Interfaces
Vol. 40, No. 1, January–February 2010, pp. 33–46
ISSN 0092-2102 | EISSN 1526-551X | 10 | 4001 | 0033

informs®
DOI 10.1287/inte.1090.0468

THE FRANZ EDELMAN AWARD
Achievement in Operations Research

Operations Research Improves Sales Force Productivity at IBM

Rick Lawrence, Claudia Perlich, Saharon Rosset, Ildar Khabibrakhmanov, Shilpa Mahatma, Sholom Weiss
IBM Corporation, IBM T. J. Watson Research Center, Yorktown Heights, New York 10598
{ricklawr@us.ibm.com, perlich@us.ibm.com, srosset@us.ibm.com, ildar@us.ibm.com, mahatma@us.ibm.com, sholom@us.ibm.com}

Matt Callahan, Matt Collins, Alexey Ershov, Shiva Kumar
IBM Corporation, Business Performance Services, Armonk, New York 10504
{callmatt@us.ibm.com, mattcoll@us.ibm.com, ershov@us.ibm.com, shivak@us.ibm.com}

In 2004, IBM introduced a set of broad operations research-based initiatives designed to improve the efficiency and productivity of its global sales force. The first solution, OnTARGET, provides a set of analytical models designed to identify new sales opportunities at existing IBM accounts and at noncustomer companies. The second solution, the Market Alignment Program (MAP), optimally allocates sales resources based on field-validated analytical estimates of future revenue opportunities in operational market segments. IBM Research developed the operations research models and initial internal websites for both solutions. The IBM Software Group initially implemented OnTARGET, which was subsequently made available to over 13,000 sales representatives across IBM sales organizations worldwide. The IBM Sales and Distribution organization deployed MAP as an integral part of its sales model to better align sales resources with the best market opportunities. We describe the development of both analytical models, and the underlying data models and websites used to deliver the solutions. We conclude with a discussion of the business impact, which we estimate as hundreds of millions of dollars annually for the combined initiatives.

Key words: customer relationship management; customer targeting; buyer behavior; data mining; wallet estimation.

IBM is a multinational computer technology, software, and services company, with approximately 380,000 employees worldwide and 2008 revenue of $103.6 billion. Over 50 percent of IBM's revenue comes from services, including strategic outsourcing, business-transformation outsourcing, business consulting, systems integration, and application management. The IBM Software Group (SWG) develops, markets, and sells software, including Information Management (database and content management), Lotus (collaboration tools), Rational (application development), Tivoli (security management), and WebSphere (Web applications). The IBM Systems and Technology Group (STG) manufactures computer servers under the product brands of System x, System i, System p, and System z, and also develops storage products. The IBM Sales and Distribution (S&D) organization is the primary client-facing sales group; it provides coordination of sales resources for IBM client accounts and sales territories. SWG and STG also include salespeople with expertise in specific product brands. At the end of 2007, IBM had approximately 40,000 employees in sales-related roles.

Improving the productivity of such a large sales force can be an effective operational strategy to drive revenue growth and manage bottom-line expenses in today's challenging economic climate. At one level, driving growth requires that sales professionals be provided with leading-edge tools to identify better leads and hence close more deals. At a higher level, given that highly productive salespeople are a constrained resource, sales executives must optimally deploy the available sales force in the best revenue-generating accounts. Both objectives benefit from developing and applying leading-edge operations research to enhance decision-making capabilities at key points in the sales organization and the sales process.

In 2004, IBM initiated a broad, analytics-based initiative to improve sales productivity at both levels. The first solution, OnTARGET, was developed in response to a request from SWG for a novel, analytics-based approach to help its sales professionals identify companies, which we refer to as *whitespace*, that had not purchased previously from IBM. (IBM customers are companies, not individuals. Sales accounts typically include multiple locations within a specific company; sales territories contain multiple sales accounts.) This objective was quickly extended to include existing IBM customers, and subsequently to also include models for IBM server offerings. In 2005, the second solution, the Market Alignment Program (MAP), was initiated with S&D to develop a quantitative approach to guide the deployment of IBM sales professionals to customer accounts with higher expected future revenue. In contrast to a conventional approach in which salespeople are largely allocated to accounts with the highest recent revenue, the MAP process recognized the need to develop analytical models to estimate the realistic revenue opportunity at each company, i.e., the amount of a specific product group that IBM could realistically hope to sell to a specific customer.

Hiring the best sales representatives is an obvious first step in optimizing sales productivity; however, we increasingly recognize (Ledingham et al. 2006) that realizing the true potential of any sales force requires that sales representatives (reps) and executives have relevant information technology (IT)-based tools and solutions. During the past decade, we have seen the development of a number of customer relationship management (CRM) systems (Zikmund et al. 2002, Berry and Linoff 2004) that provide integration and management of data relevant to the complete marketing and sales process. Sales force automation systems (Morgan and Inks 2001) enable sales executives to better balance sales resources against identified sales opportunities. We generally, but not uniformly (Speier and Venkatesh 2002), accept that such tools improve the overall efficiency of the sales process; however, major advances in sales force productivity require access to both relevant data and informative, predictive analytics derived from these data.

Figure 1 clarifies the roles of various IBM groups in the development, deployment, and use of OnTARGET and MAP. The Business Analytics and Mathematical Science organization within IBM Research, working closely with the Business Performance Services (BPS) group, developed the analytical models for both solutions. BPS played an integral role in defining the underlying business objectives of the initiatives, and worked closely with the other IBM organizations to ensure integration within their respective business environments. The sales teams that sell software and server products use OnTARGET; the sales management teams use MAP to align sales resources within S&D, SWG, STG, and Services.

Overview of the Solutions

One of the biggest challenges in any operations research (OR)-based project is mapping the set of high-level business objectives into specific tasks. OnTARGET and MAP share a common set of four required tasks (Figure 2).

Data Model

IBM maintains an internal view of its customers, including past transactions, assignments to sales territories, and other account-specific information. Past transactional data are an essential input to both solution models; we also require that these internal data be used with the respective firmographic data, such as company industry, estimated annual sales, number of employees, and organizational structure. Matching each company in the IBM internal database view to its equivalent entity provided by an external data provider (e.g., Dun & Bradstreet (D&B)) is very challenging, particularly when we include millions of companies worldwide, as is the case for our solutions. A challenge for larger customer organizations is to ensure that the organizational structure reflected in the IBM internal representation is consistent with the external view. For example, we must be certain that if a corporation has 10 subsidiaries in our internal transactional view, the corporate firmographic information correctly reflects this organizational view.

