Chapter 9

Business Decisions

Chapter Outline

What You Will Learn in This Chapter

- How do businesses make decisions?
- How do you make a good decision? Why do people make bad decisions?
- How do you find and retrieve data to analyze it?
- How can you quickly examine data and view subtotals without hundreds of queries?
- How does a decision support system help you analyze data?
- How do you visualize data that depends on location?
- Is it possible to automate the analysis of data?
- Can information technology be more intelligent?
- How do you create an expert system?
- Can machines be made even smarter?
- What would it take to convince you that a machine is intelligent?
- What are the differences between DSS, ES, and AI systems?
- How can more intelligent systems benefit e-business?
- How can cloud computing be used to analyze data?

Citigroup

How do you use information technology to make better decisions? Citigroup is one of the largest banks in the world. And the world part is important: the company operates in 101 countries. Running a global bank requires making thousands of decisions—from basic questions about approving loans to structuring mega-deals with huge corporations. Banks obviously adopted information technology early—to handle basic transactions. After all, money today is really just numbers in a computer. Banks, including Citigroup, have been slower to adopt technology to make decisions; but that reluctance has been changing in the last few years.

Citigroup is faced with competitive pressures as well as disruptions from economic fluctuations. It has attempted to reduce risk by spreading into multiple business areas, including insurance (Travelers), brokerage (Salomon Smith Barney), and investment banking. The company is one of the largest issuers of credit cards in the United States. With stagnant growth of the U.S. credit industry, Citigroup has expanded into South America and even China. Evaluating customers to minimize risk is a key aspect of the credit card industry.

Introduction

How do businesses make decisions? Figure 9.1 shows that as a manager, you will have access to huge amounts of data. How do you analyze it to understand what it means? How can information systems help you make better decisions? These questions are difficult to answer, but there is a much bigger underlying question. Why would companies need to hire you as a manager? Executives already have access to databases containing integrated data for the entire company. ERP systems can provide detailed data on any aspect of the business. EIS systems can show charts, summaries, and detailed data. Now, what if computer systems can be built to analyze the data and make decisions? What job will you have?

Analyzing data and making decisions depends somewhat on the discipline you choose and your background in statistics and operations management. A key aspect in many introductory business courses is to teach you the basic models used within the discipline. It is unlikely that you are an expert in any particular discipline at this point, and you might not be aware of the variety and power of the statistical tools available to analyze data. Nonetheless, you need to learn how information technology can be applied to many business decisions. This chapter focuses on the tools available and illustrates them with relatively simple decisions. Even if you do not fully understand the discipline-specific models, you need to learn what tools are available.

Can computers make decisions? How can information systems help managers make decisions? Some business problems are straightforward. In these cases, developers simply create a set of rules or procedures that the computer can be programmed to follow. As long as the business behaves in a predictable manner, the rules apply and the computer can handle the details. However, many business problems are less structured and cannot be solved so easily. In addition, problems often involve data that is not well defined. For example, it is straightforward to create a computer system to handle inventory because the computer can easily keep track of item numbers and quantity sold. Consider the more difficult prob-

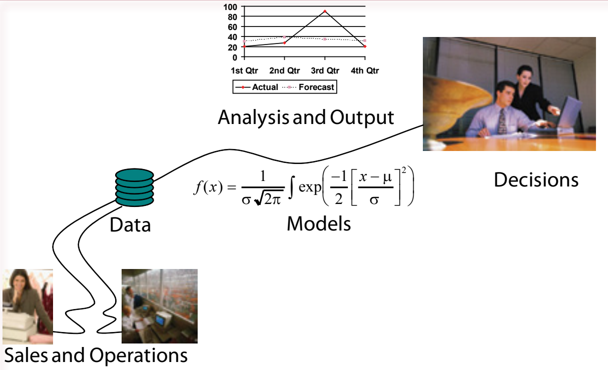

Figure 9.1

Tactical decisions often require complicated analysis. Problems utilize forecasts, optimization, and in-depth analysis. Information systems provide support through data, modeling, and presentation tools. Managers use information system tools to build, evaluate, and maintain various models.

lem faced by a manager who has to decide where to locate a new plant. Some attributes are measurable, such as distance from suppliers, cost of land, and taxes. Other features are difficult to quantify: quality of the labor force, attitudes of government officials, and long-run political stability of the area.

Many problems involve nonnumeric data and complex interrelationships among the various factors. Without computers, businesses often call in experts or hire consultants to help solve these problems. Special software programs called **expert systems (ESs)** provide many of these features. From the beginning, researchers and computer designers have known that humans perform some tasks much better than computers can. These differences led researchers to investigate how people solve problems and investigate how humans think. The research into techniques that might make computers "think" more like humans is known as **artificial intelligence (AI)**. There is some question as to whether it will ever be possible to build machines that can think the same way humans do. Nonetheless, the research has led to some useful tools that perform more complex analysis and can solve difficult problems. These tools attempt to mimic the processes used by humans in a simpler form that can be processed by a computer system.

The answer to the big question is yes, computers can make decisions in some situations. In other cases, the computer is a tool that helps you analyze the data. The level of support provided depends on the type of problem and on your skills as an analyst. As a manager, it is your responsibility to identify decisions that can be handled by machine systems and to recognize when these systems do not work.

It is Hard to Make Good Decisions

How do you make a good decision? Why do people make bad decisions? Most businesses have evolved over time. In many cases, their business processes have been built and patched in response to various changes in the industry. The firms that made better decisions and changes survived, while the others failed. But the existing process might not be the most efficient. Consider the apparently simple process of farming. Farmers feed the animals and then sell

Trends

Through the 1970s, computers were largely used to assist with transaction processing. Support for making decisions was generally limited to the basic reports produced from the data. Computers were too expensive and programming too difficult to be used by every manager. As personal computers became commonplace through the 1980s, managers began transferring data from the corporate central computers to their personal machines. Spreadsheets made it easier to analyze data, evaluate models, and create charts. In the 1990s, networks, improved spreadsheets, and better ties to databases made it possible to build more complex, interactive models and create forecasts.

Along with technology, improvements were made to modeling and analytical tools. Scientific advancements made it possible to add more intelligence to software tools. Data mining systems use statistical tools to semiautomatically evaluate data, searching for important relationships and clustering or classifying groups of data. Expert systems evolved from early work on artificial intelligence. Focusing on narrow domains, these tools encode the rules of an expert to analyze data and suggest solutions. Today, thousands of expert systems are used to improve decisions and provide quick results 24 hours a day.

The study of human brains yielded clues that led to the development of neural networks. Today, neural networks are widely used in pattern matching applications. Humans are good at pattern recognition, and neural networks dramatically improve the ability of machines to perform these tasks.

Writers and researchers have long wondered whether machines can become intelligent. No one is close to an answer yet, but the new technologies mean that today's systems are more intelligent and can handle more complex problems than machines a few years ago. As the range of solvable problems increases, managers need to understand the capabilities and limitations of each method.

them. The hard part is that the animals can be fed and housed many different ways—each with different costs. Should the animals be fed high-protein food that costs more and grows bigger animals faster, or should they be fed simple diets and take more time to mature? In the 1970s and 1980s, experts created software that analyzed these questions from the standpoint of minimizing the cost of feeding the animals. Using optimization methods, they were able to substantially reduce the production costs. But some experts have found it is possible to do even better by focusing on profits across the entire industry chain. The same types of problems apply to any business process, because most companies solve individual problems first because they are easier. As tools become better, it is important to expand the perspective and look at broader decision problems.

Even if you do have a system for making a better decision, you need to convince managers to use it. Many managers distrust new technologies and different answers, because they see an element of risk. A few companies have established a culture that focuses on continual improvement and growth. In these companies, managers are encouraged to explore new ideas and replace the existing processes.

Reality Bytes: Google's Self-Driving Cars

Following the successes of the DARPA Grand Challenges of the late 2000s, some companies have expanded research and development of self-driving cars. Interestingly, Google has been one of the leaders. Perhaps driven by ties to Stanford (which one the second Grand Challenge), and by the importance of mapping in Google. In 2010, Google told the world that it had been using off-the-shelf components to test run self-driving cars a total of 150,000 miles—with almost no human intervention. However, a little know aspect of Google's methodology is that humans first drive the cars around the test area—largely to provide a more specific map of the route, stoplights, school zones, and so on. Chris Urmson, technical lead for the Google project, and a leader from the Carnegie Mellon team that won the 2007 DARPA Urban Challenge, notes that "There are things that right now are a challenge for us. For instance, if most of the world stayed the same but the lanes are shifted—so the physical road didn't move, but for whatever reason, the department of transportation decided we should drive a half lane to the left—that would probably confuse the car today." The challenge lies in unexpected events—getting the software to recognize and react to them using some level of common sense. Ultimately, the tradeoff is an interesting question in ethics: Does a machine have to be perfect? The National Highway Traffic Safety Administration (NHTSA) noted that in 2008 the U.S. experienced 5.8 million car accidents, with 1.6 million of those resulting in personal injuries and 34,000 in deaths. If machine-drivers can cut that rate in half, would we still blame the developers for the remaining accidents?

Adapted from Nick Chambers, "Hands-Off Training: Google's Self-Driving Car Holds Tantalizing Promise, but Major Roadblocks Remain," *Scientific American*, May 23, 2011.

Figure 9.2

Sample decision. Do you invest your money in Company A or Company B? Be careful, it is a trick question.

Figure 9.3

GM 1984 models. Note the body similarities. Even the base color offerings were the same. How could GM throw away its division strategy and offer virtually identical models? Source: Wikipedia.

Human Biases

Assume you have money to invest in the stock market. Someone shows you two companies. As shown in Figure 9.2, Company A's share prices have risen by 2 percent per month for the last year. The other's share price was flat for five months but has increased by 3 percent per month since then. Which stock do you buy? But wait a minute. How can you possibly decide based on the little information you have? It sounds silly, but people make these decisions with minimal data and no logical analysis every single day. When people make decisions this way, the results are going to be inconsistent and dangerous. But it is so much easier to make a snap decision, and so much harder to do the research and complex analyses to make the right decision.

