses are available to close this gap. Most analytical software vendors provide extensive training in the use of their software, and a large proportion of these vendors will even customize the content of the course to the customers' needs. Furthermore, university courses are available, often in the evening, to provide theoretical background.

Still, there are few courses that cover software from multiple vendors while supplying enough background to ensure the correct application of the underlying methods.

Furthermore, companies possess a critical need to make the material relevant to their own business problems (Carraway and Clyman 1997), teaching through their own companies' examples and case studies to give the material context. This latter feature is perhaps the most important so that the analysts know not only how to use the techniques, but also when to use them, creating new opportunities to use the techniques in their own enterprise rather than simply repeating the same type of analysis that others have used.

To fill this gap in training options, Capital One turned to the operations research faculty at Virginia Commonwealth University (VCU). VCU is a public research university located in Richmond, Virginia, where Capital One has established a significant presence. VCU recognizes its role as an urban university, welcoming opportunities to combine academic and real-world education. Thus, we have been well placed to partner with Capital One to provide not only a suitable location and facility for training, but also knowledge of the latest instructional methods for teaching analytical material.

In this paper, we discuss a successful partnership between Capital One and VCU. We use the term *partnership*, as the training team includes both faculty from VCU and analysts from Capital One. We found the combination of internal and external trainers to be critical to the success of the venture. The external trainers brought an up-to-date knowledge of pedagogy to the venture; the internal trainers brought business relevance and access to interesting data and case studies. The close geographical proximity allows for a tighter partnership, with Capital One and VCU personnel working together to

develop and teach the courses and to improve the offerings continually based on official and unofficial feedback.

One project following an early course, which included assistance both from members of the training team and analysts who had taken the courses, has already led to a savings of \$2 to \$5 million annually in staffing costs at one of Capital One's incoming mail centers. Capital One has been ranked as the Best Training Company in America amongst financial companies by Training magazine, citing the return on investment from this training course as an example.

We begin our discussion in Section 2 by outlining our basic training approach. We then describe each of the three training courses in more detail in Section 3. In Section 4 we describe successful outcomes of the course and follow with our own thoughts on why the training has been so successful in Section 5. We do not claim to have created a new pedagogy for training, although we endeavor to apply the latest that others are developing. We do not claim to have created a new type of university-corporate relationship, although the partnership formed has led to significant benefits to both parties. We do claim, however, that the success of the venture makes the story worth telling and the elements of success worth examining.

## ***2. The Pedagogical Approach***

We designed training courses for Capital One analysts to teach them how to solve business problems using three key analytical techniques: forecasting, optimization, and simulation. In fact, as most problems are not solved using just one technique, analysts need to know how to combine the various techniques to solve their business problems. As an example, optimization models often use regression or time series based forecasts as

inputs. Furthermore, the systems they seek to optimize often include stochastic delays that are best modeled using discrete event simulation. Thus, while the analysts are not encouraged to take the courses in a particular order, we make every effort to integrate techniques developed in any course with skills taught in other courses.

Capital One's strategic objectives were a vital element in the development of learning goals for these courses. Capital One's goal was to have its own analysts running successful projects using these key techniques, thus each analyst should:

1. know how to use the techniques correctly
2. know how to *implement* them using appropriate software packages; and
3. know *when* to use them.

Most of the analysts attending the training had been through one or two quantitative analysis courses in business or engineering schools and had some basic knowledge of the techniques and software. It is the last learning goal—knowing when to use each technique— that is hardest to teach.

Our solution to ensuring comprehensive coverage while still targeting Capital One's analytical needs is the use of both internal and external trainers. Business relevance is a key aspect of the training. The importance of using relevant examples and applications is well recognized among educators in business schools (Carraway and Clyman 1997), so it is surprising that corporate training in analytical methodologies is often based on examples from baseball and basketball statistics. While these examples might be appealing to those interested in sports, they are not relevant and do not help an analyst see how to apply the methods in his own work. Therefore all course examples and exercises used in our training are financial or logistical in nature, and nearly all of them

are in the context of Capital One operations, with the applications and data being drawn from various Capital One business areas.