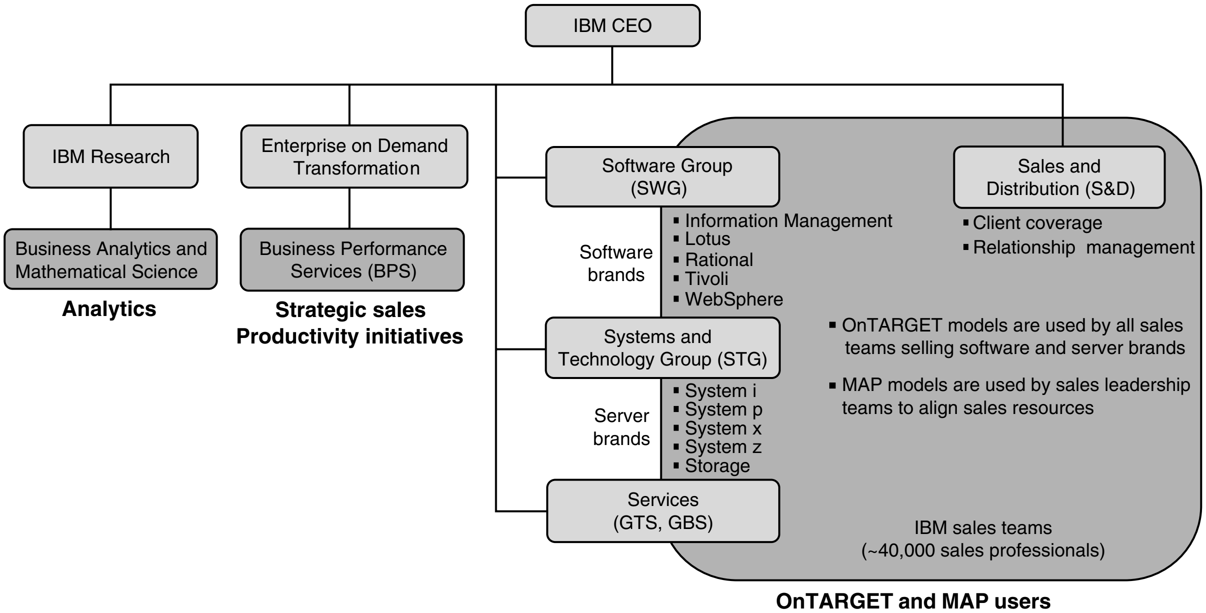

Figure 1: The organization chart shows a simplified, high-level view of IBM's organizational structure.

Although not immediately part of the specific OR methods, designing a flexible data model is an essential component of our overall solutions and required significant effort. In addition to supporting the data representation ultimately required for input to the predictive models, the data model must support frequent (i.e., quarterly) updates to the source data. Users of our Web-based tools also want to be able to view relevant summaries of the underlying transactional history and firmographic information in conjunction with the results of our predictive models. Hence, our data-model Web deployment must also support easily constructed data summaries.

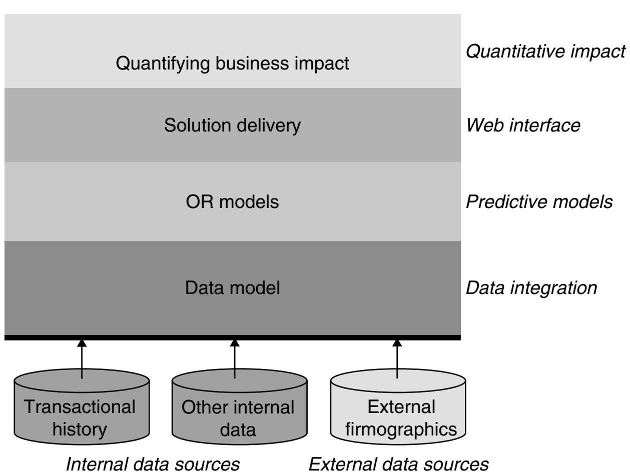

Figure 2: Four components are inherent in both OR solutions.

OR Models

Taking a high-level statement of a set of business objectives and refining them into a well-defined analytical problem that optimization approaches or predictive-modeling algorithms can address is a major challenge in the modeling process. Another important, intermediate step is processing the massive amount of available data to yield a relatively small set of explanatory features that are used as direct input to the predictive algorithm. We discuss these issues further in the respective OnTARGET and MAP sections.

Solution Delivery

Ultimately, the impact and value of any OR initiative depend on the extent to which decision makers have access to insights generated by the technical solution. The only practical means to distribute this information is via well-designed Web-based tools that facilitate easy distribution of results, seamless updates, and a hyperlinked report structure that supports compact views of both the model results and the relevant underlying data.

Quantifying Business Impact

Large companies like IBM have complex sales processes. Although this complexity provides many opportunities to apply leading-edge OR methods to improve productivity, the magnitude of the sales operation suggests that isolating and quantifying the impact of new process-improvement initiatives can be difficult. From a statistical perspective, it is not sufficient to merely observe that a key performance indicator (e.g., revenue) has improved after the introduction of a technological enhancement. Any number of other factors (e.g., the external business climate) could contribute to the observed performance, and rejecting other hypotheses that might account for this behavior is typically not possible. Given this observation, an important component of any OR initiative is a mechanism to allow a quantitative assessment of the impact of injecting OR technology into a complex business process.

The subsequent discussion of the business impact of our sales initiatives requires a high-level understanding of the IBM sales-opportunity tracking process. When an IBM salesperson identifies a potential sales opportunity, the specifics of the proposed deal are entered into a sales-opportunity database, in which all open opportunities are tracked as they progress through a set of stages, ideally culminating in a deal that generates revenue for IBM. This process, the sales pipeline, allows sales executives to view snapshots of the current dollar amount of the sales opportunities, an important indicator of future revenue. One indication of the quality of the sales pipeline is the yield, which is the fraction of sales opportunities that ultimately closes as revenue. Improvements in the opportunity-identification process can result in increased yield; even small improvements in the yield can significantly impact revenue. In later sections, we discuss the financial impact of both sales initiatives in this context.

