On optimizing the selection of business transformation projects

N. Abe
R. Akkiraju
S. Buckley
M. Ettl
P. Huang
D. Subramanian
F. Tipu

To compete and thrive in a changing business environment, a business can adapt by initiating and successfully carrying out business transformation projects. In this paper we propose a methodology for the optimal selection of such transformational projects. We propose a two-stage methodology based on (1) correlation analytics for identifying key drivers of business performance and (2) advanced portfolio-optimization techniques for selecting optimal business-transformation portfolios in the face of resource constraints, budget constraints, and a rich variety of business rules. We illustrate our methodology through a case study from the electronics industry.

INTRODUCTION

Businesses today are challenged by a continually changing business environment. They have to cope with global competition, keep pace with advances in technology, comply with government regulations, and, at the same time, keep expenses under control. In this context, enterprises are under pressure to undertake transformational initiatives that enable the implementation of new business strategies. Although it is clear that funding for such transformational projects should be directed to those that produce the best financial results, there is no reliable methodology for selecting these projects. In practice, business leaders conduct strategic planning exercises to define and prioritize qualitative links between financial metrics and operational objectives. Business transformation funding is then allocated for the top-priority operational objectives emerging from these exercises. However, relying solely on the qualitative links between metrics might lead to inaccurate estimations of the impact of implementing business transformation projects. This, in turn, might result in the selection of a suboptimal mix of projects. The following related issues appear to be at the heart of the adaptability of companies to current business environments: (1) the ability of management to understand the relationships between various value drivers and (2) the ability of management to select the right mix of business transformation initiatives that maximize business impact while minimizing risk. We show here how a thorough understanding of the interrelationships between financial and

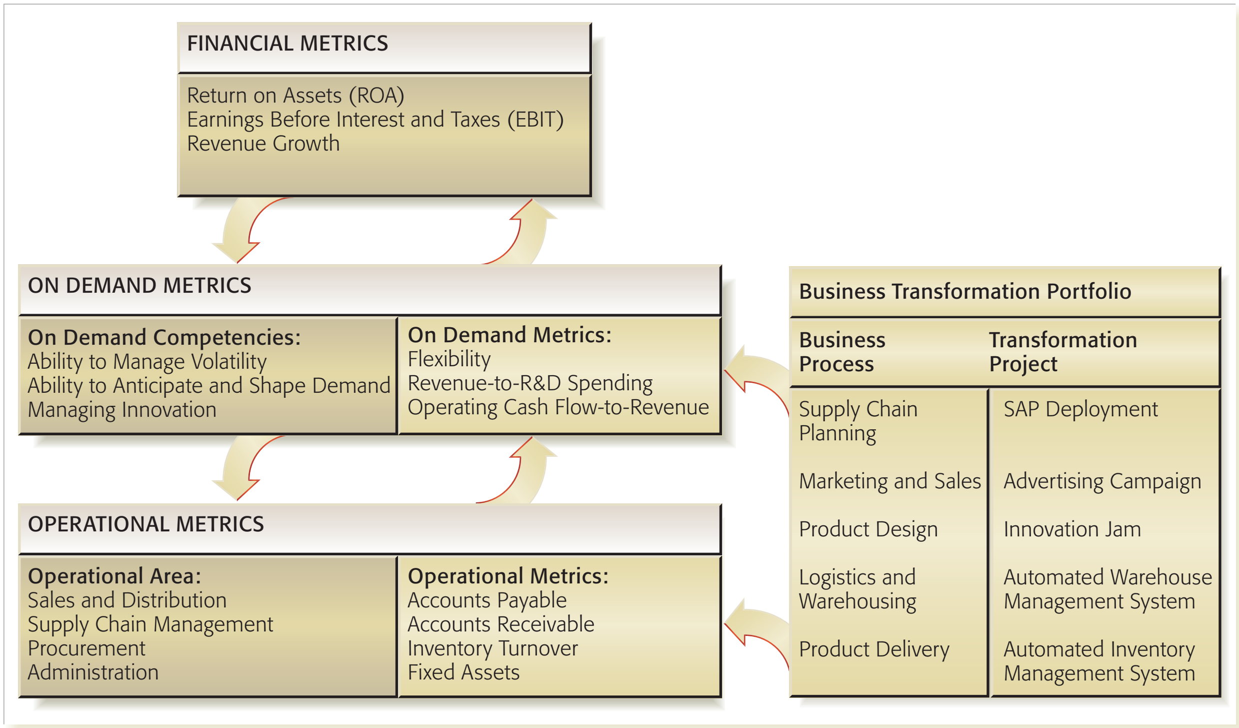

Figure 1
Relationships between business transformation projects and financial, on demand, and operational metrics

operational metrics can help make business transformation decisions; furthermore, we propose a set of metrics that business executives can use in order to evaluate transformational projects.

Traditionally, companies have measured their business performance through financial and operational metrics. Operational metrics form the foundation for measuring the business performance of a company and are typically monitored on an ongoing basis. Financial metrics represent the financial performance of an enterprise and are tracked by investors on a quarterly basis. However, new realities are forcing companies to go beyond traditional thinking in how they measure themselves. For example, in the past companies measured themselves against the return they captured for themselves in a business partnership. Now, as more and more specialized firms enter the global market, companies are seeking collaborative partnerships in various areas, including collaborating on core business functions, such as product development or research. Managers are starting to focus on building sustainable business partnerships rather than on short-term results. The new global realities are making it necessary to think about and evaluate business initiatives differently.

These new developments in the business environment require an entirely new set of metrics, those that measure the ability of companies to adapt to these changing times. These new measures quantify the ability of a company to anticipate and shape demand, manage volatility, manage value networks and business ecosystems, and manage collaborative innovation. These metrics, called *on demand* metrics, represent the resilience of an enterprise to environmental change and its long-term financial performance. On demand metrics are intended to measure the ability of a company to adapt to changing times over a longer horizon. They help identify and measure the enduring impact of business transformation activities. Leaders must determine the state of the business with respect to on demand capabilities, and they must identify the areas of greatest opportunity to improve financial and operational performance through business transformation.

Figure 1 shows an example metrics network consisting of financial, on demand, and operational metrics. The top level represents financial metrics, such as Return on Assets (ROA) and Earnings Before Interest and Taxes (EBIT), that are reported by publicly traded companies in Security and Exchange Commission (SEC) filings and company annual reports. The middle level comprises metrics that represent the on demand competencies of a company. For example, the return that a company derives from its research and development (R&D) investments is an indicator of its ability to effectively manage innovation and can be quantified by the ratio of the company's revenue to its R&D spending. The bottom level represents operational metrics that are organized by operational competency, such as sales and distribution, supply chain management, and procurement. Inventory turnover is an example of an operational metric that measures supply chain efficiency. Because business transformation projects usually aim to improve the performance of specific business processes, they will directly impact a subset of operational or on demand metrics pertaining to the targeted business process. An example is the deployment of an SAP** supply-chain management solution which may be part of a large supply-chain transformation initiative to optimize supply chain processes and metrics. A direct result of the SAP solution deployment might be a higher inventory turnover at the operational level and an increase in operating cash flow at the on demand level.

Although managers often know intuitively which operational metric impacts which financial metric, they have difficulty when they have to quantify these relationships. To accomplish this, it is necessary to track the metrics over time and use these data to analyze and correlate operational with financial metrics. In the case of a cause-and-effect relationship, the causality in the metric network helps identify the operational metrics that have the most impact on the underperforming financial metrics. This enables executives to channel business transformation investments to those projects that have the most impact on the financial performance of the enterprise. Based on measured changes in operational metrics, the metric network can also be used proactively as a sensor to predict changes in financial metrics, providing useful and potentially critical trends to executives. Once customer problem areas are identified and business value drivers established, the proposed methodology uses them in conjunction with metrics relationships to predict company performance.

The rest of this paper is organized as follows. In the next section we discuss related work, which includes work on business performance measurements and portfolio optimization. In the section that follows we present our methodology, which includes metric correlation analysis and portfolio optimization steps. The proposed methodology combines value driver analysis and continual calibration of metrics relationships for prioritizing investments in business transformations. After we introduce on demand metrics, a new set of metrics that helps evaluate the ability of a business to adapt to change, we present a case study from the electronics industry that illustrates our approach. In the last section we summarize our results.

RELATED WORK

In this section we examine related work in two areas: business performance measurements and project portfolio optimization.

Kaplan and Norton present a framework for measuring the effectiveness of an organization through the balanced scorecard, a model that integrates four perspectives: financial, customer, business process, and learning and growth.[1–3] However, a balanced scorecard does not specify the relationships among metrics either within a perspective or across perspectives. Other approaches to defining metrics for standard business processes include the Intangible Assets Monitor,[4] the Capability Maturity Model**,[5] Total Quality Management,[6] Intellectual Capital Rating[7] and industry efforts within the American Productivity and Quality Center (APQC).[8] Although all these offer different perspectives, they do not provide any method for discovering the value drivers and the relationships among the metrics. Traditionally, financial models such as net present value (NPV) models have been used to drive the investment decisions of a company;[9] however, the assumptions made in calculating the value of investments are typically based on qualitative estimates of how metrics impact one another. Chan defines a framework for measuring the performance of supply chains consisting of measurements such as cost, resource utilization, flexibility, visibility, and innovativeness.[10] Although this work presents a multi-attribute decision technique to evaluate the importance of different performance measurements, it does not cover the impact of these metrics on financial performance.