Consider a true example: designing a new automobile. The automobile industry, epitomized by GM in the early 1980s, presents an interesting example of making decisions. GM demonstrated from the 1940s forward that people buy cars based on design and appeal—not simply a collection of features. In particular, GM demonstrated that customers can be heavily persuaded by style and advertising. Assuming you have money, what kind of car would you buy? Sporty, luxurious, flashy, utilitarian, big, small? What color? How many doors? Now ask a few friends or relatives what they would buy. Will all of the answers be the same? Not likely. Now think about the problem from the perspective of an automobile manufacturer such as General Motors. What features are car buyers going to demand in two or three years? This classic marketing problem is difficult to solve. For years, GM used its multiple divisions to create separate identities that appealed to different segments of the population. Designers within each division focused on the preferences and lifestyles of their specific target. Most of that structure fell apart in the early-1980s with the introduction of a completely new line of cars based on the A-body. At that time, the GM divisions introduced a new car model from four main divisions. As shown in Figure 9.3, all four cars (Oldsmobile Cutlass Ciera, Pontiac 6000, Chevrolet Celebrity, and Buick Century) were virtually identical.

Acquisition/Input		
Bias	**Description**	**Example**
Data availability	Ease with which specific instances can be recalled affects judgments of frequency.	People overestimate the risk of dying due to homicides compared to heart disease.
Illusory correlation	Belief that two variables are related when they are not.	Ask any conspiracy buff about the death of JFK.
Data presentation	Order effects.	First (or last) items in a list are given more importance.
Processing		
Inconsistency	Difficulty in being consistent for similar decisions.	Judgments involving selection, such as personnel.
Conservatism	Failure to completely use new information.	Resistance to change.
Stress	Stress causes people to make hasty decisions.	Panic judgments and quick fixes.
Social pressure	Social pressures cause people to alter their decisions and decision-making processes.	Majority opinion can unduly influence everyone else: mob rule.
Output		
Scale effects	The scale on which responses are recorded can affect responses.	Ask a group of people to rate how they feel on a scale from 1 to 10. Ask a similar group to use a scale from 1 to 1,000.
Wishful thinking	Preference for an outcome affects the assessment.	People sometimes place a higher probability on events that they want to happen.
Feedback		
Learning from irrelevant outcomes	People gain unrealistic expectations when they see incomplete or inaccurate data.	In personnel selection you see how good your selection is for candidates you accepted. You do not receive data on candidates you rejected.
Success/failure attributions	Tendency to attribute success to one's skill and failure to chance.	Only taking credit for the successes in your job.

Figure 9.4

Biases in decision making. Without models, people tend to rely on simplistic "rules of thumb" and fall prey to a variety of common mistakes. These errors can be minimized with training and experience in a discipline. They can also be minimized by having computer systems perform much of the initial analysis.

Model Building

Understand the Process
Models force us to define objects and specify relationships. Modeling is a first step in improving the business processes.

Optimization
Models are used to search for the best solutions: Minimizing costs, improving efficiency, increasing profits, and so on.

Prediction
Model parameters can be estimated from prior data. Sample data is used to forecast future changes based on the model.

Simulation
Models are used to examine what might happen if we make changes to the process or to examine relationships in more detail.

Figure 9.5

The four primary reasons for building and using models. Descriptive, graphical, and mathematical models can be used for each of these purposes. However, mathematical models tend to be emphasized for optimization and simulation.

The available colors were even the same—particularly maroon. In effect, GM was assuming that millions of customers all wanted the same car.

Designing cars, or any product, is a difficult decision problem. In good situations, like GM, you have tons of data available. You have sales data, data on competitors, surveys, and focus groups. But, the results are meaningless if you cherry-pick data to match your preconceived ideas. "*Oh look, here are 5,000 people who want larger engines.*" If you search hard enough, you can find data to match any opinion you want; but it will not accurately represent the opinions of the population.

In response to these problems, Barabba and Zaltman, two marketing researchers working with GM, analyzed decision making at General Motors and noticed that several common problems arose. In summary, they found that people are weak at making decisions. For example, people place too much emphasis on recent events, they tend to discard data that does not fit their prior beliefs, they follow rules of thumb instead of statistical analysis, and they choose outcomes based on wishful thinking. As shown in Figure 9.4, all of these problems and more influenced the decisions of designers at GM. In particular, they found that the designers tended to discuss ideas with their bosses in an attempt to identify management preferences that would help get a particular design approved. So cars were designed to the preferences of a few managers, instead of to the needs of customers. Despite attempts to improve, the fiasco eventually forced GM to eliminate the Oldsmobile division. The books written by Barabba and Zaltman discuss even more examples and human biases in decision making.

Before you think that businesses (and GM) have solved the design problem, go look at current designs. In 2008, the *Wall Street Journal* examined cars produced by GM and Ford and found similar overlap in styles within each company. And GM's overall market share continued to decline at least through 2009, when the company was forced to file for bankruptcy protection—when the company sold off or eliminated several other divisions. The main point to remember is that making decisions without a good model and process leads to poor decisions. Sure, you might get lucky for a while (like investors in the 1990s), but ultimately you need a solid decision-making process.

Models

Models are key aspects to any decision; they are simplifications designed to help you understand and analyze a problem. Many of the models you will use in business decisions were created by academics. You will be introduced to many of these models in other business courses. As a manager, you are responsible for knowing that hundreds of models are available to help you make decisions, and for knowing which model best applies to the problem you are facing. Understanding and evaluating models is an important aspect of a business education.

Models often use drawings and pictures to represent the various objects. However, at heart they typically use mathematical equations to represent the process and the various relationships. For example, an operations engineer would model a machine as a mathematical formula that converts raw materials and labor into products. Using equations for each step of the production process, the engineer could search for ways to reorganize production to make it more efficient or to improve quality.

Models are used to help managers make decisions. Most businesses are far too complex for any single person to understand all of the details. Consequently, a variety of models may be created to present a simplified view of the business. In particular, one of the original purposes of accounting was to create a standardized model of the financial aspects of business. Another common model of business is the practice of dividing the company into functional areas. For example, a manager with experience in the finance department of one company can usually apply knowledge and problem-solving skills to finance departments in other companies and even other industries. The basic goals are summarized in Figure 9.5. Models help you simplify the world. They help you search for similarities in different situations. Models also enable managers to predict how changes might affect the business. As the decision maker, it is up to you to determine which models to use and to make sure they actually apply to the situation. Once you have selected the appropriate model, you apply whatever data you have, evaluate the results, and make the decision.

Prediction and Optimization

An important use of models is for **prediction**. If a model is reasonably accurate, it can be used to predict future outcomes. For instance, when you buy a car, you might want to know how much it will be worth in three years when you want to sell it. It is possible to estimate how the price of used cars changes over time.

Prediction first requires that you have a model that describes the situation. Then data is collected and statistical techniques are used to estimate the **parameters** of the model for the specific problem. Next you fill in values for any parameters that you already know, and the model provides a prediction. Prediction techniques such as regression and time series forecasting are used to examine the data, identify trends, and predict possible future changes. To use statistics effectively requires a model of the underlying system. For instance, to use regression methods you

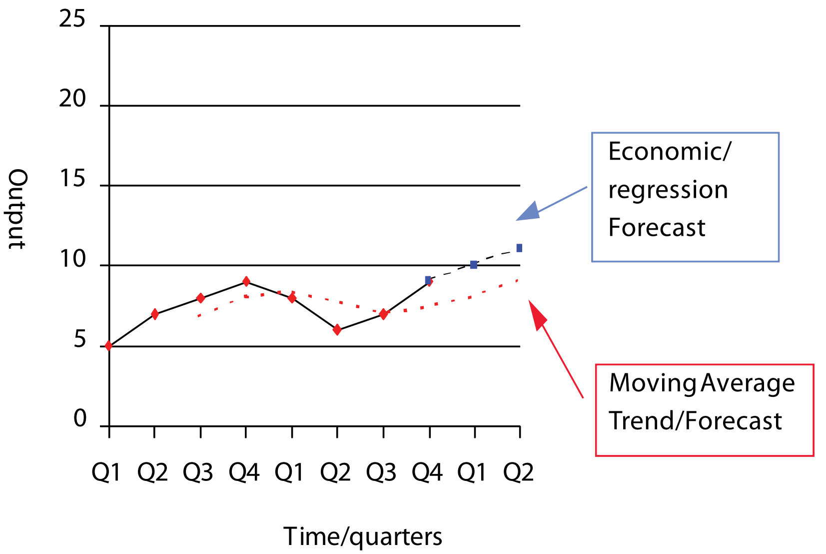

Figure 9.6

Prediction model. Several statistical techniques exist for analyzing data and making forecasts. Two common methods are regression and moving averages. Both methods require substantial amounts of data. Choosing between the two requires some expertise in statistical analysis, but many times we display both methods to show a range of possible outcomes.

first identify the dependent variable and a set of possible independent variables. These choices come from the underlying model.

Figure 9.6 illustrates how a spreadsheet can be used to display the results of a forecast. Although dedicated statistical packages contain more options, spreadsheets contain several tools to help with basic statistical regression and forecasting. Once you load the Data Analysis toolpak in Excel, you can run basic regressions, T-Tests, or ANOVA to identify how a set of independent X variable affect the dependent Y variable. You can use the trend line function to highlight averages on charts. If you are familiar with statistical functions, you can use the built-in functions to make comparisons and evaluate data. You can also purchase add-ins to perform more-sophisticated analysis of your data. When you need even more power, you can use the import and export tools in the high-end statistical software packages to extract and analyze your data. You might have to hire a statistician to help determine which tools you need and to create the initial models. But once the model and system are configured, you can run the system to analyze your results on a regular basis.

