

Predicting the Future —

Reprinted
from
DM
REVIEW

A Crystal Ball for Your Enterprise

By **Claudia Imhoff**

I recently visited New Orleans where you can have your future told by one of the many palmists and crystal ball readers in the French Quarter. I am unsure of their accuracy record, but the desire to know something about our future has great appeal. We all would like to know if we will be rich or famous, get married and have kids or stay single, and so on. Because we have no way to judge whether or not these predictions will come true, we can only live our lives to find out how correct they were.

Corporate enterprises are no different in their desire to understand their future. In many ways, it is more important than ever for companies to be able to reliably discern what to expect from their revenues and expenses, whether their customer base is healthy, their product offerings profitable, their employees productive, etc. Competition and economic hardships have made the ability to foretell where the enterprise will be — in three months, six months, a year or more — critical to its overall health and viability. Even more useful are the decisions and actions that will help the enterprises reach their financial goals.

For many companies, this fortunetelling ability (i.e., budgeting and forecasting) has mostly been performed by experienced financial analysts or their peers in other departments by creating

massive spreadsheets with manually entered historical data. These spreadsheets were then massaged, reinterpreted, tweaked or otherwise altered until the analyst (or the person responsible for the prediction) was satisfied. Then, a projection was created based on the historical trends established by the numbers. The accuracy was dependent on the accuracy of the data entry, interpretation of the numbers, and assumptions that commercial conditions would remain static and that customer, suppliers, partners, etc., were all creatures of habit.

Given these rather broad and sometimes questionable conjectures, the budgeting and forecasting process was perhaps more similar to black magic than science. Because these forecasts and budgets were

made episodically (once or twice a year), their accuracy diminished as time went by. Therefore, enterprises would find themselves “revising” the budgets or forecasts quarterly or, in extreme cases, monthly to accommodate the perturbations that inevitably occurred in their businesses causing these predictions to no longer be valid.

Much of the early success of business intelligence (BI) was based in its ability to make the budgeting process easier and faster. This was due to the increased accuracy and availability of the data by integrating it from sources across the organization in a repeatable and reliable fashion. Simple trend lines and “what if” analyses gave this process a more consistent footing but still could not deliver anything more than episodic (albeit more frequent) ruminations. What was needed was something that had a mind of its own — that could monitor events, actions and trends as they occurred — so that predictions of all sorts, not just budgets and forecasts, could be created in a more precise, timely and dependable manner.

This is the world of two emerging capabilities: predictive analytics and the even more sophisticated guided decision making. The new tools and capabilities generated from data mining and statistical analysis companies open a whole new



way of doing BI. These capabilities are part business process and part related technologies. Let's start with some basic definitions of these two capabilities. We can then examine how they are used today, what vendors offer these capabilities and what you need to know to create an environment suitable for them to provide benefit to your organization.

I make a distinction between predictive analytics and guided decision making in that predictive analytics takes you to the "deliver the prediction" stage and then lets you decide what to do. Guided decision making goes one step further to suggest the actions and reactions that should occur based on its rules engine and analysis of the situation.

Both forms of prediction use a set of processes that produce reliable conclusions about your current environment and future events. These predictions are based on linking historical and current data to possible outcomes; predictive analytics, however, do not suggest the actions the enterprise should take to either maintain the trend or alter its direction. These capabilities use data (predominantly customer data) from both internal systems and external or third-party providers such as customer demographics, attitudes and behaviors. From this massively detailed data, they create statistical, mathematical or other algorithmic techniques to generate models of customer segmentations and classifications, purchasing patterns, forecasting, profiling and propensity scoring.

It is important to note that the predictive analytic models must be combined with deep business knowledge to provide the insight into what your customers are doing, their behaviors, purchasing patterns, up-selling and cross-selling opportunities, retention and acquisition procedures, and even possible fraudulent activities. Interpretation by a human is mandatory.

Guided decision-making is similar to predictive analytics in terms of the data that is needed, the types of analyses that are possible, and the trends and patterns that are discovered. The difference from predictive analytics is the ability of the guided decision making applications to proactively manage risk, suggest key decisions and actions, test the potential actions for their likely intended and unin-

tended consequences, and to choose the best course of action. The interpretation that a human must do for predictive analytics is replaced by a sophisticated set of business rules, algorithms and/or processes applied by the application itself.

These actions along with their corresponding business rules can then be deployed to the customer-facing operational systems to identify cost savings, new and alternative revenue streams, recurring process improvements and even competitive advantages.

Predictive analytics and guided decision making have as their core functionality the ability to give your company a real crystal ball. The upshot is the better you understand your current situation and what it is likely to be in the future, the better your decisions and actions will be and the more successful your company will be.

Both of these capabilities currently apply to operational or tactical decisions but may soon reach the level of maturity to begin generating strategic decision-making capabilities for your enterprise. In any case, their current abilities are rather impressive. Here are some examples of how they are being used today:

Customer retention and acquisition. In today's economic difficulties, companies are desperate to hold on to profitable customers and attract new ones. HSBC, for example, uses its predictive analytics capabilities to develop appropriate retention and income generation strategies for its customers.¹ Their analysts are able to quickly assess customers and their transactions and apply appropriate transaction fees for all transactions.