OnTARGET: Estimating Propensity to Buy

When IBM started the OnTARGET effort in 2004, we spoke with a number of IBM sales professionals to define specific requirements for a tool that would allow them to identify new software sales opportunities. Many felt that they were being forced to use too many different tools to do their jobs and had no rigorous, analytical capability to help them prioritize which companies to pursue. As noted above, the initial focus was on targeting whitespace customers, but our initial modeling results quickly broadened the scope to include identifying new software sales opportunities at existing IBM accounts. The project formally started in January 2004; by July 2004, IBM Research had deployed a prototype solution to 35 selected sale professionals in Canada. The plan was to run the Canada pilot through the end of 2004, but positive feedback from the participants led to the accelerated deployment of a US version in January 2005. The formal development and Web hosting were transferred to an IT support group within STG in late 2005. IBM Research continues to update the predictive models quarterly.

OnTARGET Data Model

Because OnTARGET includes models for whitespace customers, its data model is organized using external company definitions rather than IBM's internal customer view. It uses D&B's DUNS information because it is available in many countries worldwide and provides a useful global organization hierarchy that specifies relationships between individual company sites, subsidiaries, and corporate headquarters. Each transaction that IBM executes with a customer is tagged with an IBM customer number that OnTARGET matches to a DUNS number to link the transactional view to the customer's D&B firmographic data.

OnTARGET Predictive Models

The goal of the OnTARGET propensity models is to differentiate customers (or potential customers) by their likelihood of purchasing various IBM products. Rather than model at the individual product level, we build our models to predict purchases within broad product groups or brands (Figure 1), such as Lotus and System p. Currently, we develop separate propensity models for each SWG and STG product brand shown in Figure 1. The objective is to build propensity models that are able to differentiate the high-propensity and low-propensity customers on a brand-by-brand basis. For each product brand Y, we first divide the universe of OnTARGET companies into three distinct groups:

1. Companies that have previously purchased Y. We eliminate these companies from the propensity modeling because our objective is to predict first-time purchases.

2. Companies that have a transactional relationship with IBM but have never purchased Y. For these companies, we can utilize both IBM historical sales data and D&B information to build our existing-customer model.

3. Companies that have never purchased from IBM. For these companies, we only have the D&B information; we term this model the whitespace model.

We begin by specifying a geography, a brand Y, and a modeling problem (i.e., existing customer or whitespace). Next, we identify positive and negative examples to be used for modeling. In each modeling problem, we are trying to understand what drives the first purchase decision for brand Y, and delineate companies by their likelihood of purchase. Assuming that the current period (typically last year or the previous last two years) is t, we formulate our modeling problem as follows:

Differentiate companies that bought brand Y during period t (but not before period t) and companies that have never bought brand Y.

Of the companies that never bought brand Y before period t, some will have bought other IBM products before t. These companies form the basis of the existing-customer model for Y. The companies that never bought any IBM brand before t are the basis for the whitespace model. Thus, for the whitespace problem, our positive and negative examples are as follows:

Positive: Companies that have never bought from IBM before t, but then bought Y during t.

Negative: Companies that have never bought from IBM before or during t.

The definitions for the existing-customer problem are similar, except that a previous purchase from IBM is required for inclusion. For some combinations of the geography, brand, and modeling problem, the number of positives might be too small for effective modeling (we typically require at least 50 positive examples to obtain good models). In that case, we often combine several similar modeling tasks (where similarity can be based on geography, brand, or both) into one metamodel with additional positives. Rosset and Lawrence (2006) discuss the trade-offs involved in this approach and demonstrate its effectiveness.

Next, we define a set of variables or explanatory features to be used in modeling. The objective is to select features that are likely to provide some differentiation (signal) between positive and negative examples. For existing customers, we derive multiple features from historical IBM transactions, describing the history of their IBM relationship before period t. These features summarize past transactional information across all brands, not only the brand being modeled, because, for example, server purchases may influence future software purchases. Examples of these features are as follows:

- Total amount spent on software purchases in the two years before t;
- Total amount spent on software purchases in the two years before t, compared to other IBM customers (i.e., rank within IBM customer population);
- Total amount spent on storage-product purchases in the four years before t.

For both existing and whitespace customers, we derive features from the D&B data:

- Company-size indicators (e.g., revenue, employees) in both absolute and relative terms (i.e., rank within industry);
- Industry variables—raw industry classification from D&B and derived sector variables; and
- Company's location in the corporate hierarchy (i.e., corporate headquarters, subsidiary, etc.).

We then build a classification model (more accurately, a probability-estimation model) that uses these explanatory features to differentiate the positive and negative examples. For each example, the model estimates the probability of belonging to the positive class. For simplicity, and to avoid potential overfitting, we implemented the classification model using logistic regression. For presentation in the OnTARGET tool, these continuous scores are binned from 1 to 5, with bin distributions specified such that only 15 percent of existing-customer examples receive the highest rating of 5. For the whitespace model, only

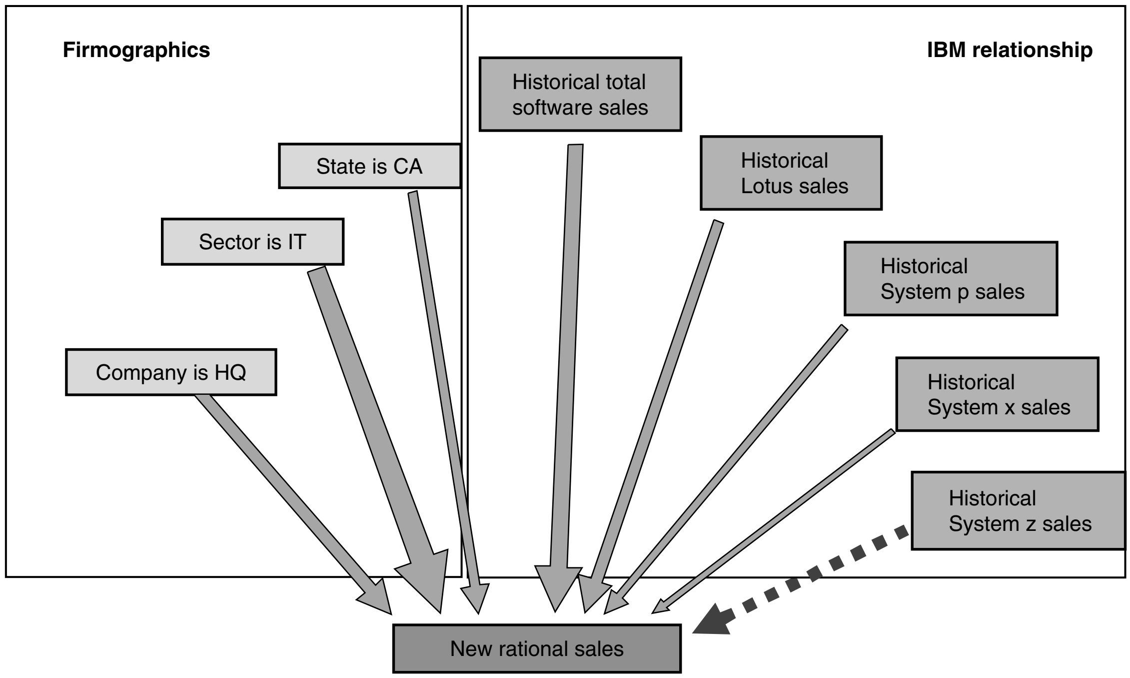

Figure 3: The diagram illustrates predictive relationships between some derived variables and new rational software sales to existing customers.