Optimization

Optimization evaluates a model in terms of the inputs or control parameters and searches for the best solution. Optimization requires a detailed mathematical model. Several tools such as linear programming are used to find optimal values. Some optimization models have resulted in substantial savings in cost or increases in profit. Tasks that are repeated hundreds or thousands of times can often benefit through optimization modeling.

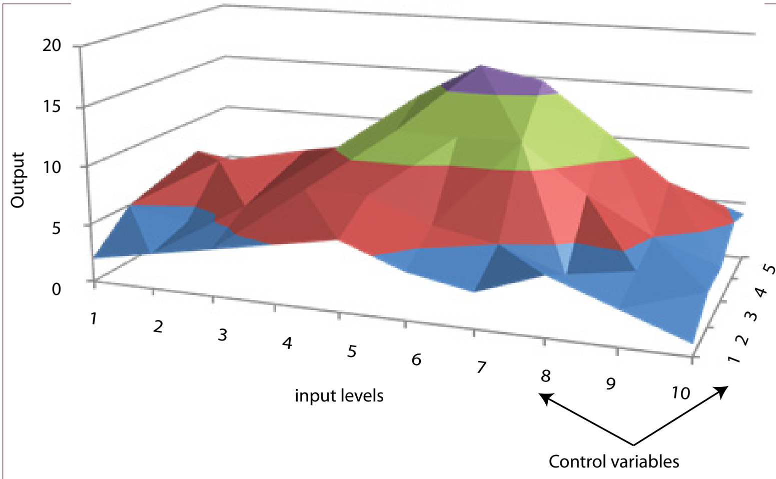

Figure 9.7

Optimization model. Optimization requires a mathematical model. Output is defined in terms of various input variables. The maximum point can be found by altering the values of the input variables. Special software, such as linear programming tools, can help evaluate complex models.

Figure 9.7 shows a simple optimization model with two input variables and a single output. Optimization typically requires a mathematical model that describes the complex relationships between the variables. You will often use statistical techniques to estimate the underlying model parameters, and then turn to optimization tools to find the best operating point. Linear programming is an optimization tool that has demonstrated considerable success in solving some relatively complex problems. As the name implies, basic tools assume that all of the relationships are linear, but modern tools can support a limited amount of nonlinearity, including quadratic terms.

The Solver tool that is sold with Excel is a limited version of a commercial optimization product, but even the basic version is capable of solving moderately-large problems. If your model grows too large, you can upgrade to the full version or purchase software from other companies. One of the strengths of the linear programming (and Solver) approach is that it is designed for constrained optimization problems—which are common in business. Operations research and management classes explain these problems in more detail. The problems are a little tricky to set up, so you should consult an expert if you lack experience with these problems. The key is to identify the output goal (perhaps profit maximization or cost minimization), and then list all of the input variables that affect that result. Once you have defined the optimization equation, you need to specify all of the constraints that affect the problem. For example, a production plant has a maximum amount that can be produced; and people work a limited number of hours a week producing a defined amount of output per person and machine. Once you understand optimization, you will recognize that it can be used to solve a variety of problems, including complex pricing decisions, worker scheduling, product mix, and input selection.

Reality Bytes: Make Decisions by Putting Down the Cell Phone and Going for a Walk

If you think choosing what to eat for dinner is tough, stay away from combinatorial auctions. Consider an airline that wants to bid on landing slots at an airport. The slots can be bid on individually or in bundle of combinations. Data includes the passenger loads, weather, and connecting flights. Angelika Dimoka, director of the Center for Neural Decision Making at Temple University, uses these types of complex puzzles to study the ability of humans to make decisions. She used functional MRI (fMRI) to examine brain patterns while volunteers worked on these puzzles. At the start, increasing the amount of data caused activity in the dorsolateral prefrontal cortex region of the brain to increase. This region is responsible for decision making and control of emotions. However, as the amount of data increased, activity suddenly dropped as the humans reached overload. At that point, frustration and anxiety soar as the emotions take over. Other research into human aspects of decision making show that unconscious systems guide human decisions and that decisions that require creativity benefit from patience—allowing the brain to think about the problem passively. But interrupting the brain with increasing amounts trivia makes it more difficult for the brain to operate. It is one of the reasons pilots and first responders train for emergencies and maintain checklists of critical tasks. Studies, such as those by Shenna Iyengar of Columbia University, show that providing more choices and more data can cause people to give up and simply opt out. For instance, employees who faced retirement plans with more options tended to opt out more often than those with fewer choices. The brain needs time to evaluate data, integrate the elements, and determine which items to ignore. So, collect the data, use analytical tools and models that were developed for the problems, try to identify the most important factors, then take a break and let the brain evaluate the choices for a while.

Adapted from Sharon Begley, "I Can't Think!" *Newsweek*, March 7, 2011.

Simulation or "What-If" Scenarios

Simulation is a modeling technique with many uses. Once a model is created, managers use simulation to examine how the item being studied will respond to changes. With simulation, various options can be tested on the model to examine what might happen. For example, engineers always build models of airplanes and engines before they try to build the real thing. The models are much cheaper. In fact, most engineers today start with mathematical computer models because they are cheaper to create than physical models and can contain more detail. Moreover, they can perform experiments on models that would not be safe to perform in real life. For example, an engineer could stress a model of an airplane until it broke up. It would be dangerous and expensive to try such an experiment on a real plane. Similarly, a business model could examine what would happen if prices were increased by 20 percent, without worrying about losing real money. Of course, the model is only useful and valuable if it accurately reflects real conditions.

Most simulation models are mathematical instead of descriptive models, because they are easy to evaluate. Mathematical models contain parameters, or variables that can be controlled by the managers. For instance, you might use a spreadsheet to create an accounting model of an income statement and balance

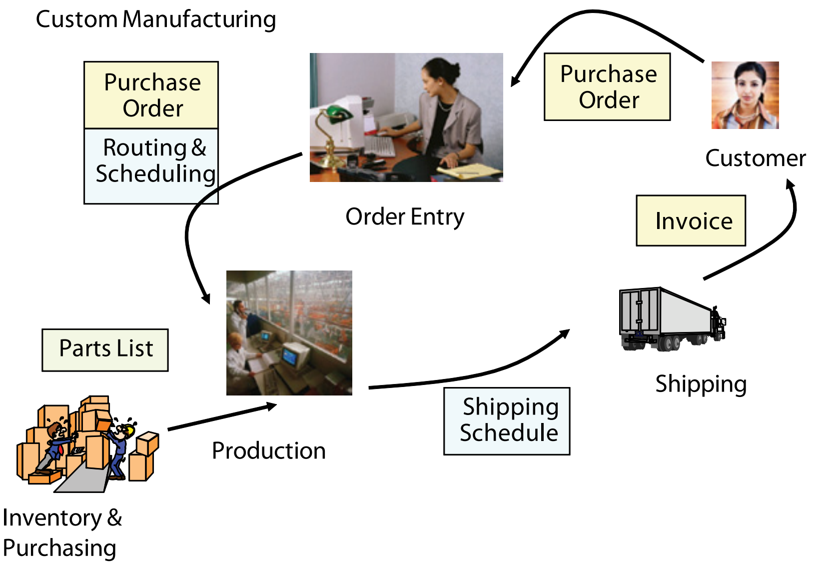

Figure 9.8

Object-oriented simulation. A simple example of custom manufacturing. Each object (parts list, purchase order, etc.) and each process are defined in detail by the modeler. The simulation system generates orders, makes shipments, and orders inventory according to programmed rules. The simulator collects a variety of statistics that can be displayed graphically or in summary reports.

sheet. When you create the spreadsheet, production quantity and price of your products are controllable parameters that affect the income and profits of the firm. You could use the model to investigate decisions, like the effect on profits if you increase production. Costs will increase, but so will revenue from sales. The net result depends on the specific details of the firm and the model. Spreadsheets are often used to analyze small models and graph the results. More sophisticated simulation packages can support more complex analysis and will automatically create graphs and pictures to show interrelationships.

More complex models provide more opportunities for simulation. In part because they are difficult to solve and understand. In a production system, a more detailed model might enable you to investigate alternatives such as increased overtime, hiring another shift, building additional plants, or subcontracting the work to another firm.

Object-oriented simulation tools developed in the last few years make it easy to create many simulations. As shown in Figure 9.8, you can place icons on the screen to represent various business objects. Behind each of these objects, you specify their behavior. For example, you would need to estimate the average number of customers that arrive in an hour and how long they are willing to wait for service. For physical processes like order entry, inventory, and shipping, you specify the number of transactions that can be handled in a given time. When all of the details have been entered, the system runs simulations and tracks statistics. You can change the parameters such as adding clerks to see what happens if you change the operations.

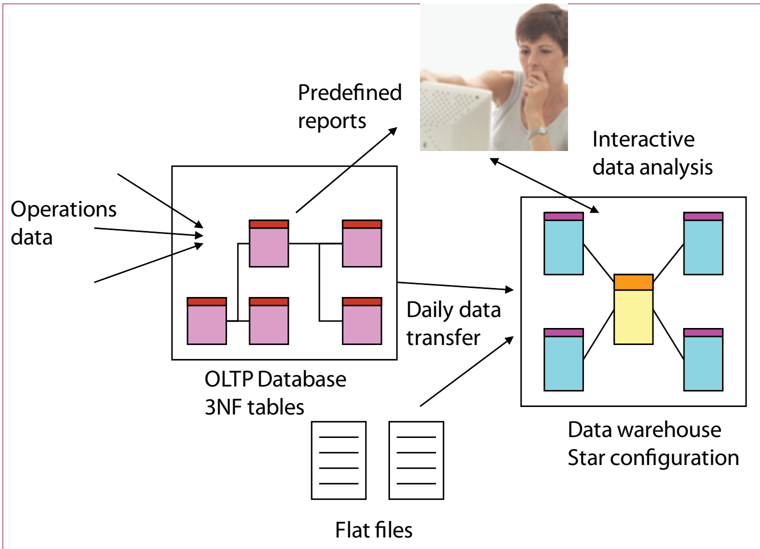

Figure 9.9

Data warehouse. A data warehouse is commonly used as a method to provide data to decision makers without interfering with the transaction-processing operations. Selected data items are regularly pulled from the transaction data files and stored in a central location. DSS tools query the data warehouse for analysis and reporting.