Another valuable feature of the training courses is the vendor-neutral approach taken to software selection. Sanders and Manrodt (2003) recently performed a survey of the type of software used in forecasting in US corporations and the relative success obtained using various types of packages. One of their major findings was that the use of commercially available software led to greater success in the application of the techniques, achieving more accurate forecasts and better overall user satisfaction. It is logical to assume that the same finding might prove true for other techniques, such as optimization and simulation. However, this still leaves the choice of which commercially available package to use.

Software vendors understandably wish to promote the use of their packages, yet as most analysts would confirm, different packages are often suited to different applications. Thus, a comprehensive coverage of suitable tools is useful to the analyst. Furthermore, the suite of packages often needed for a complete solution is usually not available from a single vendor. Attending training for multiple software packages and taking university courses for coverage of the methodology is not feasible for a corporation that wishes to implement analytical methodologies in a timely and cost effective manner.

A focus on problem modeling gives analysts appreciation of how they will use the tools covered in a course (Powell 2001). Instruction in these training courses does not consist of a piece-by-piece development of the methods, nor is it a menu item-by-menu item software demonstration. The idea is to use active learning principles in the training

environment, setting up a business case and developing the techniques to solve the problem (Kolb, 1984). The case method is widely recognized as a breakthrough in teaching quantitative methodologies in business schools (Böcker 1987; Bodily 1996). It provides relevance and gives learners the chance to steer the discussion towards what they want to learn (Corner and Corner 2003). However, this approach requires participants to prepare the case before class, whether by reading it or preparing solutions for it. In the training environment, the instructor often cannot require such preparation. Like so many other OR/MS education researchers (Scott and Buchanan 1992), we do not want to resort to straightforward lecture. We adapt then, by integrating interactive discovery techniques with the discussion and presentation of information, accounting for the fact that participants will spend little time in study outside of the classroom.

Teaching material is developed in PowerPoint™, Capital One's preferred communications medium, and is both projected to the class and printed out for their reference. As the courses are taught in the university's computer classrooms, software demonstrations are also projected and the participants follow along on desktop computers to gain basic experience with the software. Despite the increased interest resulting from the relevance and applicability of the material, the participants are human and our aim is to minimize the presentation segments, keeping them to a maximum of twenty minutes before we return to active learning. We intersperse frequent, simple, reinforcing exercises to help solidify concepts before moving on to new material. In this way, we incorporate independent study and thought in the classroom, rather than through out-of-class assignments. After the primary coverage and reinforcement exercises, the analysts feel confident in their ability to apply the techniques, yet many attendees feel uncertain of

how to apply them in their own work. The examples and exercises are chosen carefully to illustrate key points and to ensure appropriate complexity: they can be solved in a reasonable period of time. Unfortunately, the applicability of the techniques to solve actual problems faced by the analysts is often not obvious. With this in mind, and given that they now have some knowledge of primary analytical techniques, we turn to the Learning Labs, which use a modified case method.

The “cases” used in the Learning Labs are scenarios drawn from Capital One business problems. In fact, they are often drawn from problems faced by analysts who are attending the course. As cooperative learning has been shown to improve the understanding of concepts and interest in the material (Lasdon and Liebman 1998), we encouraged the analysts to work in teams on the cases, drawing assistance from the trainers when necessary and from each other to understand the business context. As teams encounter obstacles, the class is brought back together to discuss the problem.

In many Learning Labs, we develop a group solution in a discussion format for one line of business and then the participants perform their own analyses for other lines, essentially making a real-time case. In this manner, the participants can see the kinds of questions the instructors ask in order to arrive at a better understanding and a reasonable solution approach. If the techniques covered in the course are not suitable, then we can explain why and explore what other techniques might be appropriate. If the techniques are suitable, then participants can perform the relevant analysis, making mistakes and iterating until they achieve a reasonable solution. The Learning Labs, for many, are the most useful and enjoyable parts of these training courses, as they are not finding tidy solutions to pre-prepared teaching examples. Learning Labs not only reinforce the

second learning goal of knowing how to use the techniques, but move participants toward obtaining the third learning goal: knowing *when* to use them.