IBM and FinListics Solutions Inc. worked together to identify hierarchical dependencies between metrics within specific industries.[11] The authors developed qualitative hierarchies of performance metrics for a number of industrial sectors. For example, the financial metric *Inventory Turnover* depends on *Raw Materials Inventory, Work In Progress Inventory*, and *Finished Goods Inventory*. At the next lower level, *Raw Material Inventory* depends on *Forecasting Accuracy, Raw Material Inventory Policy, Replenishment Frequency, Supplier Performance*, and *Raw Materials Lot Sizing*. Although the hierarchies developed in this effort enable recommendations of product and services, the hierarchies are static, and metrics dependencies are not being quantified in any way.

Among commercial efforts in this area, Performancesoft Inc.[12] recognizes the need for considering leading indicators, such as the ratio of the revenue of a firm to its R&D investments (or revenue-to-R&D spending), along with lagging indicators, as well as the relationships between the lead and lag indicators. Its software product helps identify cause-and-effect relationships within a balanced-scorecard framework through the use of strategy maps.[1] Synergex Corporation offers VisualSmart**, a product that enables a metrics-based management process.[13] VisualSmart supports historical trend analysis of metrics to establish baseline values or performance ranges for a company's metrics. VisualSmart also provides visual dashboards to monitor performance metrics and offers drill-down analysis and exception management. Hyperion Solution Corporation's Performance Score Card** and Metrics Builder** products offer business performance management functions and historical trending analysis to discover relationships among metrics in an effort to identify key value drivers.[14] Our approach differs from these approaches in two ways. First, our on demand metrics, which are predictive of the long-term performance of a company, are used in addition to traditional operational and financial metrics. Second, we combine trend and metrics correlation analysis with project portfolio-optimization techniques for making investment decisions in transformational projects.

In the area of portfolio optimization, the concept of managing enterprise-wide projects as an investment portfolio is well-established in the industry.[15] Existing tools in general can be grouped into two categories according to the kinds of decisions they support: strategic or operational. Strategic tools often provide an interactive process to capture project information and then assign scores to projects to support decision making. Saaty's Analytic Hierarchy Process and the Analytic Network Process (AHP and ANP) are well-known approaches in this category.[16,17] Dickinson, Thornton, and Graves developed a dependency matrix approach for project prioritization at Boeing Corporation.[18] A common feature of these strategic tools is that they can accommodate multiple criteria and help the user reach a balanced decision. Operational tools focus on operational issues such as scheduling and resource allocation. These tools often employ sophisticated optimization technologies to find the optimal solution in terms of a few predefined metrics, such as financial returns, project time spans, and resource utilization. A comprehensive description of the underlying technology can be found in Reference 19. Our methodology combines both approaches. We provide support for multi-criteria decision making, and, if the user requires it, operational resource allocation capability. Our methodology can also handle dynamic portfolio selection rules submitted by the user. This feature, which is not available in commercial software tools, is necessary for effectively supporting project selection.

METHODOLOGY

The methodology we propose involves a two-stage process. In the first stage we apply advanced regression techniques in order to characterize the relationships among these metrics, which enable us to understand how the operational and on demand metrics impact financial performance. The second stage involves a portfolio optimization model in which several transformational projects are simultaneously considered for funding. If one can estimate how each of these transformational projects maps to operational and on demand metric values, then *transform regression* techniques (see the description of this advanced regression technique in the next section) can be used to estimate how each transformational project impacts financial metric values. We develop a multiperiod mixed-integer linear-portfolio-optimization model that contains binary variables for the selection of projects over a finite planning horizon. The constraints of the portfolio model ensure that the projects have sufficient resources.

Metric correlation analysis

The goal of the metric correlation analysis is to characterize the impact of various operational and on demand metrics on a financial metric of interest. The impact can be naturally formulated as a regression problem of estimating the changes in the financial metric as a function of changes to the lower-level metrics. In attempting to solve this problem, we face two technical challenges. The first is the presence of complex nonlinear relationships among the financial metrics, which makes it difficult to solve the regression problem satisfactorily with the standard techniques of linear regression. The second is due to the correlations that exist among the explanatory variables (i.e., the lower-level metrics), which may mask the causal relationships of interest between explanatory and target variables that are essential for conducting the analysis.

We address the former issue by using transform regression, which allows us to model complex nonlinear relationships between the explanatory variables and the target variable.[20] For the latter issue, we augment this regression technique with a causal modeling technique based on a particular simple subclass of Bayesian networks called *dependency trees*.[21] Specifically, dependency tree modeling is applied on the nonlinear transforms of the original explanatory variables, which are intermediate outputs of the transform regression, in order to capture the correlation structure among these transforms. We can then use the discovered structure to help weed out some explanatory variables that may be masking other variables of more direct interest. We describe each of these techniques in more detail in the following subsections.

Transform regression

The transform regression algorithm is an advanced regression method that goes beyond traditional regression methods, such as stepwise linear regression. It is inspired by the gradient-boosting method of Friedman, Hastie, and Tibshirani.[22] Its main advantages over existing methods are that: (1) it applies a nonlinear transformation to the explanatory variables in its modeling process and thus handles nonlinear dependence and interactions among variables; and (2) it enjoys superior predictive accuracy as compared to other existing tools and methods.

Transform regression is loosely motivated by the Kolmogorov Superposition Theorem and applies it in the context of gradient boosting (see Reference 23, for example). The Kolmogorov Superposition Theorem states that every continuous function can be expressed as the sum of a relatively small number of functions, each of which being a linear combination of transformations of the input variables. Gradient boosting is a new technique that was obtained by generalizing Freund and Schapire's renowned AdaBoost procedure for classification.[24] Intuitively, at each stage an estimator is used to approximate the input function, and in subsequent stages, the residuals from the previous stage are approximated using the same estimator, so as to minimize the estimation error with respect to the residuals. The process is then continued until near convergence. The final output model is the weighted additive model consisting of all the models obtained in the respective stages. Transform regression performs gradient boosting, using in each stage a linear function of nonlinear transforms, thus resulting in a final model in the form of the Superposition Theorem.

The implementation departs in a number of ways from the theory outlined above. First, the nonlinear transform is obtained by using a particular regression method called the Linear Regression Tree (LRT) algorithm, which is one of the functions available in the ProbE regression engine.[20] Specifically, each transform is obtained as an LRT on the raw variable in question. In addition, instead of allowing the function of these transformed variables to be arbitrary at each stage, it is restricted to be the simple sum of all the transforms, except that these transforms are allowed to depend on the outputs of all models from the preceding stages, realizing the desired richness in expressive power of the resulting model class. In particular, this is realized by obtaining the transform of each variable as an LRT, allowing as input variables the variable in question as well as the outputs of all previous models.

Table 1 Dependency forest estimation algorithm for continuous variables

BEGIN

1. Let $T := \{\}$ and $V := \{x_i, : i = 1, 2, \ldots, n\}$;
2. Calculate the value $\theta(x_i,x_j)$ for all node pairs (x_i,x_j): $\theta(x_i,x_j) = I(x_i,x_j) - (\log N/2N + \log_2 n)$;
3. Sort the node pairs in descending order of $\theta(x_i,x_j)$ and store them into queue Q
4. While $(\max\{\theta(x_i,x_j) : (x_i,x_j) \in Q\} > 0)$ do
5. Remove arg max $\{\theta(x_i,x_j) : (x_i,x_j) \in Q\}$ from Q;
6. If x_i and x_j belong to different sets W_1 and W_2 in V then
7. Replace W_1 and W_2 with $W_1 \cup W_2$ and add edge (x_i,x_j) to T;
8. Output T as the set of edges of the dependency forest

END

Transform regression, available as the regression function in DB2* Intelligent Miner* Modeling, Version 8.2, is equipped with a visualization module that includes feature importance information.[25] The feature importance information is obtained by performing variable perturbations for each variable using the underlying regression model. Thus, the feature importance score reflects the expected incremental change in the target variable corresponding to a random perturbation in the associated explanatory variable. To the extent that the model captures the nonlinear effect of each explanatory variable, the feature importance also reflects such effects. However, the importance measure of a given feature fundamentally depends on the particular model produced by the algorithm, and hence, is not free of some fundamental shortcomings common in any regression methods. For example, if two explanatory variables are highly correlated, it is very likely that the regression model will include one but not the other, at least with a significant coefficient. In such a case, one of the variables will receive a high feature importance value, whereas the other will be assigned a negligible feature importance value. This is not a problem if the sole goal of the modeling is to predict the target variable, but it is a concern if the goal is to understand what variables play an important role in determining the target variable, which is the case here.