Data Warehouse

How do you find and retrieve data to analyze it? ERP and other transaction systems can provide enough data to bury you. But the transaction systems are designed to store data, not to search and analyze it. Relational databases turn out to be relatively slow when you need to analyze several gigabytes or even terabytes of data. If you do not have an ERP system, you have even greater problems trying to integrate and clean data from all of your systems. As shown in Figure 9.9, the answer is to create a separate **data warehouse** that extracts and stores the data in a clean, easy-to-analyze format. The process shown in Figure 9.9 is known as **extraction, transformation, and loading (ETL)**. Larger database management systems have specific tools and data storage methods to create data warehouses. Some companies also create specific **data marts** that are basically copies of a small portion of the data warehouse designed to feed a specific application. For instance, a financial data mart might be used by the accounting and finance department just to monitor investments and bank accounts.

Finding and cleaning the data is the most time-consuming step in most data analysis projects. Data is often stored in multiple systems and proprietary formats. Even if the data is stored in a DBMS, someone has to verify the data content and format to ensure it matches exactly with the other data. Even a number as simple as sales could be defined differently by various departments. Some departments might count sales when a customer signs a contract, but others might wait until the contract has been approved by top management. Sometimes organizations are not even aware that these differences exist—until they decide to start integrating the data.

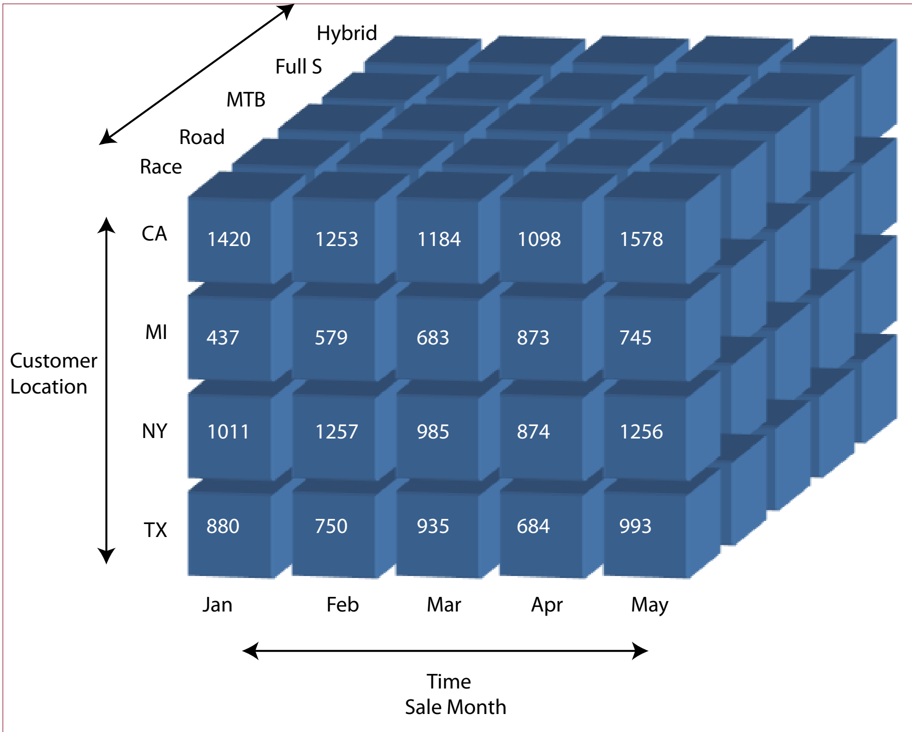

Figure 9.10

Multidimensional cube for item sales. Managers are interested in various combinations of the dimensions. For example, total item sales of dog items in the last quarter. OLAP tools rapidly provide answers to questions involving any perspective of this cube.

Documenting the data is critical because managers have to understand what each item represents and they need to be able to find specific items. **Metadata** is used to describe the source data, identify the transformation and integration steps, and define the way the data warehouse is organized. A data warehouse represents a subset of the total data in the company. In most cases, it is a static copy that is refreshed on a daily or hourly basis. This type of system is relatively easy to use; managers do not have to learn data access commands (SQL or QBE). However, it is less flexible than using a database management system. Decision makers will be unable to get additional data or to compare the data in some previously unexpected way. The success of a data warehouse depends on how well the manager's needs have been anticipated.

Online Analytical Processing (OLAP)

How can you quickly examine data and view subtotals without writing hundreds of queries? Retrieving data and examining totals from different perspectives is an important part of making decisions. When the problem is unstructured and there is no existing mathematical model, it often helps to look at subtotals of data. Sometimes it is useful to browse through the data and examine various subtotals. Which customers bought the most? Click: there are the totals. Which employees sold the most? In December? Click: see the totals

Reality Bytes: Face Recognition

The general public is easily misled by science fiction and Hollywood. It is easy to think that computers can do anything, or that breakthroughs are only a few months away. But, computers still struggle at some tasks that humans perform easily. Most of these tasks fall into the category of pattern recognition, and face recognition is just the latest example. Many photo sites use facial recognition technology to help users identify new photos. For example, Google's Picasa automatically tries to tag new photos by matching faces to existing photos, which simplifies organization and makes it easier to share them. But the system often fails. One user, Stacey Schlittenhard, noted that "All babies kind of look alike—they all have little round faces. If I label one baby as my son, it will label almost every baby as my son." She also noted that Picasa once labeled a lollipop as her friend. Yi Ma, an associate professor of electrical and computer engineering at the University of Illinois, and a researcher at Microsoft Research in China, notes that "I don't think, currently, any facial recognition system is good enough for security purposes—not even close, actually. He also believes a reliable system is at least a decade away. On the other hand, in June 2011, when Vancouver fans rioted after losing the last round of the NHL championship, several photos of people were published on the Internet. Within days, people on the Web had identified most of the rioters and two people who were caught passionately kissing in the middle of the street. Similarly, someone posted a video of a woman arguing with a subway conductor in New York and she was quickly identified by people who watched the video. In that sense, computerized facial recognition is not the threat to privacy—the prevalence of cameras and the ability to share them on the Web is removing anonymity in everyday life.

Adapted from John D. Sutter, "Why Face Recognition Isn't Scary – Yet," *CNN Online*, July 9, 2010; and Brian Stelter, "Upending Anonymity, These Days the Web Unmasks Everyone," *The New York Times*, June 20, 2011.

or view a chart. This process is a major function of **online analytical processing (OLAP)**. Most OLAP tools rely on a data warehouse to provide consistent data and fast access.

The ability to browse through the data is a useful feature of OLAP. Most decision makers want to see subtotals of data. A key method of organizing OLAP data is to identify a key fact (often sales) and then collect attributes that presumably affect the fact. As shown in Figure 9.10, OLAP tools depict this data as a multidimensional cube. Managers use specific tools to examine various sections of the data. To illustrate the process, consider a simple example from the Rolling Thunder Bicycle company database. Managers are interested in sales of bikes. In particular, they want to look at sales by date, by the model type (race, road, mountain, full suspension, and hybrid), and by the location of the customer. The fact they want to measure is the value or amount of the items sold, which is the price times the quantity. The OLAP tools enable managers to examine any question that involves the dimensions of the cube. For instance, they can quickly examine totals by state, city, month, or category. They can look at subtotals for the different categories of products or details within individual states. The tools can provide detail items that can be pictured as a slice of the cube, or they can provide subtotals of any section. Each attribute represents a potential subtotal. In terms