### **3. *The Course Designs***

To ensure that each of Capital One's learning goals is considered in our courses, we developed the following training objectives, which were the overarching considerations in our course designs:

- To inform participants of the various types of models available to them and the issues surrounding each
- To teach participants to build correct models
- To help participants understand how solutions are obtained and why some models are more difficult to solve than others
- To guide participants in selecting an appropriate model for a given scenario

To date three different training courses have been designed for Capital One.

- Time Series and Forecasting
- Simulation: Modeling and Analysis
- Optimization and Decision Techniques

We now discuss each of these courses in more detail.

#### **3.1 *Time Series and Forecasting***

The Time Series and Forecasting course is a five-day course, including a review of making decisions with statistics as well as detailed coverage of regression-based forecasting and time-series forecasting. An outline of course topics covered can be found in Table I.

**Table I: Course Outline for Time Series and Forecasting**

<b>Day 1</b>	<b>Statistics Review</b> <i>Answering questions with data</i> <i>Introduction of JMP, @RISK</i>
<b>Day 2</b>	<b>Introduction to Regression Analysis</b> <i>Simple Regression</i> <i>Multiple Regression</i> <i>Model fitting criteria (AIC, SBC, Cp)</i> <i>Stepwise regression</i>
<b>Day 3</b>	<b>LEARNING LAB I – Forecasting with Regression</b> <b>Time Series Analysis I</b> <i>Auto-correlation</i> <i>Auto-regressive and moving average models</i> <i>JMP’s time series platform</i> <i>@RISK simulation of time series models</i>
<b>Day 4</b>	<b>Time Series Analysis II</b> <i>Modeling trends</i> <i>ARIMA models</i> <i>Modeling seasonality</i> <i>SARIMA models</i> <i>Using transformations</i> <i>@RISK simulation of time series models</i>
<b>Day 5</b>	<b>Time Series Analysis III</b> <i>Exponential Smoothing Models</i> <i>Combining Time Series and Regression</i> <b>LEARNING LAB II – Forecasting with Time Series</b>

Statistical analysis is taught using the JMP™ desktop data discovery tool from SAS Institute Inc., as Capital One is the largest user of SAS products on the East Coast and has a strategic IT initiative to promote the use of JMP over spreadsheets for statistical analysis. We also cover Monte Carlo simulation in the course, partially as a useful technique in its own right and partially to introduce various statistical concepts without requiring extensive theoretical development. Monte Carlo simulation is taught using the @Risk add-in to Excel™ from Palisade Corporation.

Most of the concepts are taught through exercises using the software and the participants are forced to explain the concepts in their own words to strengthen their

understanding. Throughout the course we discuss the risks of drawing conclusions from data. For example, the discussion of testing a single mean is motivated as an argument between an analyst and a vice president. One can decide to argue against the vice president and be shown to be wrong later or one can decide to stay quiet and be shown to be right later. What are the risks of either course of action? When discussing the various criteria for predictive models, we use an example in which the true model is known, but we give participants limited data. Under these conditions, the participants always find models that fit the data well and that indeed give good predictions, but they miss some terms in the model. This is a useful demonstration that they can build models that are good predictors, but they will not necessarily find the exact “true” model. While there are risks and mistakes can be made even using careful statistical analysis, the process is usually robust and the predictive inputs to their decisions are much improved.

Analysts ask many questions of data sets. To reinforce this questioning in a way that is more natural for inexperienced statisticians, we re-cast hypothesis testing in terms of a standard form of question: “I assume that ...; do the data support this?” with the answer coming in the form of a p-value. We adopt this standard form of questioning throughout the course to provide a consistent framework across all forms of statistical tests: “I assume the population mean is \$100...” for testing a single mean or “I assume the means for the five groups are equal...” for analysis of variance. This removes the need for a theoretical discussion about null and alternative hypotheses. We demonstrate the concept of the p-value using Monte Carlo simulation. Thus, we avoid theoretical discussions, but nevertheless instill an understanding of the concepts involved.