Dependency trees

The dependency-tree tool is intended, in part, to address this issue, and works with the output model of the transform regression. The dependency-tree algorithm is based on the classic maximum-likelihood estimation method for dependency trees[21] and the corresponding estimation algorithm that is based on the Minimum Description Length principle.[26] The algorithm employed in this study has been extended to handle continuous variables, whereas the original algorithm was formulated only for categorical (i.e., discrete) variables. A dependency tree is a certain restricted class of probability models for a joint distribution over a number of variables $x_1, x_2, \ldots, x_n$ that takes the following form:

$$P(x_1, x_2, \ldots, x_n) = P(x_1) \prod_{(x_i, x_j) \in G} P(x_i|x_j),$$

where G is a graph, which happens to be a tree with root x_1. A dependency forest is a finite set of dependency trees that is defined over disjoint subsets of the entire set of variables. The algorithm exhibited in ***Table 1*** is guaranteed to find an optimal dependency forest with respect to the Minimum Description Length principle.

Note that N denotes the size of the training sample, and n denotes the number of variables. $I(x_i,x_j)$ denotes the empirical mutual information between the two observed continuous variables x_i and x_j as defined in equation (1), assuming that they are both Gaussian variables:

$$I(x_i, x_j) = \frac{1}{2}\left(1 + \log\left(\frac{\sigma^2_{x_i}\sigma^2_{x_j}}{\sigma^2_{x_i}\sigma^2_{x_j} - \sigma^2_{x_i,x_j}}\right)\right) \qquad (1)$$

In our analysis, we apply the algorithm shown above to the transformed features used in an intermediate stage of the transform regression algorithm, rather than the raw variables. This can be done by extracting the correlation table information from the output model of transform regression, as this table is computed with respect to the transformed variables. In particular, they are the outputs of the first stage in the theory described in the section on transform regression; that is, each variable x_i is transformed to $h(x_i)$, where h is the output variable of the univariate LRT; only x_i is used in splitting variables (in the tree) and is the model variable in the leaf models. That is, $h(x_i)$ is a univariate piecewise linear regression model of the target variable in terms of x_i.

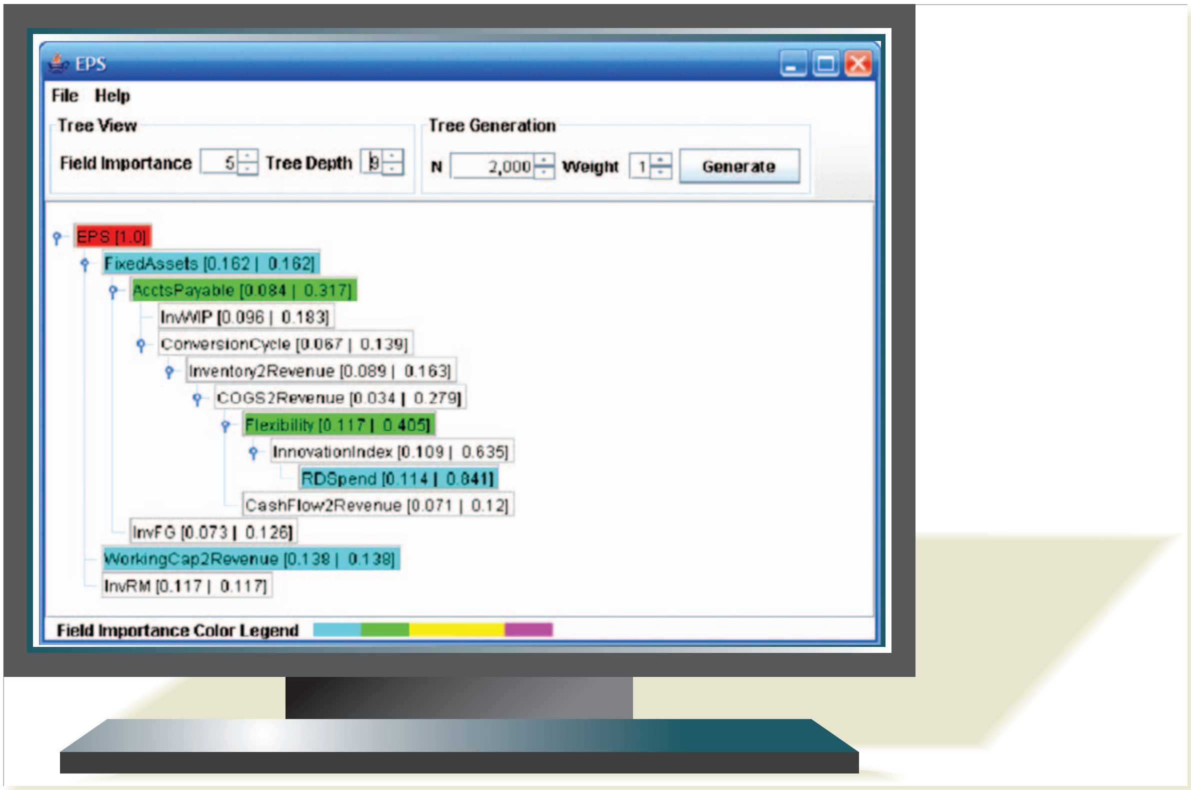

Figure 2
Dependency forest and feature importance for modeling *Earnings per Share* as target metric

Figure 2 shows the output of applying the above dependency forest algorithm to the output model of the transform regression for modeling Earnings per Share (EPS) as the target variable. The red node is the target variable. The other color-coded nodes represent variables that carry the highest feature importance in the transform regression model. In this output tree, we see, for example, that the on demand metric Flexibility is identified as one of the key value drivers for the target metric EPS, but Cash Flow-to-Revenue is not. This is indicated in the figure by the color coding (i.e., the green color of Flexibility indicates that it is among the top five features in terms of feature importance, whereas the white color of Cash Flow-to-Revenue indicates that it is not among the top five). With Cash Flow-to-Revenue having a high correlation with Flexibility, one may argue that it, too, plays a key role in the determination of the target variable. If it is the case that Flexibility is hard to control but Cash Flow-to-Revenue is conceivably easier to control, then it may make more sense to track and manage the latter than the former. This example illustrates a potential use of the dependency forests, as an additional source of information to the feature-importance-information output produced by regression modeling.

Dependency forests are one of many potential techniques one might use to analyze the correlation structure among variables. One possibility would be to use tools for structured learning of less restricted classes of graphical models (also known as Bayesian networks[27]). The other possibility is to use some other method of visualizing the information present in a correlation matrix for the explanatory variables, such as multidimensional scaling.[28]

Portfolio optimization

Portfolio selection is typically driven by multiple stakeholders with differing, sometimes conflicting, interests. The purpose of portfolio selection is to find a balanced portfolio that reconciles all these criteria. Our portfolio optimization system supports an interactive process that enables the stakeholders to generate, analyze, and compare different scenarios

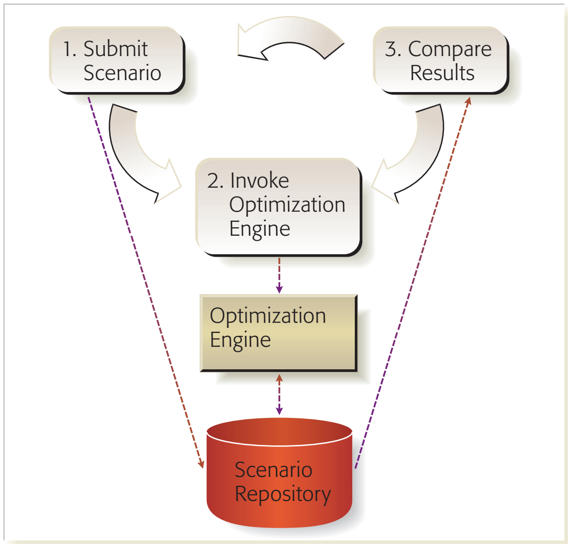

Figure 3
Portfolio optimization process

and to make the final portfolio decision. This process consists of three steps as illustrated in *Figure 3*. In step 1, analysts prepare and submit a scenario to the scenario repository. In step 2, they invoke the optimization engine, which takes the submitted scenario as input, recommends a portfolio, and saves the recommendation into the scenario repository. Finally, in step 3, the saved recommendation is compared with existing recommendations corresponding to previously submitted scenarios. If the analysts wish to examine additional scenarios after step 3, they can do so by restarting the process from step 1.

The optimization engine, the central component of our portfolio optimization system, maps a submitted scenario into a mathematical programming problem, solves that problem, and saves the result. The optimization engine can handle an arbitrary number of user-defined metrics; an arbitrary number of user-defined, time-phased resources; and many different kinds of project selection rules. This flexibility is achieved by formulating the mathematical programming problem in such a way that any arbitrary collection of user inputs is first mapped into a rich set of predefined system features and then automatically injected into the formulation.

