Partnerships in Training

Jason R. W. Merrick, Jill R. Hardin

Department of Statistical Sciences and 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 supplies credit products in the US. Since 1995, it has followed an information-based strategy, collecting and analyzing data to improve its business decisions. To continue this successful strategy, it must train its analysts in analytical techniques. In a partnership between Capital One and the Department of Statistical Sciences and Operations Research at Virginia Commonwealth University, we developed three training courses in forecasting, optimization, and simulation. These courses have been well received by the analysts and have led to considerable early return on investment with one class project forecasted to reduce costs at the firm's Richmond, Virginia mail center by more than $2 million annually.

Key words: professional: OR/MS education; decision analysis: applications.
History: This paper was refereed.

Credit card companies are in an ideal position to exploit operations research and management science (OR/MS) methods, such as 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 them to make decisions about offering credit cards, credit limits, and interest rates (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 forecast diverse quantities, such as their customers' nonpayment rates, incoming mail and phone-call volumes, and response rates to mail solicitations. They use the forecasts in making complex decisions on marketing strategies, delinquency cures, call center staffing, and other matters that require sophisticated analysis. Many of their internal systems, such as mail and call centers, are dynamic in nature and must be robust to the variable incoming volumes (Saltzman and Mehrotra 2001).

Capital One direct markets 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 franchises of financial services. Its use and management of data is the key to its success, making research critical to maintaining continued growth. In one current research program, it seeks to evaluate new data sources and analysis techniques with the aim of introducing new methods into its business-decision processes. It must cultivate its talent pool to improve employees' knowledge of decision techniques, including forecasting, optimization, and simulation, to facilitate the improvement of the models it uses to manage accounts.

Capital One commonly compares its hiring process for analysts to the National Basketball Association's (NBA) recruiting process. The NBA often drafts high-school athletes, focusing more on talent and athleticism than on making sure its draft choices have four years of college playing. Capital One hires MBAs from top business schools who have had one or two quantitative courses and then trains them to perform high-level analyses. While most of these MBA graduates have been exposed 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 methods. Although Capital One hires some graduates with

PhDs and master's degrees in statistics, it hires few with degrees in OR/MS.

Instead of hiring advanced OR/MS graduates to act as internal consultants, Capital One trains all of its analysts to use applicable analytical methods. However, it recognizes that its employees need further knowledge in the analytical methods it seeks to leverage. Many training courses are available. Most vendors of analytical software provide extensive training in the use of their software, and many will customize the courses to meet customers' needs. Furthermore, universities offer courses, often in the evening, to provide theoretical background. Still, few courses cover software from multiple vendors and supply enough background to ensure their students apply the underlying methods correctly. Furthermore, companies need to make the material relevant to their business problems (Carraway and Clyman 1997) through courses based on the company's examples and case studies. Analysts then learn how to use the techniques and when to use them, creating new opportunities to use the techniques in their own enterprises rather than simply repeating the kind of analyses others have used.

Capital One turned to the operations research faculty at Virginia Commonwealth University (VCU). VCU is a public research university in Richmond, Virginia, where Capital One has a large number of offices. As an urban university, VCU welcomes opportunities to combine academic and real-world education. In the partnership between VCU and Capital One, VCU provides the facility for training and instructors knowledgeable about teaching analytical material.

The training team consists of faculty from VCU and analysts from Capital One. The VCU trainers bring an up-to-date knowledge of pedagogy to the venture; the Capital One trainers bring business relevance and access to interesting data and case studies. The close geographical proximity allows for a tight partnership, with Capital One and VCU personnel working together to develop and teach the courses and to improve the courses continually in response to official and unofficial feedback.

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.

The Pedagogical Approach

We designed training courses for Capital One's analysts to use three analytical techniques—forecasting, optimization, and simulation—to solve business problems. In fact, because analysts do not solve problems using just one technique, analysts need to know how to combine various techniques to solve problems. For example, optimization models often use forecasts based on regression or time series as inputs. Furthermore, the systems analysts seek to optimize often include stochastic delays that are best modeled using discrete-event simulation. Thus, while we do not encourage analysts to take the courses in a particular order, we try to integrate techniques developed in any one course with skills taught in other courses.