Coverage of time series analysis focuses on teaching participants to recognize the signatures of the models and to iteratively fit models until the residuals show no further auto-correlation. We start with the simplest auto-regressive and moving average models, before adding trends and seasonality. The participants simulate each type of model to aid understanding and so they can see the type of time series that each model creates. We teach the participants to look for seasonality with spectral density plots and de-trend a time series with differencing. Participants also learn to compare models using the model fit criteria introduced in the regression discussions. Every effort is made to ensure that course examples and exercises reflect real business problems that the analysts are likely to encounter. For example, in an early offering of this course an analyst came to the instructors to request help with prediction of an important quantity that the analyst had to present to company executives on the evening of Day 4. Examining his data, the instructors found that the data had both trend and seasonality and required a transformation to achieve good predictions, making it a perfect fit for the material to be covered on Day 4. We rewrote the notes for that day, using his data for one line of business as the example. He then gave an apparently excellent presentation that evening with a much better forecast than they had achieved previously. We then used other lines of business with different time series structures as exercises to reinforce these concepts.

At the end of the course, the participants have a conceptual and applied knowledge of forecasting using both regression-based and time series-based models. Several participants also leave the class with specific models that they can use in their work.

### 3.2 Simulation Modeling and Analysis

The Simulation Modeling and Analysis course was taught over three days to assist Capital One’s Center of Excellence in Simulation. An outline of course topics covered can be found in Table II. This course covered discrete event simulation modeling (ranging from simple to advanced), analysis of simulation results, and making decisions using simulation. Two packages, Arena™ and Simul8™, were chosen for the course. Capital One had several projects underway using Simul8™, however they were also interested in Arena™ from Rockwell Software, as the interface resembled the business process maps that many teams were developing.

**Table II: Course Outline for Simulation Modeling and Analysis**

<b>Day 1</b>	<p><b>Basic Simulation Modeling with Simul8</b>  <i>Building blocks of simulation</i>  <i>Building a simple model</i>  <i>Basic output analysis</i></p> <p><b>Intermediate Simulation Modeling with Simul8</b>  <i>Controlling the flow of entities</i>  <i>Sharing and scheduling resources</i>  <i>Variable arrival patterns</i></p>
<b>Day 2</b>	<p><b>Advanced Simulation Modeling with Simul8</b>  <i>Programming with Visual Logic</i>  <i>Modeling costs and profits</i></p> <p><b>Simulation Modeling with Arena</b>  <i>Re-creating the basic model in Arena</i>  <i>Re-creating the intermediate model in Arena</i></p>
<b>Day 3</b>	<p><b>Input Analysis</b>  <i>Fitting distributions to data</i>  <i>Stat::Fit for Simul8</i>  <i>Input Analyzer for Arena</i></p> <p><b>Making Decisions with Simulations</b>  <i>Selecting the Best System &amp; Process Analyzer</i>  <i>Optimization with Opt Quest and Simul8</i>  <i>Optimization with Opt Quest and Arena</i></p>

We wanted participants to begin building models early in the course. The course started with building a simple model in Simul8, as we believed the start up time would be shorter than with Arena. After introducing the first modeling scenario, we quickly introduced the building blocks in Simul8 and had the participants build their first model. This allowed them to have a running animated simulation in the first two hours of the course. We then discussed whether the simulated system would operate constantly or close down and introduced how to perform analysis for steady-state and terminating systems. The first analysis exercise had the participants vary the number of replications, the length of each simulation run, and the length of the warm-up period. We then performed our first proper analysis, deciding whether to purchase a new faster Resource based on processing, staff, and machine costs.

The second model was a call center, of obvious interest to Capital One. The scenario included different types of calls, multiple stages of call processing, shared and scheduled resources, and variable call volumes throughout the day. Each of these features requires the use of further options, but can be performed with the same basic building blocks in Simul8. The participants were surprised at how quickly they could build relatively complex models. This section was followed by multiple exercises involving making modifications to the modeled system, performing analysis, and making decisions with the model.

After a day and a half, the participants could build complex models in Simul8 and perform the appropriate analysis to make good decisions using these models. This made the introduction of Arena simpler, as we could build the same basic and intermediate

models introducing the equivalent features. This often led to discussions of which situations were easier to model in Simul8, and which were easier to model in Arena.