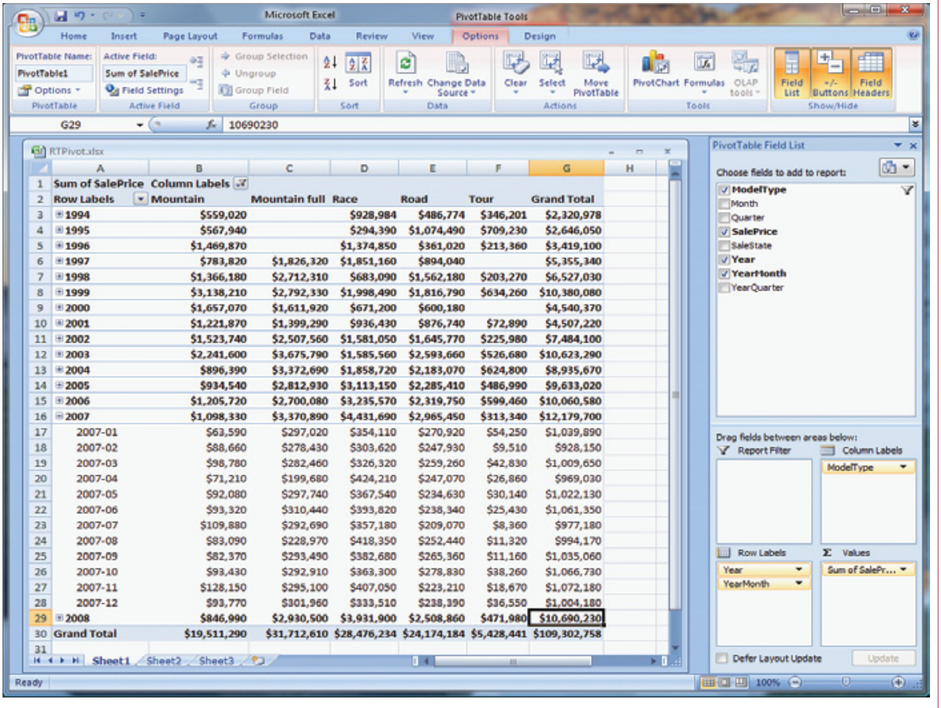

Sum of SalePrice	Column Labels					
Row Labels	Mountain	Mountain full	Race	Road	Tour	Grand Total
1994	$559,020		$928,984	$486,774	$346,201	$2,320,978
1995	$567,940		$294,390	$1,074,490	$709,230	$2,646,050
1996	$1,469,870		$1,374,850	$361,020	$213,360	$3,419,100
1997	$783,820	$1,826,320	$1,851,160	$894,040		$5,355,340
1998	$1,366,180	$2,712,310	$683,090	$1,562,180	$203,270	$6,527,030
1999	$3,138,210	$2,792,330	$1,998,490	$1,816,790	$634,260	$10,380,080
2000	$1,657,070	$1,611,920	$671,200	$600,180		$4,540,370
2001	$1,221,870	$1,399,290	$936,430	$876,740	$72,890	$4,507,220
2002	$1,523,740	$2,507,560	$1,581,050	$1,645,770	$225,980	$7,484,100
2003	$2,241,600	$3,675,790	$1,585,560	$2,593,660	$526,680	$10,623,290
2004	$896,390	$3,372,690	$1,858,720	$2,183,070	$624,800	$8,935,670
2005	$934,540	$2,812,930	$3,113,150	$2,285,410	$486,990	$9,633,020
2006	$1,205,720	$2,700,080	$3,235,570	$2,319,750	$599,460	$10,060,580
2007	$1,098,330	$3,370,890	$4,431,690	$2,965,450	$313,340	$12,179,700
2007-01	$63,590	$297,020	$354,110	$270,920	$54,250	$1,039,890
2007-02	$88,660	$278,430	$303,620	$247,930	$9,510	$928,150
2007-03	$98,780	$282,460	$326,320	$259,260	$42,830	$1,009,650
2007-04	$71,210	$199,680	$424,210	$247,070	$26,860	$969,030
2007-05	$92,080	$297,740	$367,540	$234,630	$30,140	$1,022,130
2007-06	$93,320	$310,440	$393,820	$238,340	$25,430	$1,061,350
2007-07	$109,880	$292,690	$357,180	$209,070	$8,360	$977,180
2007-08	$83,090	$228,970	$418,350	$252,440	$11,320	$994,170
2007-09	$82,370	$293,490	$382,680	$265,360	$11,160	$1,035,060
2007-10	$93,430	$292,910	$363,300	$278,830	$38,260	$1,066,730
2007-11	$128,150	$295,100	$407,050	$223,210	$18,670	$1,072,180
2007-12	$93,770	$301,960	$333,510	$238,390	$36,550	$1,004,180
2008	$846,990	$2,930,500	$3,931,900	$2,508,860	$471,980	$10,690,230
Grand Total	$19,511,290	$31,712,610	$28,476,234	$24,174,184	$5,428,441	$109,302,758

Figure 9.11

Microsoft PivotTable report. Pivot tools make it easy for managers to examine cube data from any perspective, to select subsets of the data, to perform calculations, and to create charts.

of database queries, think of an OLAP cube as a collection of GROUP BY statements. However, the browsing tools are interactive and show managers the results instantly. For example, you can use checkboxes to select specific states. You can begin looking at total sales of all bikes, and then look at the detail by model type. With a few clicks, you can choose a specific slice of the cube, such as one state, and then look at data for a single month or model type, or rotate the cube and look at totals by month.

Although most DBMS vendors (including Microsoft, Oracle, and IBM) provide OLAP cube browsers, Microsoft provides the PivotTable interface that works with almost any DBMS or spreadsheet. A PivotTable is an interactive interface to a multidimensional cube. A PivotTable is created on the user's machine—most users will build pivot tables inside of Microsoft Excel. This tool has several options and provides a great deal of flexibility for the user.

A PivotTable report for Rolling Thunder Bicycles is shown in Figure 9.11. By clicking on a row or column dimension, managers can see detail or subtotals. They can also select specific items to include in the subtotals. Managers even have the flexibility to drag the dimensions around—to move them from columns to rows, or to change the order of the summations. The four windows in the lower-right-hand corner specify where each item will be displayed, and you can drag-and-drop the columns across those windows. Additional options provide other statistics, such as averages.

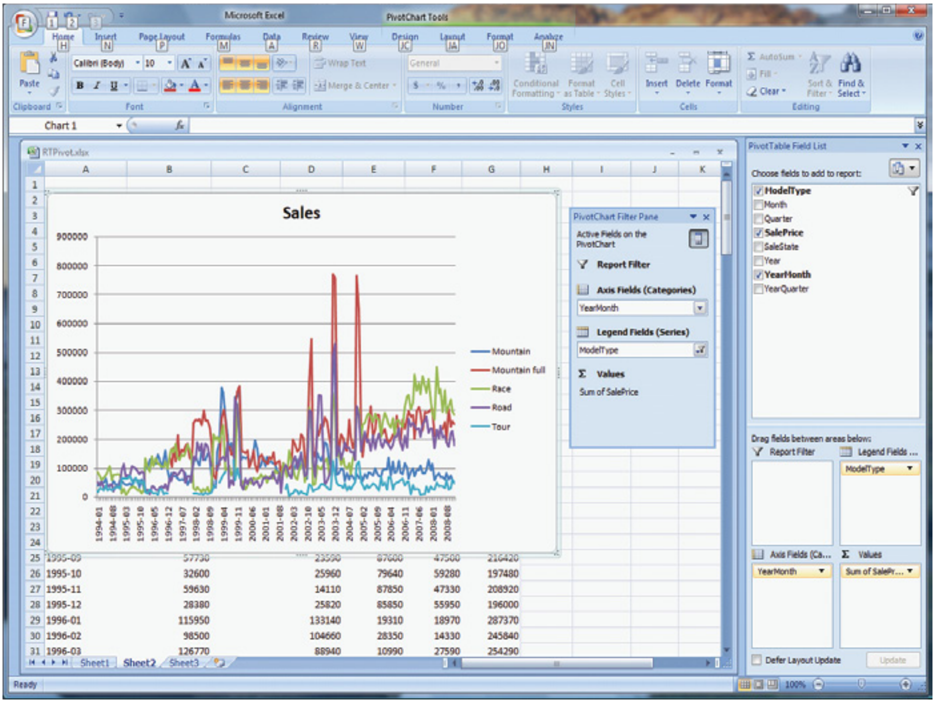

Figure 9.12

Microsoft Pivot Chart. A Pivot Chart uses the same type of data as a PivotTable, but displays the data graphically. It is easy to change the type of chart or to change the data displayed by clicking options or dragging attributes to different sections.

The Pivot Chart graphics option makes it easy to create charts—using the Excel interface that is familiar to most business managers. As shown in Figure 9.12, the system uses the same data as the PivotTable, but creates interactive charts to display the data visually. Once the connection to the data is made, you select the basic chart type and pick the attributes to be displayed. You, or your manager, can change the chart type easily, and even drag data columns to look at new relationships. The options are similar to those in the standard Excel charting package. But remember that the system is interactive, and you can quickly look at any relationship of interest. The data connection is also dynamic, so that as the data changes in the underlying database, you can refresh the chart or PivotTable to get up-to-date values without altering the spreadsheet or chart design.

Several vendors provide interactive cube browsers. Microsoft and Oracle have versions built into their DBMS reporting systems so that you can embed them on a Web report—giving managers Web access to these tools. The PivotTables and Charts built into Excel are convenient and accessible to most organizations without spending additional money. Even if your data is stored in other Excel spreadsheets, you can build PivotTables and Charts that extract the data from other spreadsheets and provide the same interactive data browsing. Additionally, you can use the statistics, functions, and programming capabilities of Excel to analyze the extracted data.

Internal	Purchase	Government
• Sales • Warranty cards • Customer service lines • Coupons • Surveys • Focus groups	• Scanner data • Competitive market analysis • Mailing and phone lists • Subscriber lists • Rating services (e.g., Arbitron) • Shipping, especially foreign	• Census Income Demographics Regional data • Legal registration Drivers license Marriage Housing/construction

Figure 9.13

Common marketing data sources. There are three primary sources of marketing data: internal collections, specialty research companies, and government agencies. Detailed data is available on the industry, customers, regions, and competitors.

Decision Support System

How does a decision support system help you analyze data? Once you have the data, you often need to analyze it statistically, or by applying discipline-specific models. A **decision support system (DSS)** consists of three basic components to help you analyze data: (1) data retrieval, (2) model evaluation, and (3) visualization of the results. Today, a data warehouse or OLAP cube is often used as the data source. The model is often developed by experts (usually consultants) and evaluated in a spreadsheet. The visualization component generally consists of charts, but more sophisticated time lines and schedules are used for complex problems.

To understand the value of a DSS, it is easiest to work with a couple of examples. Thousands of examples exist, but they often require detailed knowledge from a specialized discipline. On the other hand, most managers will need some familiarity with marketing and with human resources management. The examples are relatively small and you will cover more complex models in other business classes. The principles are the same, and you can often use the same tools. Along the same lines, if you want to apply these models to real business problems, you will have to collect more data and add more features to the models.

Marketing Forecasts

Marketing departments are responsible for market research, sales forecasting, management of the sales staff, advertising, and promotion. In some firms they also process orders and manage the design of new products and features. Processing orders is essentially a transaction-processing task. The others involve tactical or strategic questions that are more complex, so we will focus on those tasks.