5 percent receive a rating of 5, reflecting the observation that selling into a noncustomer account is generally more difficult.

An Example of an OnTARGET Model

We can examine the resulting models and, to some extent, interpret them as scorecards that describe the effect of different variables on the likelihood of converting a company into a customer for brand Y. Figure 3 shows the predictive relationships in a customer model for the Rational software brand. The solid arrows signal a positive effect (i.e., higher values of this variable correspond to increased propensity); dashed arrows signal an adverse effect on propensity. The arrow width indicates the strength of the effect, as measured by the magnitude of the regression effect. We show only statistically significant (measured by *p*-value) effects in the figure and see several interesting effects:

- Industrial sector (IT), geography (California), and company corporate status (i.e., headquarters) seem to have a strong predictive effect. This seems consistent with Rational being an advanced software-development platform that medium-sized IT companies in California (likely to be high-tech industry leaders) might be interested in purchasing.
- The size of total prior IBM software purchases (relationship) seems to be a strong indicator of propensity to buy, as does strong Lotus usage.
- Although the total size of prior IBM non-software purchases does not have a strong effect, some specific non-software brands seem to be important. System p (and System x to a lesser extent) usage seems to encourage Rational sales, and System z usage seems to discourage them. We might explain the final observation by the particular nature of the software relationship with System z customers, who often manage their software relationship with IBM in conjunction with their System z relationship.

Such scorecards are interesting because they identify expected relationships (and hence build confidence with end users) and uncover less obvious relationships that can provide additional marketing insight.

OnTARGET Model Validation

Using 10-fold cross validation (Hastie et al. 2009) is a conventional, statistical approach to validating the accuracy of a classification model. The objective of this validation is to measure the model's effectiveness in placing good sales prospects at the top of a list sorted by the classification-model output. Such models are often compared to a random ordering of the prospects and (of course) shown to be better because it is difficult to lose to a random model. Salespeople usually perform better than a random model. Given no other information, a simple model that salespeople might follow would be to prioritize existing customers based on past revenue and to prioritize whitespace companies based simply on company size as measured by annual sales. It is straightforward to capture such a strategy in a model simply by sorting existing customers by the size of previous revenue and whitespace companies by annual sales. We refer to such models as Willy Sutton models, after the famed bank robber who robbed banks because "that's where the money is." Extensive cross validation of the OnTARGET models has demonstrated that our models generally outperform Willy Sutton models, often by large margins for software products such as Rational, when propensity to buy depends on many factors, not only company size. Lawrence et al. (2007) provide details on the cross-validation approach.

A more interesting evaluation is to judge the models by their actual success in predicting new sales. We have been able to do this by considering new sales recorded in 4Q 2006 and investigating the scores that our previously built models in 3Q 2006 assigned to these sales. These sales were not visible in the data at the time we built the models; however, they were most likely initiated before the models' results were available and thus not affected by these results. Hence, we are getting a clean evaluation of the models' success in identifying actual sales as high-propensity opportunities. Because of space limitations in this paper, we refer to Lawrence et al. (2007) for details of this analysis, in which we demonstrate that our models again generally outperform Willy Sutton models, with lifts (essentially the ratio of model performance to random) ranging from 5 to 12 for the 10 product brands examined.

OnTARGET Model Automation

OnTARGET is currently used in 22 countries. We do not build separate models for each country, but we instead build models for eight distinct geographic areas formed by aggregating smaller countries. For each geographic area, we build existing-customer and whitespace models for each of the 10 software and server brands, resulting in 160 new models in each quarter. It is therefore critical to have a reliable, repeatable, and automated modeling methodology. Initially, we did all OnTARGET modeling manually; however, as we became confident in our methodology, we gradually moved to the fully automated process currently implemented. The main characteristic of this process is having a large collection of possible predictive variables and selecting some part of it for each prediction model, based on statistical considerations and our past experience regarding the importance of different variables in different models. Currently, the primary manual intervention in this process is examining the final output and evaluating the models, thereby ensuring that data changes, bugs, or other unexpected phenomena have not adversely affected the predictive performance.

OnTARGET Website

We provide the propensity models described above, plus the underlying transactional and firmographic data, to IBM sales professionals via an IBM internal website. As of December 2008, the OnTARGET website included information for more than two million companies worldwide. Given this large number of companies, a key requirement was to enable a user to quickly locate a subset of companies that meet specific requirements, such as size (as measured by employees or annual sales), industry, and location. Using the Web-based tool, building an initial list of current IBM customers within a specific IBM sales territory, or finding a set of whitespace companies within a given state or province, is straightforward. Within either subset, a user can then build a targeting list consisting only of the companies with a high propensity to buy a specific brand, based on the output of the propensity models. Clicking on any company in the list loads a company-detail page that shows firmographic information and historical transactional summaries, plus the outputs of the propensity models for the software and server brands mentioned above. We worked closely with IBM salespeople to design a flexible Web interface because the impact of an OR solution ultimately depends on delivering actionable insight to end-user decision makers.

OnTARGET Business Impact

Tool Adoption

IBM sales professionals have repeatedly stated that they are not interested in using tools that do not help improve their overall productivity and efficiency. Because salespeople are not required to use OnTARGET, its use indicates its perceived value. In October 2005, there were 1,381 users in the United States. By September 1, 2007, this number had grown to 9,862 users across 22 countries in all major IBM geographies. As of December 2008, 13,784 people within IBM had access to the tool. In 2008, there were approximately 60,000 logins to the website by 6,454 unique users. (Only about half of the approximately 40,000 IBM salespeople are in roles that would benefit from the use of OnTARGET. Many of those granted access since 2005 have moved to other positions that no longer require access to the website.) OnTARGET users downloaded over 235,000 company-detail reports in 2006, the last year for which we have these data.