In terms of decision making, only Arena is sold with an add-in for selecting the best simulated system, called Process Analyzer. The discussion of selecting the best system from a small list of options was useful, as it allowed us to test the number of simulations necessary to make defensible decisions. We then moved on to optimization. Luckily, both Simul8 and Arena use Opt Quest for simulation optimization. Thus, we could build the same simulation and optimization model and discuss the few interface differences between the two implementations.

At the end of three days, the participants could build complex models in two simulation packages, perform appropriate input and output analysis, and build optimization models to make decisions.

### *3.3 Optimization and Decision Techniques*

The Optimization and Decision Techniques course lasts six days and covers spreadsheet optimization modeling, including basic linear models, models with integer decision variables, and non-linear models. We also introduce stochastic optimization and simulation optimization to integrate with the other two training courses. An outline of course topics covered can be found in Table III. The main software package used is Frontline System's Premium Solver™, but we also introduce @Risk, JMP and Simul8 for integrated analysis and teach Risk Optimizer for turning @Risk models into stochastic optimization models.

Rather than simply presenting a “procedure” for modeling linear programs, we encourage participants to think through the process by use of examples.

**Table III: Course Outline for Optimization and Decision Techniques**

<b>Day 1</b>	<b>Linear Programming Models</b> <i>Formulating LP models in Excel</i> <i>Solving LP models with Solver</i> <i>Language, assumptions, &amp; properties of LP</i> <b>Learning Lab 1</b>
<b>Day 2</b>	<b>Solving LP Models</b> <i>Solution outcomes &amp; Solver messages</i> <i>Concepts of the simplex method</i> <i>Algebraic models</i> <b>Learning Lab 2</b>
<b>Day 3</b>	<b>Integer Programming</b> <i>Using binary variables</i> <i>Modeling logical constraints &amp; fixed costs</i> <i>Concepts of branch and bound</i> <i>Choosing an appropriate IP model</i> <b>Learning Lab 3</b>
<b>Day 4</b>	<b>Network Models</b> <i>Types of network models</i> <i>Solving network problems with Solver</i> <b>Non-linear Models</b> <i>Non-linear models</i> <i>Local vs. global solutions</i> <i>Meta-heuristics</i> <b>Learning Lab 4</b>
<b>Day 5</b>	<b>Stochastic Optimization</b> <i>Optimization models with uncertainty</i> <i>Probability models for specific situations</i> <i>Fitting distributions to data</i> <i>Solving with @Risk and Risk Optimizer</i> <b>Learning Lab 5</b>
<b>Day 6</b>	<b>Stochastic Optimization</b> <i>Forecasting random inputs with JMP</i> <i>Modeling work flow in Simul8</i> <i>Optimizing work flow with OptQuest</i> <b>Learning Lab 6</b>

To illustrate, before building our first LP model, participants are presented with a scenario that gives them cost and benefit information for placing weekday and weekend ads on a premium sports channel (Figure 1). They are then asked to determine how best to utilize existing resources to maximize return. While the scenario presented is

somewhat simplistic, it is a reasonable place to start with analysts new to building LP models. Rather than using this example to illustrate the structure of an LP, we begin by asking analysts: how would you approach this problem? After giving them time to reflect and discuss, the conversation inevitably comes back to decisions. Analysts then realize that defining the decisions to be made is a central part of the solution approach, whatever it may be. In defining those decisions, many analysts begin with statements like, “We must decide how many ads to place.” Such statements then lead to a discussion of the importance of defining clear, specific decisions: “We must decide how many weekday ads to place and how many weekend ads to place.” Once the decisions to be made have been clearly articulated, we encourage the participants to develop strategies that lead to maximum-effect decisions. This ultimately leads us to development of an LP model. Once the participants have developed the concepts on their own, we provide them with more structured statements of the concepts and vocabulary.

---

Capital One has an advertising budget of \$7.2 million to spend in the Greater Richmond area. The company decides to use this budget for advertisement on a premium sports channel. They can choose to place daily ads (Monday through Friday) or weekend ads.