We considered Capital One's strategic objectives in developing the learning goals for the courses. Capital One wanted its own analysts to run successful projects using these three 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 courses had taken one or two quantitative analysis courses in business or engineering schools and had some basic knowledge of the techniques and software. The last learning goal—knowing when to use each technique—is hardest to teach.

Our solution to ensuring comprehensive coverage while targeting the firm's analytical needs is to use both Capital One and VCU trainers. Business relevance is a key aspect of the training. Business educators recognize the importance of using relevant examples and applications (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 analysts to see how to apply the methods in their own work. Therefore, all the course examples and exercises we use are financial or logistical in nature, and nearly all of them are based on Capital One operations, with the applications and data drawn from various Capital One business areas.

Another valuable feature of the training courses is the vendor-neutral approach we take to selecting software. Sanders and Manrodt (2003) surveyed types of software US corporations use in forecasting and the relative success obtained with various packages. One of their major findings was that the use of commercially available software in applying the techniques improved the accuracy of the forecasts and user satisfaction. Researchers might obtain the same findings for other techniques, such as optimization and simulation.

Software vendors understandably wish to promote the use of their packages, yet as most analysts would confirm, different packages are suited to different applications. Thus, a comprehensive coverage of suitable tools is useful to the analyst. Furthermore, the suite of packages they often need for complete solutions are usually not available from any one vendor. Corporate analysts cannot attend training sessions for multiple software packages and university courses covering the methods when they need to implement analytical methods in a timely and cost-effective manner.

We focus on modeling problems to give analysts an appreciation of how they will use the tools covered in a course (Powell 2001). We do not describe the methods one by one, nor do we demonstrate software menu item by menu item. We 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 methods 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, participants must prepare the case before class by reading it or preparing solutions for it. In the training environment, instructors can seldom require such preparation. Like many other OR/MS education

researchers (Scott and Buchanan 1992), we do not want to resort to straightforward lectures. We adapt by integrating interactive discovery techniques with presentation of information and discussion, acknowledging the fact that participants have little time to study outside of the classroom.

We develop teaching material in PowerPoint, Capital One's preferred communications medium, and we project it in the class and print it out for their reference. We teach in the university's computer classrooms, so we project our software demonstrations and the participants follow along on desktop computers to gain basic experience with the software. Even though the material is relevant and applicable to their work, the participants have limited attention spans in an eight-hour-a-day, three-to-six-day training course, so we keep the presentation segments under 20 minutes and concentrate on active learning. We intersperse frequent, simple, reinforcing exercises to help participants solidify concepts before we move to new material. In this way, we incorporate independent study and thought in the classroom, making up for the lack of out-of-class assignments. After the primary coverage and reinforcement exercises, the analysts can confidently apply the techniques, yet many attendees are still uncertain about how to apply them in their own work. We choose the examples and exercises carefully to illustrate key points and to ensure appropriate complexity: participants can solve them in the available class time. Unfortunately, the analysts often do not grasp the applicability of the techniques to the problems they face. With this in mind, we go on to the learning labs, a less structured portion of the course in which we use a modified version of the case method.

The learning-lab cases are scenarios drawn from Capital One business problems, often problems faced by analysts attending the course. Because cooperative learning improves the students' understanding of concepts and interest in the material (Lasdon and Liebman 1998), we encourage the analysts to work in teams on the cases, calling on the trainers for help when necessary and on each other for help in understanding the business context. As teams encounter obstacles, we bring the class back together to discuss the problem.

In many learning labs, we develop a group solution in discussion for one line of business and then

have the participants perform their own analyses for other lines, essentially making a real-time case. They observe the kinds of questions the instructors ask to understand the problem and find a reasonable solution approach. If the techniques we cover in the course are not suitable, we can explain why and outline other techniques that might be appropriate. If the techniques are suitable, the participants can perform the relevant analysis, making mistakes and iterating until they achieve a reasonable solution. For many participants, the learning labs are the most useful and enjoyable parts of the training courses, because they are not seeking tidy solutions to prepared teaching examples. Learning labs reinforce the second learning goal, knowing how to implement the techniques, but they move the participants toward the third learning goal: knowing when to use them.