Productivity Gains

The primary OnTARGET users are face-to-face and call-center sales personnel looking for the best potential opportunities in their territories. An OnTARGET user-base survey identified the average productivity gain as two hours per week, which we attribute to the user being able to quickly create focused targeting lists and not having to use multiple tools to access additional data and research prospective customers. Clearly, this improves overall productivity because sales reps can spend more time focusing on the sales process. We can estimate an actual dollar savings using an average sales rep burdened cost. If the 6,454 users in 2008 each saved two hours per week over 50 weeks per year, the resulting savings attributable to this enhanced productivity tool is easily over $10 million annually.

Revenue Impact

Using OnTARGET, a sales rep accesses all information for a specific company from a single page. These data include the scores produced by our predictive models, previous sales to this company, and other useful information. We log each page that a user touches; hence, we know if that user has accessed information for a given company. When a new sales lead is entered in the sales-opportunity database, we check whether this company has been touched in OnTARGET during a 90-day window prior to the entry date. If it has been touched, we mark this lead as having been "OnTARGET influenced." We also track the total number of opportunities that are ultimately marked as "won" (i.e., produced revenue for IBM).

Using this logging capability, we classify opportunities entering the sales-opportunity database into two disjoint populations: those influenced by OnTARGET (i.e., touched within the preceding 90 days) and those not touched within this window. We can then examine the dollar amounts associated with these two populations of sales opportunities for both software and servers in North America and Europe. For each population of opportunities, we compute the yield as the fraction of the total dollar amount entered in the opportunity database that is subsequently marked as won. Taken over this broad set of opportunities, the yield for the OnTARGET-influenced opportunities is 1.6 percent higher than the non-OnTARGET yield. The higher close rate suggests that using OnTARGET results in higher-quality sales leads.

We can assess the impact of OnTARGET by looking at what would have happened had OnTARGET not been available. We postulate that the opportunities touched by OnTARGET would have closed at the 1.6 percent-lower yield rate. Although a 1.6 percent difference appears modest, it translates to a dollar-revenue impact of $468 million in 2008 for North America and Europe. We should note that the OnTARGET yield could be higher for other reasons; e.g., those who choose to use OnTARGET are inherently better sellers. Although we do not have the necessary information to eliminate such hypotheses, our discussions with the sales teams suggest that the 2008 OnTARGET revenue impact is several hundreds of millions of dollars.

MAP: Estimating Realistic Customer Opportunity

A major challenge that IBM and all sales-oriented companies face is aligning sales resources with the best revenue-generating opportunities. Sales resources are often assigned to customer accounts based on the distribution of recent revenue, thereby missing significant revenue-growth opportunities. A better approach is to make allocation decisions based on unbiased, objective estimates of the forward-looking revenue opportunity at each account. In this context, revenue opportunity is not simply an estimate of the revenue that might be generated at an account in the near future using the existing sales presence; more specifically, it is a measure of the potential revenue (i.e., wallet) based on what we could realistically expect to sell to this customer if we were to intensify our sales efforts. As we discuss later in the *MAP OR Models* section, this is an important distinction, and estimating this forward-looking opportunity is a challenging data-analytics problem that requires an OR-based solution for two reasons. First, analytics can provide a fact-based, objective estimate and hence avoid inevitable human bias. Second, these estimates are required for 25 product brands (OnTARGET addresses the 10 software and server brands shown in Figure 1; MAP expands this number to 25 brands by including a number of services offerings) and hundreds of thousands of IBM sales accounts; only computational models can provide this scalability.

MAP Business Process

In 2005, IBM initiated MAP to build a rigorous, fact-based process to improve the alignment of IBM sales resources with the best market opportunities. At a high level, the MAP process consists of three steps:

Step 1. Development of an analytical model to estimate the realistic revenue opportunity for each major IBM product brand at each existing IBM account.

Step 2. Validation of the model opportunity estimates via comprehensive workshops conducted with hundreds of sales teams covering a particular sales territory or product brand.

Step 3. Reallocation of sales resources based on analysis of imbalances between current resource assignments and validated revenue opportunity within IBM coverage units and brands.

The opportunity estimates from the OR models are available on an IBM internal website, which is used during the workshops (Step 2 above). One workshop objective is to reach agreement on reasonable revenue-opportunity numbers for the IBM accounts that the sales team covers, using the analytical estimates as an objective starting point. The sales team can either accept the model results or provide modified numbers and reasons for modifying them. The tool captures these validated opportunity numbers and stores them in the MAP database for subsequent processing, as we describe below.

MAP Data Model

The MAP data model is similar to the OnTARGET model. It joins each company's firmographic data with that company's history of past IBM transactions. Unlike OnTARGET, which includes whitespace companies, MAP focuses exclusively on existing IBM accounts. Therefore, its data model retains an IBM internal view of companies rather than the D&B view that OnTARGET uses. The details of the data model are secondary in this discussion, but we should note that a major challenge was to ensure that the view (e.g., past annual IBM revenue) of the account as seen in the Web-based tool is completely consistent with the view of the IBM sales team covering this account.

MAP OR Models

For the MAP workshops, we need an unbiased, realistic estimate of the true revenue opportunity at each account to drive an informed discussion with each sales team. We discuss customer opportunity in the context of IBM as a seller of IT products to a large collection of customers for whom we wish to estimate the opportunity. We considered three nested definitions:

1. The total spending by this customer in a particular group of IT products or services, taken over all IT providers—i.e., the customer's total IT spending (by product group), which we denote as *TOTAL* opportunity.

2. The total attainable (or served) opportunity for the customer. This corresponds to the total expenditure by the customer in IT areas that IBM's products and services cover. Although IBM serves nearly all areas of IT spending (software, hardware, and services), its products do not necessarily cover all needs of companies in each area. Thus, *SERVED* opportunity is smaller than total IT spending.

3. The "realistically attainable" opportunity, as defined by what the "best" customers spend with IBM. This differs from *SERVED* opportunity because it is unrealistic to expect customers to spend their entire budget with IBM. We refer to this as *REALISTIC* opportunity.