An enormous amount of data is available for market research. Figure 9.13 presents some of the common data available for marketing purposes. Internally, the marketing department maintains records of sales and basic customer attributes. With some firms, there can be a longer distance between the firm and the final customer. For instance, manufacturers typically sell products to wholesalers, who place the products in individual stores, where they reach the final customer. In these cases it is more difficult to identify customer needs and forecast sales. There

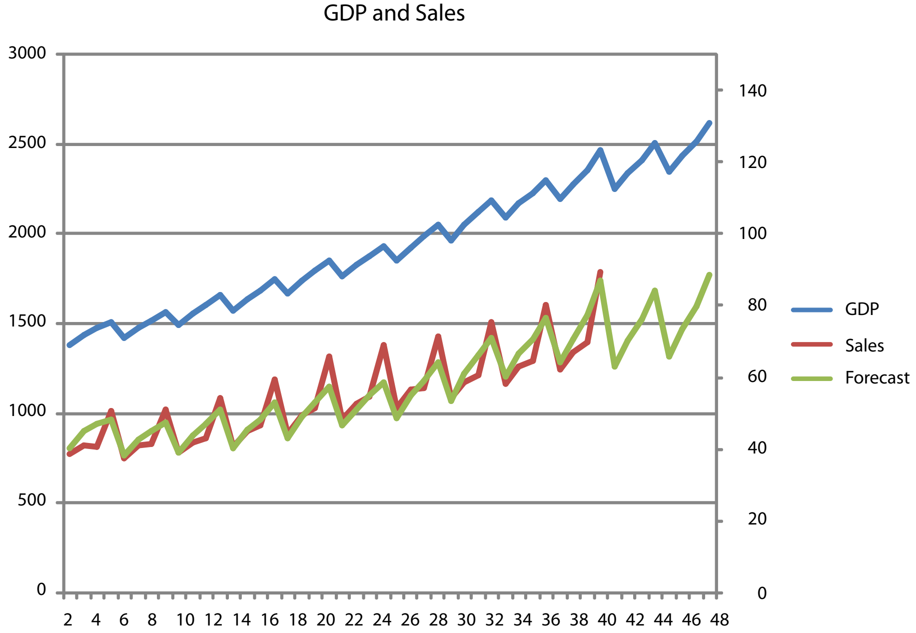

Figure 9.14

Sales forecast. Note the seasonal peaks in the fourth quarter. The points beyond quarter 40 are forecasts based on time and the relationship to income (GDP). This forecast requires that GDP predictions be made for each future quarter, but these values can often be obtained from government forecasts.

will be delays in receiving sales data because the retailers and wholesalers typically place bulk orders. Furthermore, it is more difficult to identify customer preferences because their purchases are filtered through other companies. Marketing departments also have access to data that is collected by other firms. In a manufacturing environment, marketers might get raw sales data from the wholesalers and retailers. On the retail side, with the pervasiveness of checkout scanners, it is now possible to buy daily and hourly sales records from thousands of stores in various cities. This data contains sales of your products as well as rivals' products.

Marketing is often asked to forecast sales. Several different methods can be used, but a straightforward approach is to begin with statistical forecasts. Consider a simple example shown in Figure 9.14 for a fictional store that sells consumer products nationwide. The sales estimate is based on economics where sales are dependent on time and on consumers' income. As consumer income increases, they will be more likely to purchase the company's merchandise. Gross domestic product (GDP) is often used as a proxy for consumer income.

Notice the seasonal peaks in sales for each fourth quarter. It is important to capture this holiday sales effect. Since national GDP and household income have this same effect, you can build a model based on the relationship of your sales to GDP and time. The process is described in Figure 9.15, and all of the steps can be performed in a spreadsheet. Regression provides the estimate of the coefficients that describe the relationship between sales and GDP and time. The R-squared value is over 90 percent and all of the coefficients have high t-values, so it is a strong relationship. If you look closely at the forecast values for the fourth quarters, you will notice some errors that indicate the model is missing some elements.

Data:Quarterly Sales and GDP for 16 years.

Model: Sales = b0 + b1 Time + b2 GDP

Analysis: Estimate model coefficients with regression

	Coefficients	Std. Error	T-Stat
Intercept	-48.887	13.631	-3.586
Time	-0.941	0.294	-3.197
GDP	0.067	0.011	6.326

Output: Compute Sales prediction. Graph forecast.

Figure 9.15

Forecasting process. All of these steps can be performed by a spreadsheet. To forecast the quarterly GDP values, simply split them into four columns, one for each quarter. Then use the spreadsheet's linear forecast tool to extend the columns.

Forecasting the GDP is a little tricky, so most people just use government forecasts. But be careful to get nonseasonally adjusted values (which are difficult to find), so they show the quarterly cycle. Unfortunately, the U.S. Bureau of Economic Analysis has been reporting only seasonally adjusted data since 2004. You can also forecast the quarterly values yourself. The easy way to preserve the cycle is to forecast the quarters independently (all first quarters, all second quarters, and so on). Then plug these values in for the forecast, multiply by the estimated coefficients, and graph the result. It might not be quite as accurate as a full time-series estimation technique, but any business student can make it work, and it is better than guessing or wishful thinking.

Human Resources Management

An important HRM task in any organization is the need to allocate raises. Using a merit pay system, each employee is evaluated on the basis of factors related to his or her job. Typically, each manager is given a fixed amount of money to allocate among the employees. The goal is to distribute the money relative to the performance appraisals, provide sufficient incentives to retain employees, and meet equal employment opportunity guidelines. Many of these goals are conflicting, especially with a finite amount of money available. To set the actual raises, managers need to examine the raw data. On the other hand, a graph makes it easier to compare the various goals.

A few specialized software packages can help you determine merit raises. However, as shown in Figure 9.16, it is possible to create a small system using a spreadsheet. A spreadsheet that can display a graph alongside the data tables is particularly useful. Assume that the company wishes to give a certain portion of the raise based on the average performance ratings. The amount of money per point (currently $100) can be changed. Each person can be given an additional market adjustment raise. The total departmental raises cannot exceed the allocated total ($10,000).

				Merit Pay				raise pool		10000				
	perf.			pct	salary range (000)			current	merit	market	total		New	
Name	R1	R2	R3	perf	low	high	avg	salary	100	adjust.	raise	raise pct	salary	Pct Range
Caulkins	9	7	6	73%	28.4	37.5	36.4	35.8	733	800	1533	4.3%	37.3	0.98
Jihong	3	6	7	53%	15.4	18.9	16.3	17.9	533	100	633	3.5%	18.5	0.90
Louganis	8	7	7	73%	26.7	30.2	28.9	29.5	733	850	1583	5.4%	31.1	1.25
Naber	9	8	8	83%	19.5	23.2	21.4	19.8	833	1030	1863	9.4%	21.7	0.58
Spitz	3	4	3	33%	17.3	22.4	18.4	17.5	333	600	933	5.3%	18.4	0.22
Weissmuller	5	4	6	50%	32.5	60.4	45.2	53.2	500	2955	3455	6.5%	56.7	0.87
Department	6.2	6.0	6.2		23.3	32.1	23.8	21.7	3665	6335	10000	5.7%	30.6	0.83
Corporate	5.0	6.0	5.0		124.3	124.3	18.9	18.9		Remain	0			

Figure 9.16

Merit pay analysis. With a merit system, salary increases should be related to performance evaluations (denoted r1, r2, r3). Managers are typically given a fixed pool of money to distribute among the employees. Employee raises should be based on merit evaluations, current salary, the salary range for the job. Market adjustments are often paid to attract workers in high-demand fields. A spreadsheet can be used to model the effects of various policies. In this example, the manager has allocated $100 for each merit percentage point. The rest of the money will be given as market adjustments. The effects of the adjustments can be seen in the graph displayed in the next figure.

The goal is to fill in the market adjustment column so that the raises match the performance appraisals. As illustrated by the graph in Figure 9.17, the manager can evaluate both absolute dollar raise or the percentage increase. The total departmental raises should be equal to $10,000. By displaying the graph next to the last columns in the spreadsheet, it is possible to watch the changes as you enter the data. This immediate feedback makes it easier to set the raises you prefer. Use of some type of DSS analytical system is helpful for identifying and minimizing illegal discrimination in salaries.

As shown in Figure 9.16, you can use the Office 2010 conditional formatting feature to create a simple chart to show where each person's salary falls within the defined range for the job category. The additional chart is somewhat difficult to read, but it does provide useful information. Also, if the system has access to data across the organization, it is possible to statistically analyze the raises assigned by each manager to ensure raises are given in a non-discriminatory fashion. Such a system will probably be unlikely to catch small deviations, but if systematic differences across managers or individuals appear, HR managers can investigate further. The main point of this example is that managers can visually see the effects of their decisions.

Many, many other problems can be analyzed with decision support systems. The key to a DSS is that you need a model to analyze. Typically, an expert will create the model and build a DSS that loads current data. Managers still need to understand the model and the results, because they are the ones who ultimately analyze the data and make the decisions.

Geographical Information Systems

How do you visualize data that depends on location? Many aspects of business can benefit by modeling problems as geographical relationships.