The Course Designs

To ensure that we considered all of Capital One's learning goals in our courses, we developed the following training objectives for our courses:

—To inform participants of the various types of models available and the issues surrounding each model,

—To teach participants to build correct models,

—To help participants understand how to obtain solutions and why some models are more difficult to solve than others, and

—To guide participants in selecting an appropriate model for a given scenario.

We have designed three training courses for Capital One:

—Time series and forecasting,

—Simulation modeling and analysis,

—Optimization and decision techniques.

Time Series and Forecasting

"Time series and forecasting" is a five-day course, including a review of making decisions with statistics and detailed coverage of regression-based forecasting and time-series forecasting (Table 1). We teach statistical analysis using the JMP desktop data discovery tool from SAS Institute Inc. (Capital One is the largest user of SAS products on the East Coast of the US, and it promotes the use of JMP over spreadsheets for statistical analysis). We also cover Monte Carlo

Day 1	Statistics review
	Answering questions with data
	Introduction of JMP, @Risk
Day 2	Introduction to regression analysis
	Simple regression
	Multiple regression
	Model-fitting criteria (AIC, SBC, Cp)
	Stepwise regression
	Learning Lab 1—Forecasting with regression
Day 3	Time-series analysis 1
	Autocorrelation
	Autoregressive and moving-average models
	JMP's time-series platform
	@Risk simulation of time-series models
Day 4	Time-series analysis 2
	Modeling trends
	ARIMA models
	Modeling seasonality
	SARIMA models
	Using transformations
	@Risk simulation of time-series models
Day 5	Time-series analysis 3
	Exponential smoothing models
	Combining time series and regression
	Learning Lab 2—Forecasting with time series

Table 1: We teach the "Time series and forecasting" course over five days. In the first half, we cover regression-based forecasting and end with a learning lab. In the second half, we cover time-series forecasting and end with a second learning lab.

simulation in the course, partly as a useful technique and partly to introduce various statistical concepts without extensive theoretical development. We teach Monte Carlo simulation using the @Risk add-in to Excel from Palisade Corporation.

We teach most concepts through exercises using the software. We force the participants 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, we portray the discussion of testing a single mean as an argument between an analyst and a vice president. The analyst can decide to argue against the vice president and be shown to be wrong later or 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 give good predictions, but they miss some terms in the model. This useful demonstration shows that they can build models that are good predictors, but they will not necessarily find the exact true model. While they face risks and can make mistakes even when 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 their questioning in a way that is natural for inexperienced statisticians, we recast hypothesis testing in terms of a standard form of question: "I assume that . . . ; do the data support this assumption?" 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. We then avoid the theoretical discussion about null and alternative hypotheses. We demonstrate the concept of the p-value using Monte Carlo simulation. Although we avoid theoretical discussions, we nevertheless instill an understanding of the concepts involved.

In covering time-series analysis, we focus on teaching participants to recognize the signatures of the models and to iteratively fit models until the residuals show no further autocorrelation. We start with the simplest autoregressive and moving-average models, before adding trends and seasonality. The participants simulate each type of model so they can see the type of time series each model creates. We teach them to look for seasonality with spectral density plots and to detrend a time series with differencing. They also learn to compare models using the model fit criteria we introduced in discussing regression. We try to ensure that the 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 asked for our help in predicting an important quantity to be presented to company executives on the evening of Day 4. Examining his data, we 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. Apparently, he gave an excellent presentation that evening with a much better forecast than any they had obtained previously. We used other lines of business with different time-series structures as exercises to reinforce the concepts.