Sources such as D&B provide the annual company revenue (*COMPANY_REVENUE*) for all companies. We also know the mean annual historical sales (*IBM_SALES*) of IBM's customers. In principle, we expect the following relationship to hold for each company:

$$IBM\_SALES < REALISTIC < SERVED$$
$$< TOTAL < COMPANY\_REVENUE.$$

Note that we expect *IBM_SALES* to approach the *REALISTIC* opportunity for companies in which IBM is the dominant IT provider. Therefore, the notion of realistic is defined relative to our best customers, the ones who spend as much as we could hope that they spend with IBM. The beauty of this definition is that it implies an analytical framework that allows us to estimate and evaluate the *REALISTIC* opportunity based on the observable IBM revenue, although the opportunity under the other definitions remains a largely unknown and unobservable quantity. Hence, because of its alignment with the business objective and its analytical elegance, we use *REALISTIC* opportunity as the MAP revenue-opportunity definition.

To make this notion of opportunity more concrete, let us consider a customer X and imagine that we have (say) 1,000 identical customers exactly like X, with each customer independently making its decision about how much to spend with IBM. We could take the 95th percentile of the spending distribution (i.e., the quantity Q such that 95 percent of these 1,000 identical customers spend \$$Q$ or less with IBM) as our *REALISTIC* opportunity estimate for X and its 1,000 replicas. In practice, we do not, of course, observe multiple copies of each company; hence, our challenge is to use the available data to estimate this spending percentile for each company.

In general, the approaches for estimating the spending percentile fall into two categories:

• *Local* approaches, which start with the idea we describe above (of having 1,000 copies of X), approximate it by finding companies that are "similar" to X, and estimate X's opportunity as the 95th percentile of the IBM sales of this "neighborhood."

• *Global* optimization models, which attempt to describe the 95th percentile as a function of all the information we have about our customers. The most commonly used approach is quantile regression, which directly models the quantile (or percentile) of a response variable Y as a function of explanatory features (or predictors) X. In our case, the response Y is the IBM spending, and X includes both firmographic variables from D&B and features extracted from IBM historical transaction data. Unlike conventional regression techniques, which minimize the squared loss, quantile regression minimizes the quantile loss, which is an asymmetrically weighted absolute loss function.

We explored several existing approaches for quantile estimation, such as linear quantile regression (Koenker 2005). We also developed novel modeling techniques, including a k-nearest-neighbor approach that directly follows the definition of identifying similar customers, quantile regression trees, and others (Breiman et al. 1984, Langford et al. 2006, Merugu et al. 2006, Perlich et al. 2006). We evaluated the different models against the expert feedback that was collected in the initial round of MAP sales team workshops in 2005; based on this analysis, we selected linear quantile regression for subsequent MAP modeling. We estimate the currently deployed models using historical revenue and firmographics data as independent variables (similar to the OnTARGET model in Figure 2) and the recent IBM revenue in a given brand as dependent variables (response), subject to minimizing quantile loss as noted above. The appendix provides more detail on the quantile-regression formulation.

In 2008, we added one refinement to the above model after observing a dominating effect of the historical IBM revenue in the revenue-opportunity estimation. The result was that some large companies, with very modest recent IBM revenue, were assigned opportunities that were significantly lower than would otherwise be expected. To remedy this, we extended the model process and built two independent models, one including and one excluding past IBM revenue. The latter assigns larger opportunities to large customers independently of the amount of previously generated IBM revenue, provided they share other important characteristics of our best customers, particularly size. Our final opportunity estimate is the maximum prediction from both models.

One essential but unknown parameter in this algorithm is the actual quantile. We initially suggested 95 percent as a potential value. The choice of a higher quantile implies higher estimates of the revenue opportunity relative to existing revenue, which is particularly appropriate for emerging counties with rapid growth. Over the past three years, we have refined this parameter to reflect worldwide growth projections from external sources and internal market estimates. Working with the worldwide IBM marketing teams, we reached agreement on the ratio of future revenue opportunity to current revenue for each brand and each country and set the quantile to preserve this ratio. Growth markets may have opportunity-to-revenue ratios of up to 4.5; established markets will have ratios closer to 2. We find the quantiles required to preserve these ratios via a simple search, yielding quantiles between 75 percent and 99 percent.

MAP Business Impact

Large companies like IBM have complex sales processes and may concurrently run multiple initiatives to improve performance in the changing marketplace. This complexity and the magnitude of the sales operation imply that isolating and quantifying the impact of one particular process-improvement initiative can be difficult. The MAP workshop process was deployed initially in 2005 and, based on the validated opportunity obtained in this process, a small number of US sales resources were realigned in 2006. This initial deployment involved a significant focus on implementing mechanisms that ultimately allow us to assess quantitatively the impact of the injected OR technology. Once the overall benefits were demonstrated during 2005–2006, the MAP process was integrated within the broader IBM sales model, and the focus shifted toward execution of the broad alignment initiative. As a result, since 2006, we have fewer quantitative (bottom-up) measurements of impact but a larger, top-down view of the overall business impact.

Given the complexity and the changes over the years, we will discuss the impact of MAP along different metrics for the periods during which they were available: sales pipeline growth, quota attainment, and revenue growth and impact. However, before we present the analysis and results of these metrics, we must provide more detail on the MAP process to clarify important notions and measurements.

MAP Sales Force Alignment Process

During the MAP workshops, the sales teams review the analytics-based opportunity predictions and have the opportunity to overwrite them based on their knowledge of the market and competitive environment. During the initial MAP deployment in 2005, we observed that about 45 percent of the predicted opportunities were accepted directly, 17 percent were increased, 23 percent were decreased, and 15 percent were set to zero. These validated opportunities are the main input for the account segmentation step described below.

The key output of the MAP workshops is the client segmentation, which is used as input to subsequent resource decisions (Figure 4). Prior to MAP, resource allocations were based on the one-dimensional view of prior-year revenue on the horizontal axis. Note that accounts with higher future opportunity are labeled as *Invest* accounts and hence are eligible to receive increased sales resources. Conversely, accounts with relatively low future opportunity (*Core optimize*) are expected to remain flat or decrease in terms of resources. Accounts in the *Core growth* segment are examined individually for potential resource reallocations. The final segment, which we refer to as *Opportunistic*, covers smaller accounts that may be more appropriately covered via lower-cost channels, such as IBM business partners.