Next we define the notation required for formulating the problem:

- P: Number of projects, indexed by $p = 1, 2, \ldots, P$.
- M: Number of metrics, indexed by $m = 1, 2, \ldots, M$.
- T: Maximum number of time periods in the planning horizon, indexed by $t = 1, 2, \ldots, T$.
- R: Number of renewable resources, indexed by $r = 1, 2, \ldots, R$.
- $H_{r,t}$: Available resource r at time t.
- E_t: Available budget at time t.
- $v_{p,m}$: Contribution of project p to metric m.
- S_p: Time span of project p, $1 \leq S_p \leq T$.
- $a_{p,s}$: Required investment for project p in its time span, $1 \leq s \leq S_p$.
- $b_{p,s}$: Expected benefit of project p over interval T from the starting time, $1 \leq s \leq T$.
- $c_{p,r,s}$: Required amount of resource r for project p in its time span, $1 \leq s \leq S_p$.
- $X_{p,t} \in \{0,1\}$: Equals 1 if project p is selected and started at time t; 0 otherwise.
- $Y_p \in \{0,1\}$: Equals 1 if project p is selected; 0 otherwise.

We maximize the weighted sum of all user-defined metrics, where w_m is the weight of metric m as shown in the objective function (2):

$$\max \sum_{m=1}^{M} W_m \sum_{p=1}^{P} v_{p,m} Y_p \tag{2}$$

The constraints are as follows:

$$\sum_{t=1}^{T} X_{p,t} = Y_p, \forall p \tag{3}$$

$$Y_p \leq 1, \forall p \tag{4}$$

$$\sum_{p=1}^{P} \sum_{s=1}^{S_p} c_{p,r,s} X_{p,t-s+1} \leq H_{r,t} \forall r, t \tag{5}$$

$$Z_0 = 0 \tag{6}$$

$$Z_t = E_t + Z_{t-1} - \sum_{p=1}^{P} \sum_{s=1}^{S_p} a_{p,s} X_{p,t-s+1}, \forall t \tag{7}$$

$$Z_t \geq 0, \forall t \tag{8}$$

The constraints (3) and (4) ensure that each project is selected at most once. Constraint (5) guarantees that the amount of resource of each type allocated to all selected projects cannot exceed the available amount at any time t. Constraint (7) is the budget-balancing equation, which tracks the leftover budget in each period t. We assume that any leftover budget at the end of a time period carries forward completely to the start of the next time period. Constraint (8) requires that this leftover is positive in each period; that is, in each period, the budget is sufficient for funding the selected projects.

Besides the constraints listed above, we also map a variety of project selection rules to constraints. For example, the budget allocation rule "at least 20 percent of the total budget should be allocated to business transformation projects" can be represented by the following constraint:

$$\sum_{p=1}^{P} \phi_p Y_p \left(\sum_{s=1}^{S_p} a_{p,s}\right) \geq K \qquad (9)$$

The indicator ϕ_p is 0 or 1, indicating whether a specific project p belongs to the set of projects corresponding to a given rule. In the preceding example, if project p is a transformational project, then $\phi_p = 1$; otherwise, $\phi_p = 0$. Constraint (9) requires that the total budget consumed by all such business transformation projects is greater than a threshold (such as 20% of the total budget). Each time a budget allocation rule is specified by the user, an appropriate instance of constraint (9) is automatically added to the formulation. This ensures that all user-defined rules are enforced. Our system supports many other types of business rules, which are handled in a similar fashion.

In any practical application, the contribution of project p to metric m, namely $v_{p,m}$, cannot be simply characterized by a deterministic estimate (e.g., the value of the expected contribution), but must incorporate some measure of risk that captures the degree of uncertainty of the project outcome. This is also true, in general, for other model parameters such as $a_{p,s}$, the required investment for project p over its time span s; $b_{p,s}$, the expected benefit of project p over interval s from the starting time; and $c_{p,r,s}$, the amount of resource r required for project p over its time span s. Incorporating all the probability distributions into a single mathematical programming problem leads to a stochastic optimization formulation in which one may define the objective function and constraints using appropriate measures of risk, such as Value-at-Risk, or Conditional Value-at-Risk.[29] In order to account for the degree of uncertainty in the outcome of a transformational project, we have implemented a stochastic approach based on Monte Carlo sampling. First, we sample several realizations of the stochastic parameter set from the probability distributions that are derived in the transform regression (Monte Carlo samples); then, we run an instance of the corresponding deterministic optimization model for each Monte Carlo realization. Subsequently, we track the potential benefits of each proposed transformational project with the *membership fraction*, a coefficient that captures the fraction of Monte Carlo runs in which each proposed project was selected in the optimal portfolio. The membership fraction imposes a simple ordering on the set of projects, which may be further used to prioritize projects.

CASE STUDY

We demonstrate the proposed two-stage approach in a hypothetical case study from the electronics industry. As previously mentioned, we define a set of on demand metrics that help assess the readiness of the business to adapt to changing market conditions. On demand metrics are intended to measure the enduring impact of business transformation activities. The definition of these metrics is grounded in the four perspectives outlined in the balanced scorecard model,[1] that is, financial perspective, customer perspective, internal process perspective, and innovation perspective. For this case study, we select a small number of on demand metrics that can be defined using publicly available metrics. For example, to measure the ability of the company to successfully adapt its internal processes over a time horizon $[t, T]$, we define the flexibility ratio $f_{t,T}$:

$$f_{t,T} = \frac{\mathrm{Cost}_t/\mathrm{Revenue}_t}{\mathrm{Cost}_T/\mathrm{Revenue}_T} \qquad (10)$$

Intuitively, this ratio defines flexibility as the ability of an enterprise or a business unit to expand margins as revenues rise and maintain margins as revenues in successive financial quarters decline. A flexibility ratio of less than one indicates higher flexibility. It is computed using generally available and reported metrics, such as operating cost and revenue derived from SEC filings and company annual reports. In addition to the flexibility ratio, we use the following metrics to measure the on demand readiness of a company:

- *Managing Volatility*
  - Capital Expenditure-to-Revenue
  - Current Ratio
  - Working Capital-to-Revenue
  - Costs of Goods Sold (COGS)-to-Revenue
  - Selling, General, and Administrative Expenses (SG&A)-to-Revenue
  - Operating Cash Flow-to-Revenue
  - Flexibility Ratio

- *Anticipating and Shaping Demand*
  - Inventory Cost-to-Revenue
  - Inventory Turnover
  - Cash Conversion Cycle
  - Net Working Capital Ratio

- *Innovation*
  - Revenue-to-R&D Spending
  - *Business Week* Innovation Index[30]
  - Capital Expenditure-to-Revenue

By tracking the on demand metrics described above, business executives can determine if their business is on its way to achieving the on demand capabilities that it needs to adapt to changing market conditions. In addition, an understanding of these metrics and the relationships between the different tiers of metrics (operational, on demand, and financial) can help companies better estimate the impact of business transformations, which in turn helps select the optimal mix of transformational projects.

Overview

We conducted a comprehensive numerical case study in which we selected over 700 public companies: 206 companies from the "semiconductor and semiconductor equipment" sector and 527 companies from the "technology hardware and equipment industry" sector. To avoid introducing bias, we eliminated companies with annual revenues of 25 million dollars or less. For each of the selected companies, we extracted operational and financial metric data for the seven years from 1998 to 2004 from financial reports in the Standard & Poor's Compustat** database, and we computed the on demand metrics. Most of the metrics selected in our case study were reported on a quarterly basis. All metric data was normalized (i.e., the metric values were normalized by subtracting the sample mean and dividing it by the standard deviation within each industry group) and statistical outliers were discarded.

Having gathered historical metric data and having computed the on demand metrics, our first goal was to better understand the relationships between the various metrics. Specifically, we were interested in determining which subset of operational and on demand business metrics were correlated with financial performance metrics for the industry as a whole. We conducted two types of analyses: (1) transform regression modeling for detecting metrics correlations and identifying the key value drivers of financial performance; and (2) predictive modeling for estimating future values of financial metrics based on operational and on demand performance. The set of hypothetical business transformation projects that formed the basis of our analysis is shown in ***Table 2***.

The table also shows the funding requirements and resource requirements of each transformational project that are used as inputs to the portfolio optimization engine. For each project, we identified operational and on demand metrics along with their expected improvements as a result of implementing each of the transformational projects. Traditional project portfolio optimization techniques analyze a portfolio of transformational projects by evaluating the funding requirements and expected operational improvements of each project. Typically, the objective is to maximize the amount of operational improvements that can be obtained, represented as a weighted sum of the operational metrics or another higher-level operational metric.

Recognizing that the ultimate objective of executives embarking upon business transformation initiatives is to maximize the financial performance of a company, we propose to use a normalized weighted sum of financial metrics as the objective function in our approach. Using the predictive models described above, we were able to translate the expected operational improvements provided in the transformation portfolio into corresponding improvements of financial metrics. Traditional methods are incapable of incorporating financial metrics in the objective function because the linkages between operational metrics and financial metrics are extremely difficult to quantify. We next present the results of our experiments.