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

Simulation Modeling and Analysis

We taught the "Simulation modeling and analysis" course over three days to analysts working with Capital One's Center of Excellence in Simulation (Table 2). We covered discrete-event-simulation modeling (simple to advanced), analysis of simulation results, and using simulation to make decisions. We used Arena and Simul8 for the course. Capital One had several projects underway using Simul8; however, it was also interested in Arena from Rockwell Software, because

Day 1	Basic simulation modeling with Simul8
	The building blocks of simulation
	Building a simple model
	Basic output analysis
	Intermediate simulation modeling with Simul8
	Controlling the flow of entities
	Sharing and scheduling resources
	Variable arrival patterns
Day 2	Advanced simulation modeling with Simul8
	Programming with Visual Logic
	Modeling costs and profits
	Simulation modeling with Arena
	Re-creating the basic model in Arena
	Re-creating the intermediate model in Arena
Day 3	Input analysis
	Fitting distributions to data
	Stat::Fit for Simul8
	Input analyzer for Arena
	Making decisions with simulations
	Selecting the best system and Process Analyzer
	Optimization with OptQuest and Simul8
	Optimization with OptQuest and Arena

Table 2: We teach the "Simulation modeling and analysis" course over three days. In the first half, we cover modeling in Simul8, as this was considered the easier of the two simulation environments to learn. In the second half, we cover modeling in Arena and more advanced simulation techniques.

the interface resembled the business-process maps that many teams were developing.

We wanted participants to begin building models early in the course. The course started with building a simple model in Simul8; we believed that the start-up time would be shorter with Simul8 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. They had 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 we explained how to perform analyses for steady-state systems and terminating systems. In the first analysis exercise, the participants varied the number of replications, the length of each simulation run, and the length of the warm-up period. The participants then performed their 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 realistic models. We followed this model with exercises that required making modifications to the modeled system, performing an 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 analyses to make good decisions using these models. The introduction of Arena was simpler because we could build the same basic and intermediate models while introducing the equivalent features. Participants discussed 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 because we tested the number of simulations necessary with each system to make defensible decisions. We then moved on to optimization. Luckily,

both Simul8 and Arena use OptQuest for simulation optimization. Thus, we could build the same simulation and optimization model with both packages 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.

Optimization and Decision Techniques

"Optimization and decision techniques" lasts six days and covers spreadsheet optimization modeling, including basic linear models, models with integer decision variables, and nonlinear models. We also introduce stochastic optimization and simulation optimization to integrate this course with the other two training courses (Table 3). 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 using examples. For example, before building our first LP model, we present the participants with a scenario that gives them cost and benefit information for placing weekday and weekend ads on a premium sports channel (Figure 1). We then ask them to determine how best to use the 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. Instead of using this example to illustrate the structure of an LP, we begin by asking analysts: how would you approach this problem? After we give them time to reflect and discuss, the conversation inevitably comes back to decisions. The 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 such statements as "We must decide how many ads to place." Such statements 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 they have clearly articulated the decisions to be made,

Day 1	Linear programming models
	Formulating LP models in Excel
	Solving LP models with Solver
	Language, assumptions, and properties of LP
	Learning Lab 1
Day 2	Solving LP models
	Solution outcomes and Solver messages
	Concepts of the simplex method
	Algebraic models
	Learning Lab 2
Day 3	Integer programming
	Using binary variables
	Modeling logical constraints and fixed costs
	Concepts of branch and bound
	Choosing an appropriate IP model
	Learning Lab 3
Day 4	Network models
	Types of network models
	Solving network problems with Solver
	Nonlinear models
	Nonlinear models
	Local versus global solutions
	Metaheuristics
	Learning Lab 4
Day 5	Stochastic optimization
	Optimization models with uncertainty
	Probability models for specific situations
	Fitting distributions to data
	Solving with @Risk and Risk Optimizer
	Learning Lab 5
Day 6	Stochastic optimization
	Forecasting random inputs with JMP
	Modeling work flow in Simul8
	Optimizing work flow with OptQuest
	Learning Lab 6

Table 3: We teach the "Optimization and decision techniques" course over six days. We cover linear programming in the first two days. In the second two days, we cover integer and nonlinear programming. We cover stochastic programming in the final two days. Each day ends with a learning lab for which the cases build cumulatively in complexity.

we encourage them to develop strategies that lead to maximum-effect decisions. This activity ultimately leads us to develop 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.