In principle, one could develop a formal optimization approach to specify resource movements such that future revenue is optimized. However, our experience to date has been that resource-allocation decisions are heavily constrained but also require strong and easily violated assumptions of how revenue will decrease once resources are withdrawn. Decision makers possess valuable information that is not easily formalized and would be difficult to incorporate into a model. Therefore, the most defensible and practical

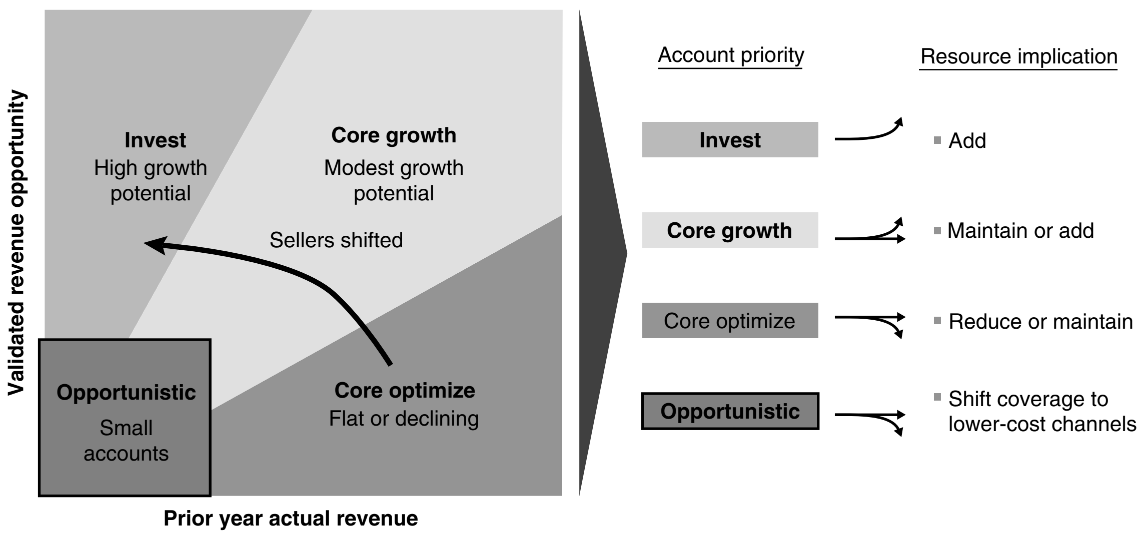

Figure 4: The diagram shows the two-dimensional segmentation that labels each account based on future opportunity relative to historical revenue performance.

approach is to provide human decision makers with the basic analytical insight (i.e., the account segmentation for each major geography) and let them make the final reallocation decisions.

MAP Adoption

Since MAP's initial deployment in the United States in 2005, IBM has rolled it out to 55 countries. During the 2006 deployment, approximately 420 MAP workshops, which involved more than 2,300 individual customer accounts across all IBM sales units, were conducted with sales teams globally. One measure of the MAP footprint is that these accounts represented 95 percent of the previous-year IBM revenue and approximately 55 percent of the total modeled revenue opportunity. As a result of the 2005 deployment, approximately 500 sales professionals were reassigned to higher-opportunity accounts in 2006, and the coverage of approximately 50 lower-opportunity accounts was shifted to lower-cost Web and call-center coverage models. During the 2008 deployment, approximately 600 MAP workshops, which involved IBM sales units covering more than 13,000 individual accounts were conducted with sales teams globally. Approximately 1,500 and 3,000 resource moves were made in 2007 and 2008, respectively, and we estimate that 3,000 more resources will have been reallocated in 2009 because of the MAP initiative. The magnitude of these resource shifts is evidence of the impact of the MAP initiative within IBM and the confidence of senior IBM executives in the viability of this analytics-based approach to drive sales productivity.

Sales Pipeline Growth

The sales pipeline, as derived through the sales-opportunity management system, is an important leading indicator of future revenue. Again, growth of the Invest account segment and the contribution of the Invest segment to the total pipeline are important indicators of MAP impact. It is also important to recognize that any impact on the sales pipeline resulting from MAP will occur over some period beyond the current quarter. We can therefore use a rolling four quarter's worth of validated pipeline as an appropriate measure. As of week 12 in 3Q 2006, the validated sales pipeline of Invest accounts (over a rolling four-quarter period) grew year-over-year at a rate of 14 percent greater than the total US sales pipeline. Because the sales pipeline is a leading indicator of revenue, the fact that pipeline growth is greater than revenue growth in the Invest accounts is further evidence of the financial impact of MAP during its early deployment.

Quota Attainment

An additional measure of impact is the performance of those sales resources that are either shifted or dedicated to Invest accounts as a result of MAP. Many sales professionals are assigned an annual revenue-based sales quota, and attainment of this quota measures individual sales performance. For the first two quarters in 2006, the year-to-date quota attainment of shifted resources because of MAP recommendations was 45 percent compared to 36 percent for resources shifted as a result of other initiatives. This 9 percent improvement in quota attainment suggests that MAP has identified accounts with greater sales opportunities and that moving resources to these accounts has yielded increased productivity.

Revenue Growth and Impact

Because resource deployment investments are expected to drive incremental revenue growth, measuring the MAP impact on revenue and growth is useful. However, we would not expect the impact of those investments to occur quickly; any shifted resource will need time to ramp up to full productivity. We can therefore assess the impact by comparing the year-over-year revenue growth across each MAP account segment (e.g., Invest). For the set of accounts in the Invest and Core growth segments, representing 2008 revenue of $53 billion, the year-over-year growth rates were 8.2 percentage points higher than in the two other account segments (Core optimize and Opportunistic), which delivered revenue of $9 billion but were deprioritized for investment. IBM senior management believes that some fraction of 8.2 percent-higher growth rate is attributable to MAP-initiated resource moves. A conservative position is that 1 percent of this growth is because of MAP. Under this assumption, we estimate the financial impact of MAP to be 1 percent of $53 billion, or approximately $500 million in 2008.

A further measure of the impact is that the MAP output influences the revenue and expense plans of all operational units. As part of their performance reviews, the sales units are expected to have executed the MAP process effectively, to be committed to shifting resource based on their MAP work, and to be able to prove through other human resources tracking that seller territories have been better aligned and optimized to support growth. IBM is anticipating significant additional benefits from MAP in the coming years.