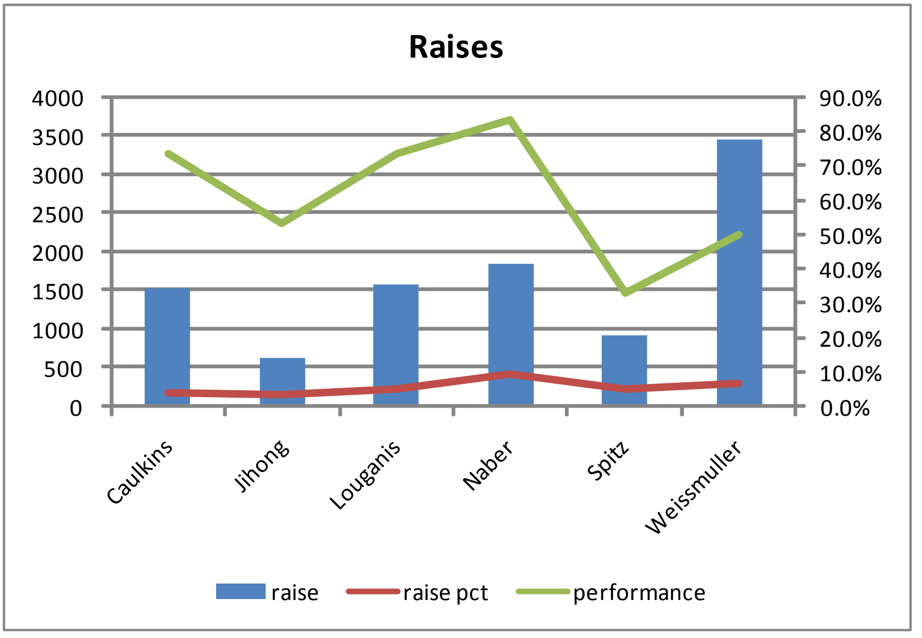

Figure 9.17

Performance evaluation. Using a separate y-axis for the two types of data and overlaying line plots on the bar chart makes this graph easier to read. If this graph is dynamically linked to the salary table, the manager can make salary changes and instantly compare the raises to the performance ratings.

For instance, to choose the site of retail outlets, you need to know the location and travel patterns of your potential customers as well as the locations of your competitors. Manufacturing can be made more efficient if you know the locations of raw materials, suppliers, and workers. Similarly, locations of distribution warehouses need to be chosen based on retail outlets, manufacturing facilities, and transportation routes. Thousands of other geographical considerations exist in business, such as monitoring pollution discharges, routing and tracking delivery vehicles, classifying areas for risk of crimes and fire, following weather patterns, or tracing migration paths of fish for commercial fishing. **Geographic information systems (GISs)** have been designed to identify and display relationships between business data and locations. Arc Info and Microsoft's MapPoint are two of many commercial GIS packages available.

A GIS begins with the capability of drawing a map of the area in which you are interested. It might be a world or national map that displays political boundaries. It might be a regional map that emphasizes the various transportation routes or utility lines. It might be a local map that displays roads or even buildings. An oil exploration company might use a map that displays three-dimensional features of a small area. A shipping company could use ocean maps that display three-dimensional images of the ocean passageways. The level of detail depends on the problem you wish to solve.

Maps and Location Data

Most maps today are built from digitized map data. Each item is stored as a geographic point or line segment; defined by latitude, longitude, and altitude. The maps are drawn mathematically from a collection of these points. Consequently, the maps can be drawn to any scale and the tools make it easy to zoom in or out at will. Most U.S. digital maps are based on data that the Bureau of the Census created for the 1990 national census, known as TIGER. The Bureau of the Census has every road and house number entered into a giant database. Because of privacy reasons, they will not sell house locations, but you can get the range of street numbers for each city block. The U.S. Department of Defense has digital data available for many areas, including international locations, and often includes elevation data. The U.S. Geological Survey topographical maps are also being converted to digital systems. However, keep in mind that the systems being mapped are constantly changing, so even digital maps often contain missing, incomplete, or inaccurate data—as the United States learned when it accidentally blew up the Chinese embassy in Belgrade because the CIA maps were out of date. Due to the popularity of online mapping systems, several private companies work to add new data and fix errors.

Several companies have integrated satellite photos into their mapping systems. Through a variety of programs, the U.S. government has accumulated and released low-resolution photos of most areas of the planet. A couple of private companies supplement those photos with images captured from vans driving through major cities. Computers can match the digital map data with the satellite photos and the drive-by photos to provide realistic views of many cities.

Once you have the base maps, the objective is to overlay additional data on the maps. For example, you can obtain census data that displays average consumer characteristics such as income, house price, and number of autos within each geographic area. The GIS could be used to plot different colors for each income level. Next you can overlay the locations of your retail stores. If you are selling a high-price item such as a Cadillac, you want to locate the stores in areas of higher income.

Although you can buy base geographical data, how do you know the location of your retail stores? Or how do you plot the locations of delivery vehicles, or police cars, or trains? The easiest answer today is to use the **global positioning system (GPS)**, which is a set of satellites maintained by the U.S. government. A portable receiver tuned to the satellites will identify your location in latitude, longitude, and elevation (if it can reach four satellites) within 20 feet. Several handheld units are available for a few hundred dollars. If you work for the Department of Defense, you can get receivers that will identify your location within a few millimeters, but you need appropriate security clearances to obtain these receivers. Currently, the government is broadcasting the higher-resolution signal to civilian receivers, but sometimes blocks it during emergencies. Civilian models that combine signals from U.S. and Russian satellites provide even better resolution, but cost about $8,000. Europe has been trying to raise the funds to build its own GPS system (Galileo) and China has announced that it intends to create its own system as well. There is probably not enough money, or orbital space, for that many satellites; but competition is good and should encourage the U.S. to keep the system open.

As a model, the GIS makes it easier to spot relationships among items. Visual presentations are generally easy to understand and are persuasive. A GIS can be an

Technology Toolbox: Browsing Data with a PivotTable

Problem: You and your manager need to analyze sales based on several attributes.

Tools: OLAP cube browsers are designed to make it easy for you to examine slices of the cube and examine a fact (sales) based on several dimensions. The big DBMS vendors sell cube browsers with their packages. However, you can also use the Microsoft PivotTable that is built into Excel.

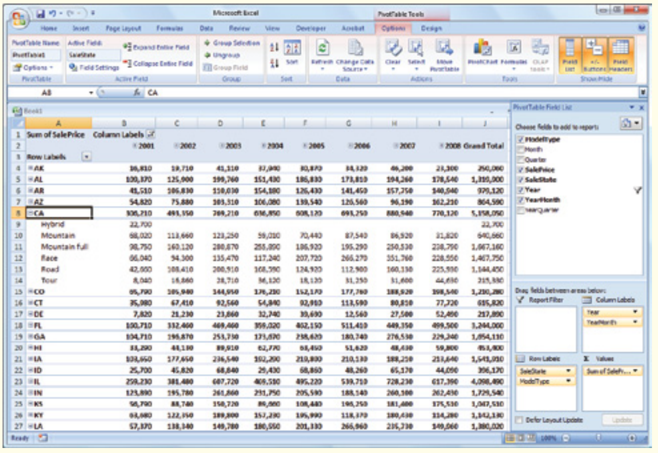

The Rolling Thunder Bicycle data provides a good example for a cube. The fact to be evaluated is the SalePrice. The more interesting dimensions are OrderDate, ModelType, and SaleState. The date presents a common problem: dates are hierarchies. You might want to examine sales by year, or you might want to drill down and see sales by quarter or month. Some cube browsers make it easy to create this hierarchy. Excel does not create it automatically, but you can use a query to convert the date into the different dimensions. The qryPivotAll query shows you how to compute the various year, month, and quarter fields.

In an Excel worksheet, use the Data/PivotTable option to begin. Select the External Data Source option because the data comes from a DBMS. Follow the basic steps to get data from a Microsoft Access Database and select your copy of the Rolling Thunder database. In the Query Wizard, find the qryPivotAll query, and select all of its columns by moving them to the right side box. Follow the Next and Finish prompts to return the data to the spreadsheet.

The structure of the PivotTable contains a space for row variables, column variables, page variables, and the main fact. Drag Year and YearMonth from the field list onto the table column location. Drag the ModelType and SaleState fields onto the row location. Finally, drag the SalePrice onto the main body of the table. You can place the remaining fields in the page location if you want to make them easy to find later.

The fun part is playing with the cube. Select a Year heading and click the minus button to collapse the detail. Do the same with the ModelType or SaleState column. Watch as the totals are automatically updated. Drag the ModelType or SaleState field heading to swap the order. Collapse or expand the field to summarize or drill down into the data. You can just as easily expand only one year or one state.

Quick Quiz:

1. How is the cube browser better than writing queries?
2. How would you display quarterly instead of monthly data?
3. How many dimensions can you reasonably include in the cube? How would you handle additional dimensions?

City	2000 pop	2009 pop	2000 per-capita income	2009 per-capita income	2000 hard good sales (000)	2000 soft good sales (000)	2009 hard good sales (000)	2009 soft good sales (000)
Clewiston	8,549	7,107	15,466	15,487	452.0	562.5	367.6	525.4
Fort Myers	59,491	64,674	20,256	30,077	535.2	652.9	928.2	1010.3
Gainesville	101,724	116,616	19,428	24,270	365.2	281.7	550.5	459.4
Jacksonville	734,961	813,518	19,275	24,828	990.2	849.1	1321.7	1109.3
Miami	300,691	433,136	18,812	23,169	721.7	833.4	967.1	1280.6
Ocala	55,878	55,568	15,130	20,748	359.0	321.7	486.2	407.3
Orlando	217,889	235,860	20,729	23,936	425.7	509.2	691.5	803.5
Perry	8,045	6,669	14,144	19,295	300.1	267.2	452.9	291.0
Tallahassee	155,218	172,574	20,185	27,845	595.4	489.7	843.8	611.7
Tampa	335,458	343,890	19,062	25,851	767.4	851.0	953.4	1009.1

Figure 9.18

Geographic sales data. We suspect that sales of hard and soft goods are related to population and income. We also want to know whether there are regional patterns to the sales.

effective means to convince management that neighborhoods have changed and that you need to move your retail outlets. A GIS can also be used for simulations to examine alternatives. For example, a GIS oriented to roadmaps can compute the time it would take to travel by different routes, helping you establish a distribution pattern for delivery trucks.

Example

Consider the problem faced by a manager in a small retail chain that has stores located in 10 Florida cities. It sells a combination of hard goods (such as cleaning supplies, snack items, and drapery rods) and soft goods (mostly clothing). For the most part, profit margins for soft goods are higher than for hard goods. However, total sales of hard goods seem to be better than those of soft goods—except in certain stores. The manager has been unable to find a reason for the difference, but a friend who has lived in Florida longer suggested that there might be some geographical relationship. The basic numbers are presented in Figure 9.18.