Our goal is not to teach participants how to execute the simplex algorithm—or even to 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. In

this discussion, we try to prepare the participants to understand messages returned by software solution packages, particularly when they encounter errors or cannot find optimal solutions. We also want the participants to be aware that algebraic modelers are available and can be more suitable than Solver for some applications. Because the participants find this topic difficult, however, we leave them to practice further on the exercises we provide on their own. We cover integer programming and network models in a similar manner. We spend some time on nonlinear optimization and conclude with a basic description of metaheuristics, such as tabu search, simulated annealing, and genetic algorithms. We aim to help the analysts to understand the kinds of methods that are used rather than teaching them precisely how they work.

In the last two days of the course, we try 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 forecasting course and then show them how to integrate these tools into an optimization model with Risk Optimizer. In the simulation course, we cover discrete-event simulation in Simul8. For the participants in the optimization course, we review building simple models and analyzing simulation output before showing how we can overlay an optimization model with OptQuest. Our focus throughout is practical problem solving, highlighting that the software finds solutions using the metaheuristics already discussed.

We structure the course to cover all topics through example, usually giving participants a chance to explore modeling and solution methods independently before we present the formal techniques. Participants gain an understanding of how complex such models can be and an interest in understanding the underlying theory.

The Learning Labs
In the learning labs, we underscore the relevance of the techniques covered in each course by applying them to current business problems. Ideally, participants bring in their own problems and all relevant data to the learning lab, introduce the scenarios to

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.

Placing a daily ad	Placing a weekend ad
• Results in 150 new customers	• Results in 250 new customers
• Costs $9,000	• Costs $16,000
• Contributes $1,500 to overhead	• Contributes $900 to overhead

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?

Optimization models

3

Figure 1: We used this basic modeling scenario to introduce analysts to the thought processes required to build optimization models. We encouraged them to develop their own approaches before the whole class discussed the problem and we introduced their first linear program.

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 "Time series and forecasting" 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 and do not always recognize its applicability in their work. Consequently, they have difficulty bringing in defined business problems that lend themselves to optimization. For the learning labs in the "Optimization and decision techniques" course, we use cases adapted from real business applications by the Capital One trainers and turned into manageable cases by the VCU instructors. Simulation presents a different set of circumstances because developing a model to solve an actual business problem would take too long for a three-day course. Learning labs were impractical for simulation, so the instructors helped the analysts with their individual projects after the course.

In "Time series and forecasting," we encourage participants to bring in their own data for the learning labs. While we prescreen the data to ensure that the problem can be solved using the techniques covered thus far, the problem owner presents the problem and data to the class. 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 fairly complex problems. Taking another line of business, the participants may then fit their own models, often with very different results, as behavior differs in the various lines. To appeal to the competitive nature of the participants, we often set up a competition with a small prize for the model with the best value of a given model-fit criterion. In many cases, we keep a portion of the data for cross-validation, and we can see how well the participants' models actually predict new data. Thus, we used the learning labs to reinforce important concepts while allowing participants to solve important business problems during the class.

In the second offering of the forecasting class, for example, 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, developing different models. In another class, the analysts developed a model to forecast daily volumes

of incoming mail. The model was a great improvement on their current forecasting model. Of course, the learning labs are not always successful in improving upon the analyses already performed within Capital One. For example, another class attempting to improve business-loss forecasting in the UK business was unable to improve on a model developed by consultants.

We structure the learning labs for the optimization course so that the analysts use the skills they acquire that day—in addition to previously learned skills—in the solution process. Each lab builds on the previous one and requires participants to consider their tools carefully to choose an appropriate approach. We take all the labs from real Capital One applications, although we simplify the early labs to accommodate participants' developing skills. In the ongoing case study, the participants develop a model to determine the optimal allocation of the associates in the collections call center to lines of business. The final model included decisions concerning site expansions and the optimal use of outsourcing to external vendors.

The Success of the Courses

To evaluate the effectiveness of the training courses and to continually improve them, we asked the participants to assess their experiences under multiple criteria. We also followed up the first nine months of training with a business-impact survey; the courses are truly effective only if the attendees use the techniques in their work and achieve improved results (a return on the investment). We had 55 responses to the course evaluations out of a possible 84 and 20 responses to the business-impact survey.