Summary and Conclusions

OnTARGET and MAP are examples of operations research solutions that were designed to address specific business challenges in the broad area of sales force productivity. Although they address different underlying issues, these solutions implement a common approach that is generally applicable to a broad class of operational challenges. Both solutions rely on rigorously defined data models that integrate all relevant data into a common database. Choices of the data to be included in the data model are driven both by end-user requirements and the need for relevant inputs to analytical models. Both business problems have a natural mapping to applications of predictive modeling: OnTARGET predicts the probability to purchase, and MAP estimates the realistic revenue opportunity. Delivering the underlying data and the analytic insights directly to decision makers (sales representatives for OnTARGET and sales executives for MAP) is crucial to driving business impact, and IBM has invested significant effort in developing efficient Web-based tools with the necessary supporting infrastructure. Both solutions have been deployed across multiple geographic regions, with a strong focus on capturing and quantifying the business impact of the initiatives. Indeed, we have field evidence that the analytical models developed for OnTARGET are predictive. The impact of MAP is perhaps more difficult to quantify, but we have growing evidence to suggest that sales force allocations made within the MAP process are leading to measurable improvements in sales efficiency. As described above, we estimate the OnTARGET 2008 revenue impact to be in the several hundreds of millions of dollars and the MAP 2008 impact to be approximately $500 million. As of this writing, we anticipate a similar impact in 2009. Finally, although we have implemented these solutions within IBM, we believe that the underlying methodologies, business processes, and potential impact are relevant to enterprise sales organizations in other global industries.

Appendix. Linear Quantile Regression Formulation and Calculation

Assume we have n training observations $(\underline{x}_1, y_1), \ldots, (\underline{x}_n, y_n)$, with $\underline{x}_i \in \Re^p$ a vector of features or explanatory variables, and $y_i \in \Re$ a numeric response. We are looking to build a model to describe the τth quantile of $P(y \mid \underline{x})$ (denote it $f_\tau(\underline{x})$). The linear quantile regression approach assumes that we can approximate this by $f_\tau(\underline{x}) \approx \sum_{j=1}^{p} \beta_j x_j$.

We obtain the estimate of the parameters $\beta_1, \ldots, \beta_p$ by minimizing the quantile loss function L_τ, defined by

$$\hat{\underline{\beta}} = \arg\min_{\underline{\beta}} \sum_{i=1}^{n} L_\tau\left(y_i - \sum_{i=1}^{p} \beta_j x_{ij}\right),$$
$$L_\tau(r) = \begin{cases} \tau r & \text{if } r \geq 0, \\ -(1-\tau) r & \text{if } r < 0. \end{cases} \quad (1)$$

Koenker (2005) and Perlich et al. (2006) provide a discussion of why this is an appropriate modeling approach for conditional quantiles.

The optimization problem implied by the formulation in Equation (1) is a linear program (LP) in the parameters $\beta_1, \ldots, \beta_p$. The simplest formulation of this LP is obtained by adding $2n$ nonnegative dummy variables for the residuals

$$\begin{aligned} \min_{\underline{\beta}} \quad & \sum_{i=1}^{n} \tau r_i^+ - (1-\tau) r_i^- \\ \text{s.t} \quad & r_i^+ - r_i^- = \left(y_i - \sum_{i=1}^{p} \beta_j x_{ij}\right) \quad \forall\, i = 1, \ldots, n, \\ & r_i^+ \geq 0 \quad \forall\, i = 1, \ldots, n, \\ & r_i^- \geq 0 \quad \forall\, i = 1, \ldots, n. \end{aligned}$$

Koenker (2005) includes a discussion of appropriate solution approaches for this specifically structured LP.

Acknowledgments
We acknowledge the important contributions of the following people to the OnTARGET and MAP projects: Jorge Arroyo, Georges Atallah, Kevin Bailie, Madhavi Bhupathiraju, Mike Burdick, Upendra Chitnis, Sheri Feinzig, Steve Garfinkle, Katherine Hanemann, Joan Kennedy, John LeBlanc, Agatha Liu, Mark Niemaszyk, Mike Provo, Diane Statkus, Shan Sundaram, Nancy Thomas, and Catherine Zarate. We are particularly indebted to John Milne for his active involvement and guidance as our coach in the Franz Edelman Competition.

S. Rosset's current affiliation is the Department of Statistics, Tel Aviv University, Tel Aviv 69978, Israel.

References

Berry, M. J. A., G. S. Linoff. 2004. *Data Mining Techniques: For Marketing, Sales, and Customer Relationship Management*. John Wiley & Sons, Hoboken, NJ.

Breiman, L., J. H. Freidman, C. J. Stone, R. A. Olshen. 1984. *Classification and Regression Trees*. Wadsworth International Group, Belmont, CA.

Hastie, T., R. Tibshirani, J. H. Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. Springer, New York.

Koenker, R. 2005. *Quantile Regression*. Cambridge, Cambridge University Press, Cambridge, UK.

Langford, J., R. Oliveira, B. Zadrozny. 2006. Predicting conditional quantiles via reduction to classification. Retrieved June 29, 2009, http://hunch.net/~jl/projects/reductions/median/median_uai.pdf.

Lawrence, R., C. Perlich, S. Rosset, J. Arroyo, M. Callahan, J. M. Collin, A. Ershov et al. 2007. Analytics-driven solutions for customer targeting and sales-force allocation. *IBM Systems J.* **46**(4) 797–816.

Ledingham, D., M. Kovac, H. L. Simon. 2006. The new science of sales force productivity. *Harvard Bus. Rev.* (September) 124–133.

Merugu, S., S. Rosset, C. Perlich. 2006. A new multi-view regression approach with an application to customer wallet estimation. *Proc. 12th SIGKDD Internat. Conf. Knowledge Discovery Data Mining, Philadelphia*, ACM, New York.

Morgan, A. J., S. A. Inks. 2001. Technology and the sales force: Increasing acceptance of sales force automation. *Indust. Marketing Management* **30**(5) 463–472.

Perlich, C., S. Rosset, B. Zadrozny. 2006. Quantile modeling for marketing. *Proc. 12th SIGKDD Internat. Conf. Knowledge Discovery Data Mining, Philadelphia*, ACM, New York.

Rosset, S., R. Lawrence. 2006. Data-enhanced predictive modeling for sales targeting. *SIAM Conf. Data Mining, Bethesda, MD*, SIAM, Philadelphia.

Speier, C., V. Venkatesh. 2002. The hidden minefields in the adoption of sales force automation technologies. *J. Marketing* **66**(3) 98–111.

Zikmund, W. G., R. McLeod, F. W. Gilbert. 2002. *Customer Relationship Management: Integrating Marketing Strategy and Information Technology*. John Wiley & Sons, Hoboken, NJ.