Credit card fraud. Brick-and-mortar credit card companies have a good handle on credit card fraud. Internet merchants are far more vulnerable to fraud, losing nearly \$1 billion a year. Now, Internet companies such as ClearCommerce can analyze suspicious transactions in a real-time scenario, yielding more protection from fraud for their customers.²

Targeting sales activities to appropriate customers. Understanding customers, their behaviors and their needs is perhaps the most universal usage of these types of capabilities. With a thorough discovery and understanding of the attributes that constitute customer behaviors, marketers can very easily target their promotions,

campaigns and offers to their profitable customers and promising prospects. Many companies, such as 1-800-Flowers.com and UBS, perform these types of analyses today.^{3,4}

Assessing constantly changing business scenarios. The oil and gas industry is in a continuous state of change due to significant price volatility, technological advances, regulatory changes and the need to discover new exploration areas. Guided analytics are being used to assess promising locations for oil exploration, streamline production processes and optimize distribution logistics in a timely manner.⁵

Better patient healthcare. It is critical for hospitals and caregivers to determine which patients are likely to become high risk. These patients should receive preventive care, thus relieving some of the burden on the healthcare system. It is also important that these caregivers determine which patients need more immediate care than others. Companies such as American Healthways are using their predictive capabilities to improve the quality of life for their patients and reduce the demands on hospitals and insurance providers.⁶

To create these environments, many experts and vendors recommend that, as a first step, you thoroughly study and understand the business processes that drive the need for predictive analytics or guided decision making. By obtaining a clear understanding of the business processes, you also gain an understanding of the expected outcomes long needed by the business community and an understanding of the role played by the application and how its outputs affect other systems or subsequent inputs. Once you have defined the processes and therefore the need for the analytic capability, the rest of the steps are relatively straightforward.

1. You should list any constraints or restrictions that limit the solution. Also, list any assumptions you make and all facts that are known. These should include limitations, assumptions or known facts about the availability or quality of data, resources you have to bring to bear on the problem and technology that will be needed.

2. Create the analytic environment based on your existing BI architecture. If you follow the classic Corporate

Information Factory architecture, the analytic capability becomes a specialized dependent data mart that receives data from the data warehouse at regular intervals. The feeds into the mart are tested, the interfaces established and the users trained on the new capability.

3. Now the application can begin to fulfill its purpose. If it is a predictive analytic application, it will begin analyzing historical records and generating forecasts or predictions of future events. You may choose to generate many possible scenarios (what-if analyses) and then evaluate their usefulness using business-defined criteria. If it is a guided decision making application, the first steps are the same as for the predictive application, with the additional step of creating recommendations and actions to be taken should certain events occur.

4. The analysts using these applications must determine which scenarios to use for the business based on his or her experiences. Once chosen, the applications are interfaced with the operational systems where appropriate.

At this point, a number of checks and balances should be put into place. First, there must be a process created to allow feedback to be incorporated into the analytic applications to indicate whether the forecasts or decisions generated are performing optimally. The ideal is for these to perform the same way an "experienced" decision-maker would have performed. This is the ultimate goal of this environment; thus it is critical to constantly monitor the results from these applications to ensure a successful outcome. Understanding the impact on marketing, risk, fraud or customer serv-


ice decisions and how these influence customer decisions is mandatory feedback into these applications.

The second check and balance to implement is the formation of control groups – customers/prospects/partners/suppliers/employees who receive no special treatment. These control groups allow you to determine the effectiveness as well as continued enhancement of these applications.

I have two final thoughts about the implementation of your predictive capabilities. These applications are only as good as the data and formats of that data being provided. In most BI environments, we tend to think only in terms of multidimensional designs. Unfortunately, star schemas and cubes of data do not lend themselves to this more sophisticated form of analysis. Predictive capabilities require a new way of thinking, new technologies and even new database designs. Depending on whose technology you purchase for your predictive environment, these designs will vary from flat files to normalized designs to floating point designs or others – all using massive amounts of transactional data. Check with your vendor about your design.

Now the data. Perhaps the best way to determine what data is needed is to start with the expected outcomes of the applications and work backwards. For example, if you want to predict customer buying behaviors based on customer segments, then you will need customer demographic, loyalty and behavioral data. Next, you should list the most favorable variables from the data that have the most influence on the preferred customer actions. Once these are known,

you can further hone the data attributes needed to develop the final models. You may want to extract extra pieces of data as well if you have the time and capacity. Even though customer loyalty and behavioral data may not be in your current environment, models that include these customer attributes are much more accurate in predicting customer actions than models containing only demographic or transactional data.⁷

The ability to study history and then foretell the future leads to the obvious conclusion that we should then be in a position to determine the proper course of action based on events occurring today. I think predictive capabilities have a natural interaction with the business activity monitoring (BAM) or business performance management (BPM) capabilities. BAM and BPM both have the ability to send alerts, alarms or other warnings based on triggers. These triggers should be based on events meeting the predicted or predefined criteria generated from our analytic engines. I predict that the vendors in these two camps are already developing interfaces for each other and that soon we will see nicely integrated predictive analyses driving tactical, actionable events. 

References:

1. See case study at www.spss.com.
2. Ibid.
3. See case study at www.sas.com.
4. See case study at www.genalytics.com.
5. See www.spotfire.com for case studies.
6. See case study at www.sas.com.
7. See www.WalkerInformation.com for more information.

Claudia Imhoff, Ph.D., is the president and founder of Intelligent Solutions (www.intelsols.com), a leading consultancy on CRM and business intelligence technologies and strategies. She may be reached at cimhoff@intelsols.com.