The participant evaluations consisted of a set of statements, to which we asked the attendees to respond on a five-point scale from strongly agree to strongly disagree. This method of evaluation is standard for all Capital One training programs, but we added questions to assess the effect of our pedagogy. Ninety-eight percent of attendees liked the balance among lectures, exercises, cases, and learning lab activities. Only six percent of attendees found the coverage unclear. All attendees agreed that the courses were interesting and that participation was encouraged (53 percent strongly agreed with both statements).

Overall, all attendees would recommend the instructors to others for training, and 55 percent would strongly recommend them. Ninety-four percent of attendees would recommend the training courses to others, while 34 percent would strongly recommend them. Respondent comments revealed that the six percent who would not recommend the training courses felt that the course did not meet their expectations in terms of content. In the follow-up business-impact survey, 90 percent of the analysts saw themselves using the techniques covered within their team, and 84 percent in the near term.

The success of the courses is revealed in other ways than participant appreciation. Capital One has gained direct return on its investment, although it has not yet estimated all such return. The learning lab sessions allow analysts to bring in their own problems and work on data in class. The team members responsible for the improved 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, one discussion concerned 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 the analyst's team's effort to develop a decision-support tool using a discrete-event simulation. The first author developed the simulation model and front-end and back-end spreadsheets to allow the analyst's team to perform tests of multiple resource-allocation scenarios. The analyst used this tool and the analytical techniques he had learned in the two courses to save more than $2 million annually in costs. Capital One used this example as part of its entry for the "Best Training Company in America" presented by *Training* magazine, placing 12th overall and first among financial-services companies. Thus, one project spawned in these training courses has already paid for all the training expense many times over.

Elements of the Success

The success of the course can be attributed largely to the partnership between champions within Capital One and the VCU trainers. The Capital One 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 experts in teaching analytical methodologies. The VCU trainers focus on excellence in teaching and apply 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 partnership that provides all the elements for success in this endeavor.

The inclusion of trainers from Capital One put the material into a business context. They translate the material into the corporate language, using appropriate media. The PowerPoint presentation style we use is in line with training that Capital One provides the analysts for preparing good presentations. We describe the problems in the corporate language, even safely using the plethora of acronyms that are so prevalent in the corporate world. The Capital One trainers are both in corporate roles at Capital One that give them a wide appreciation of the various types of analyses that are being or should be performed at Capital One. Thus, they can suggest other ways analysts can use the tools. 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 to decide which courses would be useful in their roles. These broadly experienced analysts do, however, not necessarily have time to do all of the course development.

One advantage of including people from outside Capital One in the training team is their wide view of the material. University faculty members are experienced in teaching the techniques and are aware of the latest teaching methods. They have a broad knowledge of software tools for implementing the techniques and can help analysts to select appropriate software, even beyond the needs envisioned within the company. The university also has the advantage of resources—computer labs outfitted with the technology needed for effective teaching, such as computers for all participants and instructors and projection equipment. It also has most major analytical software, allowing training on the right software, rather than the available software.

Conclusion

The training has proved popular with analysts and has been effective in propagating the techniques taught throughout Capital One. We continue to try to assess the return on investment for projects fostered in the training courses. The early indications are 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 courses' business impact, and similar partnerships are worth considering for 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 Trans. Ed.* **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.

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.

Jean Gasen, Senior Manager, Core Learning, Capital One University, Capital One Services, Inc., Attn: 12070-370, 15000 Capital One Way, Richmond, VA

23238, writes: "I am pleased to support the publication of this paper that details the partnership between Capital One Services, Inc. and Virginia Commonwealth University. The paper represents the culmination of our work to develop customized training within the company that tailors learning to the needs of associates in specialized domains of expertise.

"In developing and sponsoring training for corporate associates at Capital One, we seek to achieve multiple goals that positively impact the associates and the enterprise. In evaluating classes, we apply the four levels of training evaluation from Kirkpatrick (1994) to our training programs. The four levels from Kirkpatrick are:

(1) *Reactions*—Evaluation at this level measures how participants in a training program react to the training. It attempts to answer questions regarding the participants' perceptions—Did they like it? Was the material relevant to their work?