Findings

We investigated the development of an optimal project portfolio from the preceding set of business

Table 2 Business transformation projects portfolio used in electronics case study

Business Process	Transformation Project	Implementation Cost	Financial Head Count	Operational Head Count	Metrics	Type	Expected Improvement (%)
1 Market and sell products and services	New advertising campaign	$6M	—	—	Advertising spending	Operational	5
2 Supply chain planning	SAP production planning deployment	$6M	200	300	Inventory turnover	Operational	8
					Cash conversion cycle	On Demand	5
					COGS to revenue	On Demand	−3
					Working cap ratio	On Demand	−2
					Flexibility	On Demand	−5
3 Produce/Manufacture/Deliver product	Automated inventory management solution	$5M	—	250	Inventory (WIP)	Operational	−10
4 Manage logistics and warehousing	Warehouse management solution deployment	$6M	—	200	Inventory (finished goods)	Operational	−5
					Inventory (raw materials)	Operational	−5
5 Manage financial resources	F&A shared services center project	$2M	150	—	SG&A to revenue	On Demand	−2
6 Revenue accounting	Order-to-cash management system	$4M	150	—	Accounts receivable	Operational	−8
					Days sales outstanding	Operational	−8
					Days purchases outstanding	Operational	3
					Cash flow	On Demand	5
					Cash conversion cycle	On Demand	5
7 Manage fixed assets	Repair & maintenance transformation project	$2M	—	200	Fixed assets	Operational	−3
					Current asset liabilities	Operational	−5
8 Cash disbursements	Travel & expense reporting optimization and fraud detection	$1M	150	—	Accounts payable	Operational	2
9 Design products and services	Innovation Jam project	$2M	—	—	R&D spending	Operational	2
					Capital expenditures	On Demand	−2
					Innovation index	On Demand	5

Table 3 Relative improvements of financial metrics obtained from predictive modeling

Business Process	Transformation Project	Relative Improvements of financial metrics (%)						
		Beta	EBIT	EPS	Market Cap	P/E Ratio	Revenue	ROA
1 Market and sell products and services	New advertising campaign	−0.2	6.8	11.6	5.4	1.8	11.1	9.2
2 Supply chain planning	SAP production planning deployment	6.5	8.3	−1.7	5.2	6.2	5.9	8.4
3 Produce/Manufacture/Deliver product	Automated inventory management solution	−0.1	6.8	1.1	7.1	5.4	11.7	10.8
4 Manage logistics and warehousing	Warehouse management solution deployment	−1.9	8.7	8.1	2.7	1.8	11.3	10.1
5 Manage financial resources	F&A shared services center project	−5.4	7.2	3.3	5.4	1.5	10.5	10.3
6 Revenue accounting	Order-to-cash management system	−1.7	5.2	0.9	5.4	2.3	9.5	5.5
7 Manage fixed assets	Repair & maintenance transformation project	−8.7	11.3	8.8	5.4	1.6	8.7	0.3
8 Cash disbursements	Travel & expense reporting optimization and fraud detection	−3.6	6.8	1.1	5.4	6.0	10.0	10.3
9 Design products and services	Innovation Jam project	−0.4	8.6	1.1	5.4	2.0	0.2	12.8

transformation projects, given the quantitative relationship between the operational and on demand metrics and the financial metrics provided by the metric correlation analysis. First, we describe the results of the metric correlation analysis.

Metric correlation analysis

The transform regression technique was used to develop a regression model for the expected improvement in each of the financial metrics listed in Table 2, using the dependency tree tool to help choose the explanatory variables whenever appropriate. More specifically, the modeling was done to predict the expected annual percentage change for each financial metric in the following year, as a function of the percentage changes in the operational and on demand metrics in the current year. The models were then applied to the percentage improvements in operational and on demand metrics anticipated for each of the nine transformational projects listed in *Table 3* to obtain the projected percentage improvements in each of the financial metrics.

In multiple models including that for Return on Assets (ROA), Earnings per Share (EPS), and Revenue per Employee, it was observed that fixed assets tend both to receive a high feature importance and to exhibit high correlations with other explanatory variables, accounts payable and accounts receivable in particular. It can be argued that accounts payable and accounts receivable are easier to control than fixed assets. These results suggest, therefore, that the predictive models for these metrics may very well be underestimating the effects of the former two by the masking effect of the latter. By considering this observation, we may elect to remove fixed assets as an explanatory variable in our modeling. *Figure 4* exhibits the feature importance chart for the modeling of EPS with and without fixed assets as an explanatory variable. It is evident that the removal of this variable has led to an increased importance assigned to the accounts payable variable. This will affect the predicted impact of any transformational project that targets improvement on the metric of accounts payable.

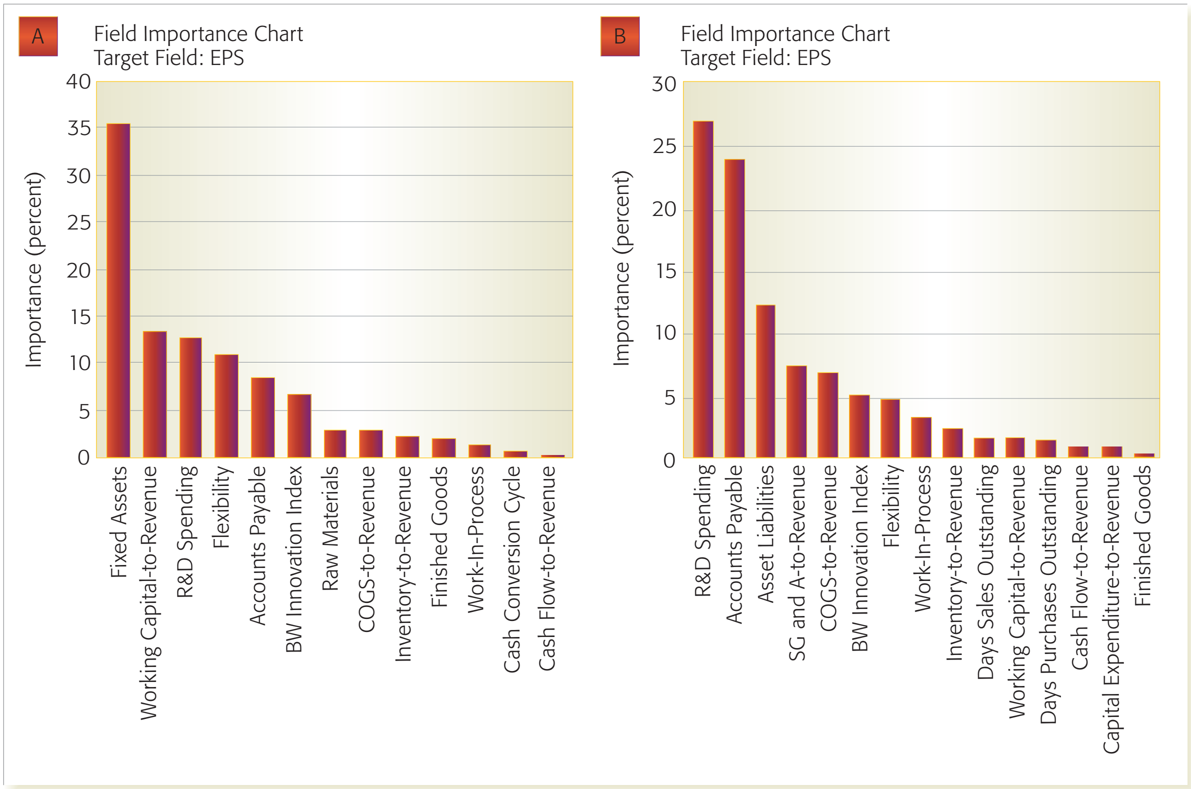

Figure 4
Feature importance information for predicting target metric EPS (A) with fixed assets and (B) without fixed assets

Having gone through these types of iterative analyses and improvements, the predictive models were applied to obtain financial metric improvement data for the various transformational projects that were considered for funding. The results are summarized in Table 3.

These estimates on the improvement of financial metrics for the transformational projects provide the necessary link to perform the subsequent project portfolio optimization for the objective of maximizing overall financial impact. Some remarks on the predictive accuracy of our metric correlation analysis are in order. One measure of the quality of our models is given by the "ranking" score given by the DB2 Intelligent Miner Transform Regression tool. The ranking score here is formally known as the *Gini coefficient*, defined as the area between the cumulative lift curve and random prediction. The Gini coefficient is 1 for the optimal ranking and 0 for the random predictor. ***Table 4*** summarizes the ranking metric values for each of the models we obtained for each of the financial metrics.