<div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p><b>Placing a daily ad</b></p> <ul style="list-style-type: none"> <li>• results in 150 new customers</li> <li>• costs \$9000</li> <li>• contributes \$1500 to overhead</li> </ul> </div>	<div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p><b>Placing a weekend ad</b></p> <ul style="list-style-type: none"> <li>• results in 250 new customers</li> <li>• costs \$16,000</li> <li>• contributes \$900 to overhead</li> </ul> </div>
---	--

Administration has dictated that at most \$1 million be spent on overhead (in addition to the advertising budget). How should Capital One utilize its resources to maximize the number of new customers?

---

© 2003 Virginia Commonwealth University. All Rights Reserved.

**Optimization** ≡

**Models**

**Figure 1: This basic modeling scenario was used to introduce analysts to the thought processes required to build optimization models. Analysts were encouraged to develop their own approaches before the class at large discussed the problem and were introduced to their first linear program.**

The goal is not to teach participants how to execute the simplex algorithm—or even provide them with a detailed description of its inner workings—but to help them understand how it takes advantage of key LP properties and why it works particularly well. At its core, this discussion aims to prepare the participants to understand messages returned by solution software packages, particularly when errors are encountered or optimal solutions cannot be found. We also want the participants to be aware that algebraic modelers are available and can be more suitable than Solver for some applications. Given the difficulty participants seem to have with this topic, however, we leave further practice on the exercises we provide to the individual participant. We cover integer programming and network models in a similar manner. We spend some time with nonlinear optimization and conclude with a basic description of metaheuristics such as tabu search, simulated annealing, and genetic algorithms. Our aim is to help them understand the kinds of methods that are used, rather than teaching them precisely how they work.

The last two days aim to integrate the optimization course with the simulation and forecasting courses. In the forecasting course, we cover regression based forecasting with JMP and Monte Carlo simulation with @Risk. We review these topics for the participants who have not taken the other course and then show how these tools can be integrated into an optimization model with Risk Optimizer. In the simulation course, we cover discrete event simulation in Simul8. Again, we briefly review building simple models and analyzing simulation output, before showing how we can overlay an optimization model with OptQuest. The focus throughout is practical problem solving, highlighting that the software finds solutions using the metaheuristics already discussed.

The course is structured so that all topics are covered through example, usually giving the participants a chance to explore modeling and solution methods independently before the formal techniques are presented. This gives the participants insight into just how complex such models can be, in turn increasing interest in understanding the underlying theory.

### *3.4 The Learning Labs*

The Learning Labs underscore the relevance of the techniques covered in each course by applying them to current business problems. Ideally, participants would bring in their own problem and all relevant data to a Learning Lab, introduce the scenario to the class, and then work with the class to develop a solution, either as a whole class, or in small groups. Most analysts have performed basic statistical analysis and take the Forecasting and Time Series course to improve the accuracy of their forecasts, thus implementing Learning Labs in this manner is natural in this course. On the other hand, many analysts have little experience with optimization, thus they do not always recognize its applicability in their work. Consequently, it is difficult for them to bring a defined business problem that lends itself to optimization. Learning Labs in the Optimization and Decision Techniques course, then, are cases adapted from real business applications by the internal trainers and turned in to manageable cases by the external instructors. Simulation presents a rather different set of circumstances, since a model to solve an actual business problem would take too long to develop in a three-day course. Learning Labs were impractical here, so the instructors assisted the analysts with their individual projects after the course.

In the Forecasting and Time Series course, participants are encouraged to bring in their own data for the Learning Labs. While we do pre-screen the data to ensure that the problem can be solved using the techniques covered thus far, the problem and data are presented to the class by the problem owner. In some cases, the participants can easily see the solution approach; in others, the instructors must use leading questions to steer them in the right direction. The participants must suggest the direction that the model fitting will take, but, working as a group, they can solve these relatively complex problems. Taking another line of business, the participants may then fit their own models, often with very different results, as behavior in the different lines is quite different. To appeal to the competitive nature of the participants, we often set up a competition with a small prize for the best value of a given model fit criteria. In many cases, we have kept a portion of the data for cross-validation and we can see how well the participants' models actually predict new data. Thus, the Learning Labs are used to reinforce important concepts while allowing participants to solve important business problems during the class.