(2) *Learning*—Assessing at this level moves the evaluation beyond learner satisfaction and attempts to assess the extent to which students have advanced in job-specific skills, knowledge, or attitude. Tests may be used here.

(3) *Transfer*—This level measures the transfer of technique and information to the learner's behavior due to the training program. Evaluating at this level attempts to answer the question—Are the newly acquired skills, knowledge, or attitude being used in the everyday environment of the learner?

(4) *Results*—Frequently thought of as the bottom line, this level measures the success of the program in terms more financially tangible and related to the business goals of the enterprise—increased production, improved quality, decreased costs, reduced frequency of accidents, increased sales, and even higher profits or return on investment.

"We feel that the training courses developed through our partnership with Virginia Commonwealth University are some of the best examples of training that we have seen here at Capital One and have successfully achieved all four of these levels.

"As discussed in the paper, the associate reactions to these training courses have been very positive, with 94% of the respondents to the post-training surveys saying they would recommend the training to other analysts at Capital One and 100% recommending the instructors. Thus, this training was a success at level one. Learning is more difficult to assess, and a conscious decision was made not to test at the end of each course. However, the associates in the course regularly developed rather complex models and analyses during the learning lab section of the training, demonstrating a high level or application of the course material to business and job-specific goals, indicating strong performance on level two.

"To assess Level 3, a business-impact survey was sent out a few months after the training to determine the degree to which the material had impacted analyses on the job. This survey revealed that 85% of the course attendees were either using the techniques or saw themselves using them in the near term, while 90% saw themselves using the techniques at some point in the future, satisfying level three.

"It is the evaluation at Level 4 that set this training apart. With a return on investment of more than $2 million annually from one follow-on project enabled by skills gained in these courses, *this training represented some of the best documented value of any training at Capital One and demonstrated the value of customized training in the corporate setting*. The documented return on investment significantly contributed to our application for *Training* magazine's Best Training Company in America in a year that saw us climb from 20th overall to 12th and placing us first among financial-services companies.

"We look forward to continuing this partnership with Virginia Commonwealth University and are pleased to support the publication of this successful venture and fruitful partnership."

Reference

Kirkpatrick, D. L. 1994. *Evaluating Training Programs: The Four Levels*. Berrett-Kochler, San Francisco, CA.

Roy Lowrance, Chief Technology Officer/Vice President, Capital One Financial Corporation, Information Technology, 1680 Capital One Drive, McLean, VA 22102, writes: "Thank you for the opportunity to relate the benefits that Capital One Financial Corporation has enjoyed by partnering with Virginia Commonwealth University to develop this

specialized training. As the Chief Technology Officer for Capital One, I am responsible for recommending new technology to the enterprise. The technology recommended reflects the company's needs and strengths. Capital One has built its success on using an *information-based strategy* for making decisions using data. This testing rigor occurs at all levels and business functions in the company. The decision-making processes that are core to our enterprise require constant infusions of new technology and ideas. Successful partnerships with universities, such as the partnership with Virginia Commonwealth University, not only offer a stream of technology and ideas, but also an important opportunity to provide specialized training that supports our ongoing efforts to research, identify, enable, and accelerate into Capital One novel technologies and techniques that improve and maintain Capital One's advantage in using data to make decisions. The content of the courses developed in this partnership has been identified by our *New Decision Paradigms* research team as valuable to our current areas of research concentration and has enabled successful work for the enterprise, with identified and confirmed financial value.

"For instance, we are developing simulation-based decision tools that enable improved management of complex systems and improved resource allocation. Through optimization based modeling, we are identifying cost savings and delivering "best" decision solutions by aligning business objectives directly to available and limited resources within multiple constraints. Making bold decisions in this data intensive context requires mastery of advanced modeling techniques. The specialized training developed through this partnership significantly improves the speed at which we, as an organization, can develop skills in a new technology or technique. It supports our business goals and exceeds our desires in terms of quality and value.

"Capital One continues to seek out and nurture such symbiotically beneficial partnerships with universities, and we expect that our partnership with VCU will deliver further documented return on investment and lead to future success."