To visualize the predictive accuracy of these models, we show in ***Figure 5*** a binned scatter plot for the two target metrics Revenue and ROA. In a binned scatter plot, the x-axis corresponds to the average predicted value given by the model and the y-axis to the average actual value, where the average is taken

Table 4 Ranking metrics, or Gini coefficients, of predictive models for financial metrics

Metric	Gini Coefficient
Beta	0.458
EBIT margin	0.376
EPS	0.371
Market Cap Growth	0.519
P/E Ratio	0.216
Revenue	0.533
ROA	0.492

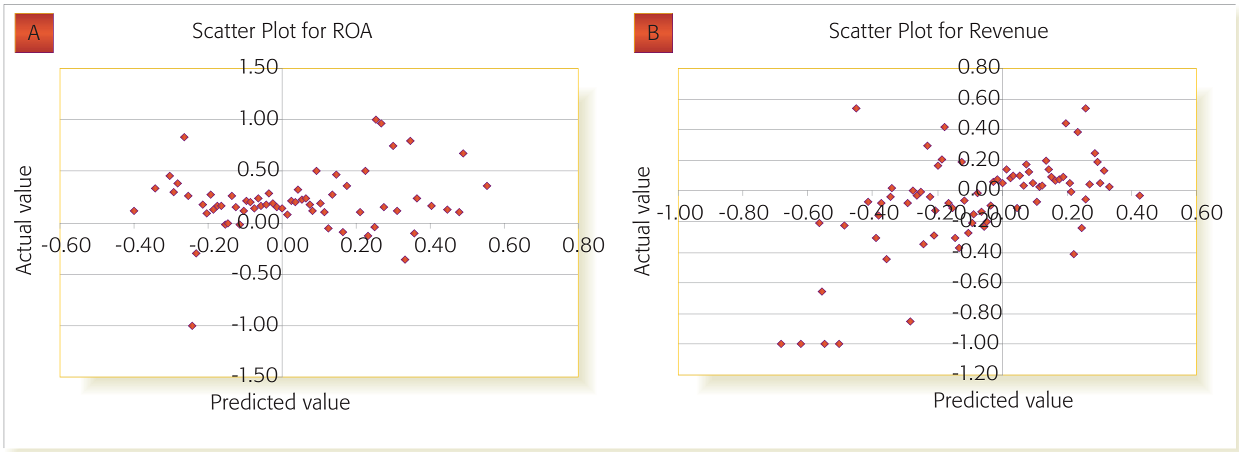

Figure 5
Binned scatter plots for (A) ROA and (B) Revenue, as given by the metrics correlation analysis models

with respect to groups of companies defined by bins of predicted values (here in 100 bins of equal size). In obtaining these scatter plots, we trained the models using data up to year 2003 and did testing on the data for year 2004. Inspecting these plots reveals that the predicted values indeed provide systematic gains over random predictions.

Although predictive accuracy is less than perfect in this exercise, we believe the predictive quality of these models indicates the potential of our approach. There are a number of possible future improvements that would most likely boost the predictive accuracy. For example, more could be done to capture and correct for seasonal effects and long-term industrial trends. Including variables that capture temporal characteristics of the industry should allow the model to become adaptive to changes in the environment. It should also help greatly to include firmographic features, such as revenue and number of employees of companies as explanatory variables in the regression modeling. In our study, all companies in the electronics industry, in some sense, are treated uniformly, and obviously significant improvement would be expected if more sophisticated treatment of various characteristics of the companies in question were used.

Portfolio optimization

We next describe the setting in which we perform the portfolio optimization, in terms of the various parameters and constraints that are pertinent to the case study. To illustrate the multiperiod nature of project portfolio optimization, we chose a portfolio execution horizon of four quarters. Each project incurs a cost for implementation and requires an estimated combination of the two resource types that are required to execute the project (financial and operational head count). The duration of each project along with its costs and resource requirements are summarized in Table 2. The available budget is treated as a consumable resource; the unconsumed amounts at the end of each quarter are assumed to carry forward to the beginning of the next quarter. The available budget is given as a time-phased profile of $10M (millions) in the first quarter, and an additional $3M in each of the subsequent quarters. The staffing resources are assumed to be renewable resources: a project that is allocated its required combination of resource types releases the allocated amounts back into the system pool of resources upon its completion. We assumed that a financial head count of 200 and an operational head count of 300 are available in each quarter.

First, we ran the optimization model in succession with each of the seven financial metrics as the sole objective of optimization, subject to budget constraints and resource constraints. The seven resulting optimal portfolios, each corresponding to one of the financial metrics as the sole optimization objective, are shown as a pivot chart in *Figure 6*.

The nine transformational projects are shown along the horizontal axis. The vertical axis records the number of optimal portfolios, out of a maximum of seven, in which each project was selected. For example, the figure shows that the project "Automated Inventory Management" was selected in four optimal portfolios, which correspond to the individual optimization of the four financial metrics, namely, Market Capitalization, Earnings per Share, Revenue, and Return on Assets, respectively. Similarly, project "Travel and Expense Reporting" was chosen in all seven optimal portfolios. A project that is chosen in a larger number of optimal portfolios has favorable implications for a higher number of financial metrics in the face of budget and resource constraints, by virtue of the metric correlation between the operational metrics and financial metrics.

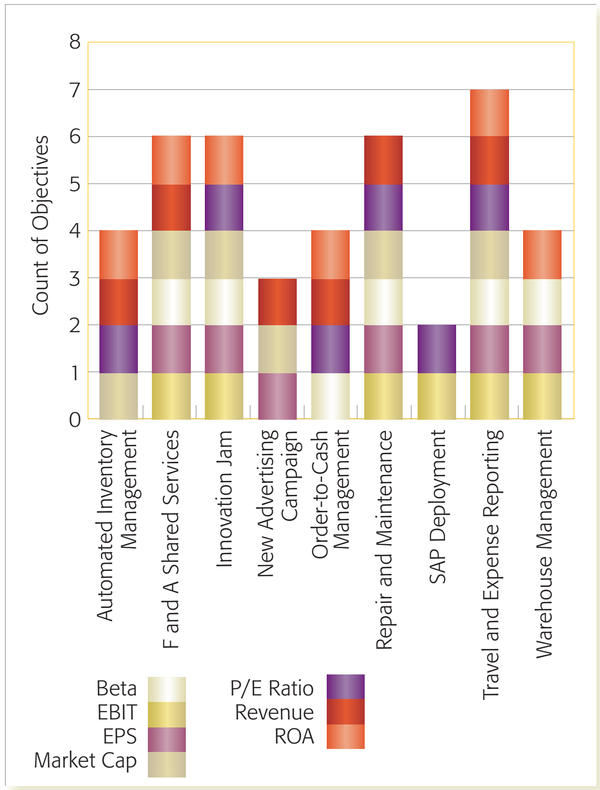

Figure 6
Optimal portfolios corresponding to each of the seven financial metrics

Comparisons to traditional portfolio analysis

In order to illustrate the value of using the results of the metric correlation analysis in portfolio decision making, we also compared a project portfolio resulting from a traditional approach that uses operational metrics with the solution resulting from our proposed approach. For the traditional case, we set up a portfolio optimization instance on the previously mentioned set of projects with a uniformly weighted linear combination of all operational metric improvements as the optimization objective. The optimal solution under resource and budget constraints is a collection of projects that delivers maximum improvement in operational metrics, which can further be evaluated for the corresponding improvement in financial metrics by using the metric correlation analysis. To arrive at a proxy financial score, we used a uniformly weighted linear combination of the corresponding financial metric improvements of the portfolio.

In our proposed approach, we solved the same problem with a uniformly weighted linear combination of all financial metric improvements as the proxy financial optimization objective. The optimal objective function corresponding to this solution can be compared with the financial score corresponding to the traditional approach. Further, the composition of projects that make up the optimal portfolio in the traditional approach and the proposed approach can also be compared. The results showed that the expected improvement of financial metrics under the proposed approach (as captured by the uniformly weighted linear combination of improvements across all seven financial metrics) is 17.9 percent higher than the corresponding improvement under the traditional approach. ***Figure 7*** shows the optimal portfolios corresponding to the traditional and proposed approaches in the form of a pivot chart. The transformational projects are shown on the horizontal axis. The vertical axis records the number of optimal portfolios in which the corresponding project was selected.