As an example, in the second forecasting class, one analyst had the task of forecasting the percentage of accounts that would not be paid on time and presenting the forecasts to his senior management team three days into the course. He supplied the data, we wrote the example into the course notes, and the analyst presented the forecast on time. We then used other lines of business in a learning lab, finding significantly different models. In another class, a model to forecast incoming daily mail volumes was developed by the class. The model was a significant improvement on their current forecasting model. Of course, the Learning Labs are not always successful in improving upon the

analysis already performed within Capital One. For example, an attempt by another class to improve business loss forecasting in the UK business was not able to improve on a model developed by consultants.

Learning Labs for the optimization course are structured so that the skills acquired that day—in addition to previously learned skills—are used in the solution process. Each lab builds on the previous one, and requires participants to consider their tools carefully in order to choose an appropriate approach. All labs are taken from real Capital One applications, though the early labs are simplified in order to accommodate participants' developing skills. The ongoing case study developed a model to determine the optimal allocation of collections call center associates to lines of business. The final model included decisions concerning site expansions and optimal use of outsourcing to external vendors.

#### ***4. The Success of the Courses***

To evaluate the effectiveness of the training, and in an attempt to continually improve each course, we asked the participants to assess their experiences under multiple criteria. We also followed up the first 9 months of training with a business impact survey, as the courses are only truly effective if the attendees are using the techniques in their work and achieving improved results (i.e. a return on the investment). We had 55 responses to the course evaluations of a possible 84 and 20 responses to the business impact survey.

The participant evaluations consisted of a set of statements, to which attendees were asked to respond on a 5-point scale from strongly agree to strongly disagree. This method of evaluation is standard for all Capital One training, but we added questions to assess the effect of our pedagogy. 98% of attendees liked the balance between lectures,

exercise, cases, and learning lab activities. Only 6% of attendees found the coverage unclear. All attendees agreed that the courses were interesting and that participation was encouraged, while 53% strongly agreed with both statements.

Overall, all attendees would recommend the instructors to others for training, and 55% would strongly recommend them. 94% of attendees would recommend the training to others, while 34% would strongly recommend the training. Respondent comments revealed that the 6% who would not recommend the training felt that the course did not meet their expectations in terms of content. In the follow-up business impact survey, 90% of the analysts saw themselves using the techniques covered within their team, 85% in the near term.

The success of the courses can be described in more than participant appreciation, however. There has been direct return on investments, although not all such return has been estimated currently. The learning lab sessions allow analysts to bring in their own problems and to work on data in class. The team members responsible for these forecasts were excited about improving the quality of their forecasting processes. Furthermore, one example of return on investment has been quantified. In the first forecasting course, a discussion took place about allocating resources in a payment processing mail center. One of the analysts involved then took the simulation course. Following the course, The third author facilitated an effort to develop a decision support tool using a discrete event simulation. The first author developed the simulation model and front and back end spreadsheets to allow the analysts to perform tests of multiple resource allocation scenarios. The analysts used this tool, along with the analytical techniques learned in the two courses, to save more than \$2 million annually in costs. Capital One used this

example as part of their application for the Best Training Company in America awarded by *Training* magazine, placing twelfth overall and first among financial services companies. Thus, one project spawned from this training has already paid for all the training many times over.

## ***5. The Elements of the Success***

The success of the course can be largely attributed to the partnership developed between champions within Capital One and outside trainers. The internal trainers include a PhD from Cornell and a Six-Sigma black belt. However, as both are financial services analysts, not educators, they would not consider themselves up to date in the best approaches to teaching analytical methodologies. The external trainers maintain a focus on excellence in teaching and have applied the latest teaching methods to improve their university courses, but they are not experts in the financial services business and do not have access to relevant data and business problems. Thus, it is the inside-outside partnership that provides all the elements of success in this endeavor.