It can be seen that the two portfolios have four projects in common, namely, Automated Inventory Management, Innovation Jam, Repair and Maintenance Project, and Travel and Expense Reporting. They also have two projects that are unique to their respective portfolios. While the traditional approach selected Order-to-Cash Management and SAP** Product Deployment, the proposed approach selected F&A Shared Services and New Advertisement Campaign. This highlights the essential insight that optimizing operational metrics alone is not sufficient if the eventual objective is to optimize the financial objectives. While Order-to-Cash Management and SAP Product Deployment projects lead to higher

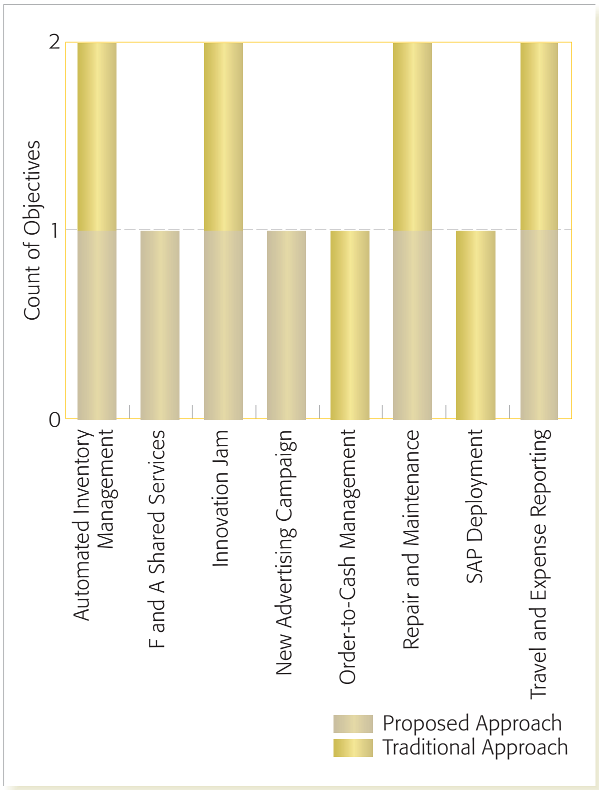

Figure 7
Optimal portfolios selected under the traditional and proposed approaches

operational metric improvements, F&A Shared Services and New Advertisement Campaign lead to operational metric improvements that have a higher correlation with financial metric improvements.

Monte Carlo simulation

This case study helps illustrate the value of using our proposed approach for portfolio decision making in the context of business transformation projects. However, we recognize that the statistical predictions of how the operational and on demand metrics transform into financial metric values of improvements is typically not a point estimate, but an interval estimate. In other words, the statistical prediction includes a variance in the estimated improvement in financial metrics that can be captured by the transform regression models. As with financial-portfolio optimization problems, it is important to incorporate the variance information as a measure of risk directly into the portfolio optimization model. This enables decision makers to identify an optimal portfolio that explicitly recognizes the range of financial improvements that could result from any set of business transformation investments.

As a first step toward incorporating the variance in the financial improvements that are delivered by the various projects, we developed a Monte Carlo analysis in which the deterministic optimization analysis is carried out for each stochastically sampled set of financial improvement metrics. We assumed that the expected values of the financial improvement predictions in Table 2 follow a normal distribution, and chose a uniformly weighted linear combination of the financial metric values as the proxy of the financial optimization objective in each deterministic optimization run. A total of 1000 Monte Carlo simulations of randomized financial improvement predictions were run through the deterministic portfolio optimizer.

The corresponding solutions are summarized in ***Figure 8*** in the form of membership fractions that can be associated with each proposed transformation project. The membership fraction tracks the fraction of the 1000 Monte Carlo runs in which the corresponding transformation project was selected in the optimal portfolio. It can be seen that the Repair and Maintenance, Travel and Expense, and F&A Shared Services projects were almost always selected in the optimal portfolio, whereas the SAP Product Deployment project was never selected. The membership fractions in Figure 8 impose a rank-ordering on the set of proposed transformational projects which may be used in a "greedy" fashion to select projects until a budget limitation or resource constraint limitation is encountered.

An alternative approach to portfolio optimization with explicit uncertainties in the financial improvement predictions in the form of variances, or similar higher-order measures of risk, can be carried out by developing a stochastic optimization formulation. Further, one could use an appropriate risk measure, such as Value-at-Risk, or Conditional Value-at-Risk, in the objective function or the constraints, or both, to identify stochastically optimal portfolios. We intend to investigate such a formulation and its application to the portfolio analysis of business transformation projects in a future study.

SUMMARY AND CONCLUSIONS

Business transformation projects are proposed and defined in terms of expected improvements to key operational metrics that describe the business process which is the target of transformation. However, the purpose of costly investments in transformational projects is to drive business value, which is measured by shareholders and stakeholders in terms of improvements in the financial metrics of the enterprise. In this study, we have highlighted the essential insight that understanding and quantifying the correlations between improvements in operational metrics and financial metrics is a crucial first step in characterizing costly transformational projects. Additionally, we have also proposed that the estimated improvements in financial metrics should be used directly in project portfolio decision making, because the final objective of business transformation projects is to deliver business value.

We have incorporated these ideas into a two-stage methodology for helping business managers select an optimal set of transformational projects by correlating business transformation investments to sustained business performance. Based on advanced correlation analysis, the methodology identifies operational metrics that play the most significant role in managing the financial performance of a company and establishes analytical metric relationships, which can be continually updated to predict company performance. The predicted improvements in the enterprise financial metrics are then directly used in a detailed portfolio optimization framework that is driven by a financial objective function to select an optimal portfolio in the face of resource constraints, budget constraints, and a rich variety of business rules. The proposed methodology could help business executives to optimally allocate business transformation funding, thereby improving a company's performance more rapidly. Traditional portfolio selection approaches that are formulated based on operational objective functions can result in portfolios that are materially different and suboptimal with respect to financial improvement. A hypothetical case study based on electronics industry data illustrates these differences and underscores the value of our proposed approach. We intend to investigate a more rigorous stochastic optimization approach in a future study. Nevertheless, highlighting the need to explicitly recognize the financial impact of operationally driven transformational projects, as well as a formal two-stage methodology to address this need, are the key contributions of this paper.

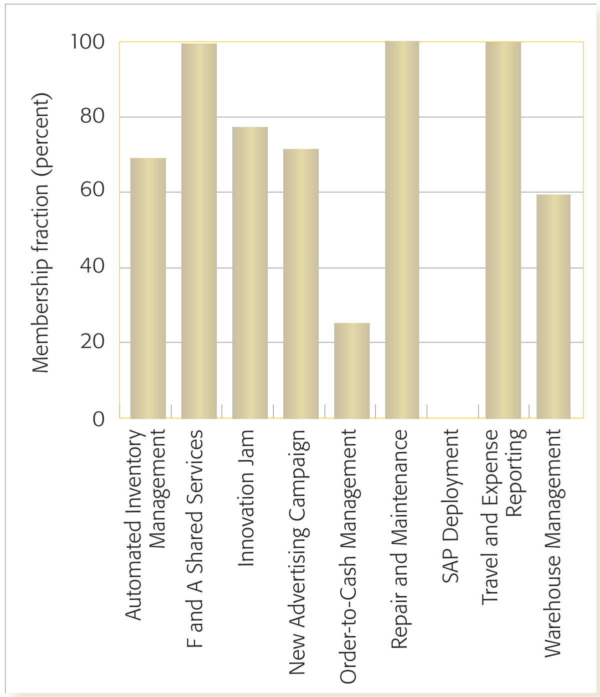

Figure 8
Membership fractions of project portfolio derived from Monte Carlo simulations

The methodology described in this paper has a number of pitfalls and limitations. First, the transform regression models assume that the correlation relationships between metrics are static, whereas in reality these relationships are not static, particularly in a rapidly changing business environment. Therefore, the operational and on demand metrics do not fully explain the financial metrics. The proposed regression models are an attempt to quantify complex interrelationships that are difficult to quantify and put to use for decision making. Although certain changes in the business environment could be modeled by introducing temporal variables that capture these changes, such as seasonal effects, the model predictions should be viewed as high-level heuristic estimates of the expected impact of operationally focused metrics on financial metrics. Second, the predictions of improvements in financial metrics are typically not point estimates, but interval estimates that incorporate a measure of risk in the outcome of a transformational project. In other words, the regression models must generate statistical predictions that include a variance in the estimated improvement in financial metrics. This requires sufficient data to allow meaningful correlation analysis between operational metrics and financial metrics. The prediction of financial metrics should be validated using suitable data partitioning into training and validation sets. And third, the portfolio optimization model should directly incorporate variance information so that the optimal portfolio selection explicitly accounts for the range of financial improvements that could result from any set of business transformation investments.

ACKNOWLEDGMENTS

The authors gratefully acknowledge contributions by Betsy Tong, Christine Kinser, and Sean Lafferty for leadership, funding, and industry expertise on metrics; Blair Binney and Dan Benoit for leadership, funding, advice, and support on portfolio optimization; and Alan King for advice on portfolio optimization. We also thank the anonymous referees for their exceptional efforts which helped significantly improve the content and presentation of the paper.

*Trademark, service mark, or registered trademark of the International Business Machines Corporation in the United States, other countries, or both.

**Trademark, service mark, or registered trademark of SAP AG, Carnegie Mellon University, Synergex Corporation, Hyperion Solution Corporation, or Standard & Poor's in the United States, other countries, or both.