The inclusion of trainers from inside Capital One allows us to put the material in a business context. The material is translated into the corporate language, using appropriate media. The PowerPoint style we use is in line with training the analysts receive from Capital One in preparing good presentations. The problems can be described in the corporate language, even safely using their own plethora of acronyms that are so prevalent in the corporate world. The inside trainers are both in roles at Capital One that give them a wide appreciation of the various types of analysis that are being or should be performed at Capital One. Thus, they can suggest other ways the tools can be used by the analysts. They also provide follow-up consulting to reinforce the application of the

techniques. Lastly, they are internal champions of the courses, ensuring that analysts in Capital One are aware of the opportunities and helping them decide which courses would be useful in their roles. It should be noted, however, that such broadly experienced analysts do not necessarily have time to perform all of the course development.

One advantage of including an outside component in the training team is a wider view of the material. University faculty members are more experienced at teaching the techniques we cover and are aware of the latest teaching methods. They have a broad knowledge of software tools for implementing the techniques and can assist in selecting appropriate software, even beyond the needs envisioned within the company. There is also an advantage in terms of resources. Training facilities can be scarce inside a company, so we use the university's computer labs. These labs are outfitted with appropriate technology for effective teaching, like computers for participants and instructors and projection equipment. They also have most major analytical software, allowing the training to include the *right* software rather than the *available* software.

## **6. Conclusion**

The training has proved popular with analysts and has been effective in propagating the techniques taught throughout Capital One. We continue the endeavor to assess the return on investment of projects assisted by the training. As we have discussed, early indications are extremely promising, and both Capital One and VCU have benefited greatly. While the pedagogy is effective in the training environment, combining inductive learning with case studies and consulting sessions, the partnership between university faculty and industry analysts has greatly improved the business impact and is worth considering in all corporate training efforts.

## **References**

- Albritton, M. D., P. R. McMullen, L. R. Gardiner. 2003. OR/MS Content and Visibility in AACSB-Accredited US Business Programs. *Interfaces* **33**(5) 83-89.
- Board, J., C. Sutcliffe, W. T. Ziemba. 2003. Applying Operations Research Techniques to Financial Markets. *Interfaces* **33**(2) 12-24.
- Böcker, F. 1987. Is case teaching more effective than lecture teaching in business administration? An exploratory analysis. *Interfaces* **17**(5) 64–71.
- Bodily, S. E. 1996. Teacher's forum: Teaching MBA quantitative business analysis with cases. *Interfaces* **26**(6) 132–138.
- Carraway, R., D. Clyman. 1997. Managerial relevance: The key to survival for OR/MS. *Interfaces* **27**(6) 115–130.
- Corner, J., P. D. Corner. 2003. Teaching OR/MS Using Discussion Leadership. *Interfaces* **33**(3) 60-69.
- Kolb, D. 1984. *Experiential Learning—Experience as the Source of Learning and Development*. Prentice-Hall, Englewood Cliffs, NJ.
- Lasdon, L., J. S. Liebman. 1998. The teacher's forum: Teaching nonlinear programming using cooperative active learning. *Interfaces* **28**(4) 119–132.
- Powell, S.G. 2001. Teaching Modeling in Management Science, *INFORMS Transactions on Education* **1**(2) 62-67 <http://ite.informs.org/Vol1No2/Powell/>.
- Saltzman, R. M., V. Mehrotra. 2001. A Call Center Uses Simulation to Drive Strategic Change. *Interfaces* **31**(3) 87-101.
- Sanders, N. R., K. B. Manrodt. 2003. Forecasting Software in Practice: Use, Satisfaction, and Performance. *Interfaces* **33**(5) 90-93.

Scott, J., J. Buchanan. 1992. Teaching management science: Hold the lectures. *OR/MS Today* **19**(October 5) 46–50.

Seal, K. C., Z. H. 2003. Przasnyski. Using Technology to Support Pedagogy in an OR/MS Course. *Interfaces* **33**(4) 27-40.

Trench, M. S., S. P. Pederson, E. T. Lau, L. Ma, H. Wang, S. K. Nair. 2003. Managing Credit Lines and Prices for Bank One Credit Cards. *Interfaces* **33**(5) 4-21.