CITED REFERENCES

1. R. Kaplan and D. Norton, "The Balanced Scorecard: Measures that Drive Performance," *Harvard Business Review* **70**, No. 1, 71–79 (1992).
2. R. Kaplan and D. Norton, "Putting the Balanced Scorecard to Work," *Harvard Business Review* **71**, No. 5, 134–147 (1993).
3. R. Kaplan and D. Norton, *The Balanced Scorecard: Translating Strategy into Action*, Harvard Business School Press, Boston (1996).
4. Intangible Assets Monitor, Value Based Management.net, http://www.valuebasedmanagement.net/methods_iam.html.
5. Capability Maturity Model, Software Engineering Institute, Carnegie Mellon University, http://www.sei.cmu.edu/cmm.
6. Total Quality Management, iSixSigma LLC, http://www.isixsigma.com/me/tqm/.
7. Intellectual Capital Rating, Value Based Management.net, http://www.valuebasedmanagement.net/methods_icrating.html.
8. APQC (American Productivity and Quality Center), http://www.apqc.org.
9. T. Koller, M. Goedhart, and D. Wessels, *Valuation: Measuring and Managing the Value of Companies*, John Wiley & Sons, New York (2005).
10. F. T. S. Chan, "Performance Measurement in a Supply Chain," *International Journal of Advanced Manufacturing Technology* **21**, No. 7, 534–548 (2003).
11. S. Buckley, S. Kapoor, K. Katircioglu, and O. Oladeji, *Using an Expert System on the Web to Recommend Solutions and Services*, IBM Research Report RC22053, Yorktown Heights, NY (2001), http://domino.research.ibm.com/library/cyberdig.nsf/papers/2B5446E12F3CF35785256B6C0070DD1E/$File/RC22053.pdf.
12. PerformanceSoft, Inc., http://www.performancesoft.com.
13. Synergex International Corporation, http://www.synergex.com.
14. Hyperion Solutions Corporation, http://www.hyperion.com.
15. D. L. Bolles and D. G. Hubbard, *The Power of Enterprise-Wide Project Management*, AMACOM, American Management Association (2006).
16. T. L. Saaty, *Multicriteria Decision Making: The Analytic Hierarchy Process*, Vol. 1, RWS Publications, Pittsburgh, PA (1990).
17. T. L. Saaty, *Decision Making with Dependence and Feedback: The Analytic Network Process*, RWS Publications, Pittsburgh, PA (1996).
18. M. W. Dickinson, A. C. Thornton, and S. Graves, "Technology Portfolio Management: Optimizing Interdependent Projects over Multiple Time Periods," *IEEE Transactions on Engineering Management* **48**, No. 4, 518–527 (2001).
19. C. Artigues, S. Demassey, and E. Neron, *Resource-Constrained Project Scheduling: Models, Algorithms, Extensions and Applications*, ISTE Publishing Company (2007).
20. C. Apte, E. Bibelnieks, R. Natarajan, E. Pednault, F. Tipu, D. Campbell, and B. Nelson, "Segmentation-Based Modeling for Advanced Targeted Marketing," *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pp. 408–413 (2001).
21. C. K. Chow and C. N. Liu, "Approximating Discrete Probability Distributions with Dependence Trees," *IEEE Transactions on Information Theory* **14**, 462–467 (1986).
22. J. Friedman, T. Hastie, and R. Tibshirani, "Additive Logistic Regression: A Statistical View of Boosting," *Annals of Statistics* **28**, No. 2, 337–407 (2000).
23. E. P. D. Pednault, "Transform Regression and the Kolmogorov Superposition Theorem," *Proceedings of the Sixth SIAM International Conference on Data Mining*, Bethesda, Maryland (April 20–22, 2006), Society for Industrial and Applied Mathematics pp. 35–46 (2006).
24. Y. Freund and R. Schapire, "Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting," *Journal of Computer and System Sciences* **55**, No. 1, 119–139 (1997).
25. A. Nicolussi and S. Grontzki, *DB2 Intelligent Miner: Comprehensive Guide to IBM DB2 Intelligent Miner, Version 8.2*, developerWorks, IBM Corporation, http://www-128.ibm.com/developerworks/db2/library/techarticle/dm-0506nicolussi.

26. J. Rissanen, "Modeling by Shortest Data Description," *Automatica* **14**, No. 5, 465–471 (1978).
27. D. Heckerman, D. Geiger, and D. M. Chickering, "Learning Bayesian Networks: The Combination of Knowledge and Statistical Data," *Machine Learning* **20**, No. 3, 197–243 (1995).
28. M. F. Cox and M. A. A. Cox, *Multidimensional Scaling*, Chapman and Hall, Boca Raton, FL (2001).
29. C. Alexander, *Risk Management and Analysis. Volume 1: Measuring and Modelling Financial Risk*, John Wiley & Sons, New York (1998).
30. P. Coy, "The Search for Tomorrow, BusinessWeek's New Index Examines Corporate R&D and Capital Spending," *Business Week*, 216–220 (Oct 11, 2004).

Accepted for publication May 18, 2007.
Published online September 25, 2007.

Naoki Abe
IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598 (nabe@us.ibm.com). Dr. Abe, a research staff member in the Data Analytics Research group since June 2001, is engaged in research in machine learning and data mining. After graduating from the Massachusettus Institute of Technology in 1984 with B.S. and M.S. degrees in computer science, he joined the Watson Research Center. He received a Ph.D. degree from the University of Pennsylvania in computer and information science in 1989. He served as a post-doctoral researcher at the University of California at Santa Cruz from 1989 until 1990, when he joined NEC Research Laboratories in Japan. From 1998 to 2000, he served as an adjunct associate professor at the Tokyo Institute of Technology. Dr. Abe has served on program committees of various professional conferences and is currently on the editorial board of the *Journal of Machine Learning Research and Data Mining* and the *Knowledge Discovery Journal*. His current research interests are in the applications of advanced machine learning methods to business analytics and optimization.

Rama Akkiraju
IBM Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, NY 10532 (akkiraju@us.ibm.com). Since joining IBM in 1995, Ms. Akkiraju, a Senior Technical Staff Member at the Watson Research Center, has been involved in research on agent-based decision-support systems, electronic marketplaces, and business process integration technologies. Her research interests are in the application of artificial intelligence techniques to solving business problems and semantic Web services. She received the INFORMS Daniel H. Wagner prize in 1998 for her work on scheduling algorithms for the paper industry. She has a master's degree in computer science from New York University and an M.B.A degree from the Stern School of Business at the same university.

Stephen Buckley
IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598 (sbuckley@us.ibm.com). Dr. Buckley, a research staff member since 1987, is currently the liason between IBM Research and IBM Corporate Development. Prior to that, he managed a group in the Mathematical Sciences department of IBM Research. His research interests include supply chain management, simulation, and business process management. Dr. Buckley received a Ph.D. degree in computer science from MIT in 1987. He also received an M.S. degree in computer science from Pennsylvania State University in 1978, and a B.S. degree in applied mathematics and computer science from Florida State University in 1977.

Markus Ettl
IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598 (msettl@us.ibm.com). Dr. Ettl manages the Supply Chain Analytics and Architecture group of the Mathematical Sciences Department. He received his doctoral degree in computer science in 1995 from Friedrich-Alexander University in Erlangen, Germany. Since joining IBM in 1995 as a research staff member, he has applied mathematical methods to supply chain-management problems and holds several patents in this field. His current research interests lie in decision support for production systems and logistics networks, and in sense-and-respond systems. Dr. Ettl received the INFORMS Franz Edelman Award in 1999. He also received two IBM Outstanding Technical Achievement Awards and several IBM Research Division awards.

Pu Huang
IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598 (puhuang@us.ibm.com). Dr. Huang, a research staff member, is involved in applying intelligent data-mining techniques and sophisticated decision-support algorithms to the development of sense-and-respond systems for business decision making in large-scale, distributed environments. His current interests are in supply chain management, stochastic optimization, machine learning, and data mining. Dr. Huang has an M.S. degree from the School of Computer Science and a Ph.D. degree from the Tepper School of Business at Carnegie Mellon University.

Dharmashankar Subramanian
IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598 (dharmash@us.ibm.com). Dr. Subramanian, a research staff member in the Mathematical Sciences department, is involved in research in operational risk modeling, portfolio analysis, and stochastic optimization. His research interests include mathematical modeling, mathematical programming, business and process optimization, process control of linear, nonlinear, and hybrid systems, continuous and discrete event simulation, stochastic optimization, and risk analysis. He received a Ph.D. in chemical engineering from Purdue University, where he focused on combinatorial decision making under uncertainty. Before joining IBM, he was a senior research scientist at Honeywell Laboratories in Minneapolis, where he conducted basic and applied research in supply-side energy portfolio optimization, cogeneration-based integrated energy systems, aircraft trajectory optimization, wireless sensor networks, and nonlinear process control.

Fateh Tipu
IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598 (fateh@us.ibm.com). Since joining IBM in 1991, Mr. Tipu, a senior engineer, has been involved in the development of computer-aided-design tools for logic design and, more recently, in research on data-mining systems. He received an M.S. degree in electrical engineering from University of Wisconsin-Madison in 1991. Mr. Tipu's awards include an IBM Outstanding Technical Achievement Award and an IBM Research Division award. He is a senior member of IEEE. ■