Because there are only 10 cities, it might be possible to identify patterns in the data without using a GIS. However, an actual firm might have several hundred or a few thousand stores to evaluate. In this case, it is much more difficult to identify relationships by examining the raw data. It is better to use a GIS to plot the data. Different colors can be used to highlight large increases in sales. By overlaying this data with the population and income data, it is easier to spot patterns. Notice in Figure 9.19 that there is a correlation between population and total sales. Also, notice that sales in the northern cities are concentrated more in hard goods than in the southern cities. Each of the radar charts in the north points to the northwest (left), while those in the south point to the northeast (right). Once you see this pattern, it becomes clear; yet it is difficult to see the pattern within the raw data, or through any other chart.

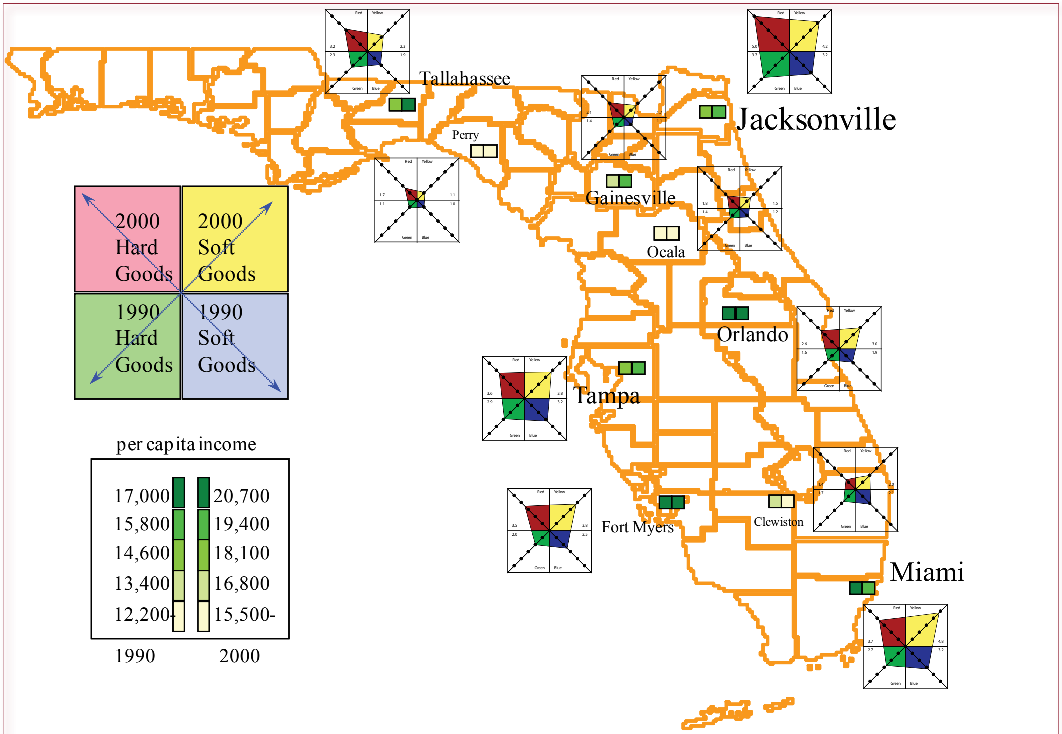

Figure 9.19

Geographic-based data. It is difficult to display this much data without overwhelming managers. Notice that the sales (radar) graphs use size and shape to highlight total sales, and the changing sales mix. Income is color-coded in a smaller graph. Notice in the sales graphs that the northern counties experienced a greater increase in sales in hard goods compared to the southern counties.

Most GIS tools use simpler methods of displaying data. Figure 9.20 shows a common approach where geographic regions (states) are shown in darker colors for higher values. The chart was created by creating a query in Access to compute sales by state for the year 2008. Microsoft Map Point was then used to draw the map and shade the values by state. Map Point also includes demographic data from the U.S. Census Bureau at relatively detailed levels that can be added to the chart. More sophisticated, and more expensive, tools such as ESRI's ArcInfo can display multiple layers of data using overlays. Google maps (and Microsoft's Bing) provide some simple GIS tools using their online systems. The benefit is that the tools are free, but the options available are limited. Typically, you can display only push-pins to highlight specific locations. However, they do support some programmatic tools so you can automatically create charts from your Web server.

Data Mining

Is it possible to automate the analysis of data? Many decision makers, particularly researchers, are buried in data. Transaction-processing systems, process control, automated research tools, even Web sites all generate thousands,

Figure 9.20

Geographic-based data. Tools such as Microsoft Map Point often display location-based data by shading geographic regions based on the relative levels of the data. This map shows sales by state for Rolling Thunder Bicycles in 2008.

millions, or even billions of pieces of data a day. Sure, you can retrieve the data into a data warehouse, run queries, or put it into a spreadsheet, but it is still difficult for humans to comprehend raw data.

Data mining consists of a variety of tools and techniques to automatically retrieve and search data for information. Originally, the term was derogatory because it represents an undirected search for relationships. Statistically, results obtained from nonscientific searches can be spurious and not repeatable. On the other hand, the tools can find minute comparisons that are not ordinarily found through traditional statistical methods. As the flood of data increases, more companies are turning to automated and semiautomated tools to help search databases and make it easier to visualize patterns. Additional models and research can then be used to investigate and validate the relationships.

Data mining tools exist for numerical and other types of data. The numerical tools tend to be based on statistical theory. The others sometimes use statistical theory but are often highly specialized. For example, event analyses look for patterns in data based on timing or sequence—such as the path viewers take through a Web site. This chapter focuses on the more generic statistical tools because they are used in many different companies and tasks. The specialty tools typically require a detailed background in a specialty discipline, so you can learn more about them when you study specific topics.

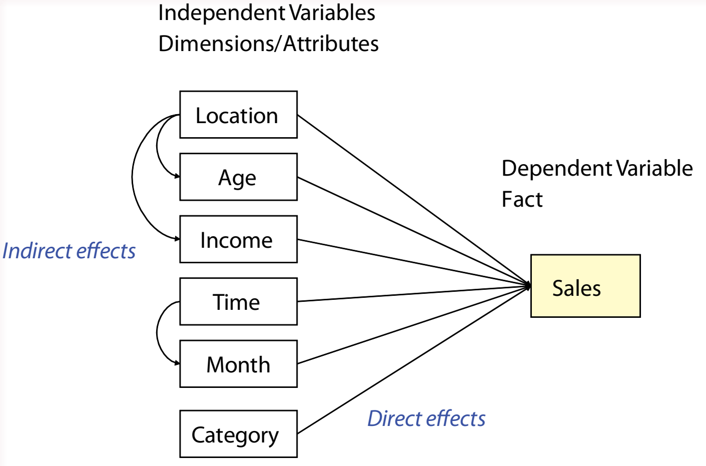

Figure 9.21

Common data mining task. Determine how independent variables affect the dependent fact variable. Models can have direct or indirect effects. Basic linear models are easy to estimate, non-linear ones are harder to estimate and more difficult to understand.

As shown in Figure 9.21, a common theme in the numerical data mining tools is the search for items that influence a given fact variable. Much like the data cube, you might choose the Sales variable as the fact to be investigated. In statistics terms, it becomes the dependent variable. You then want the system to examine various attributes or independent variables to see how they influence the dependent variable. For instance, you might have data on customer location, income, product category, time and month purchased, marketing budget, and a host of other variables that could conceivably affect the level of sales. Remember this goal as you read the rest of this section. Even if you have limited experience with statistics, remember that most of the tools are trying to identify which independent variables strongly affect the dependent variable.

Correlation is a key statistical tool that is leveraged in data mining. A data mining system can compute the cross correlation for all dimensions. High correlations provide a useful indicator of how one dimension (variable) affects another. Correlation typically represents direct effects between pairs of variables. Multiple regression is an extension of correlation where multiple dimensions (independent variables) are used to predict values for a dependent variable. The regression coefficients have long been used by statisticians to measure the importance of each attribute. Special data mining tools extend these concepts into nonlinear relationships, automatically searching for relationships between variables. For example, marketers could mine the sales data to determine which variables had the strongest impact on sales revenue, including price, quality, advertising, packaging, or collections of product attributes. Regression tools provide numerical estimates of the coefficients for direct and indirect effects, along with measures of the accuracy of the estimate (standard errors).

Clustering is another data mining technique. It tries to find groups of items that have similar attribute values. For example, a car manufacturer might find that younger buyers are attracted to one car model, while older buyers tend to pur-

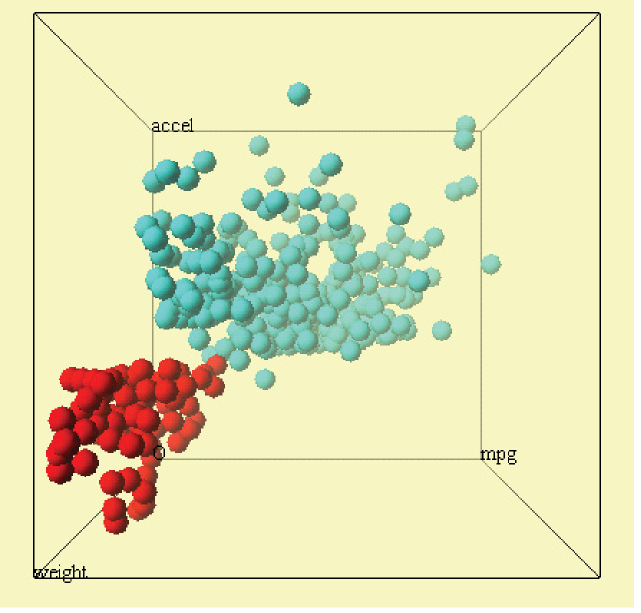

Figure 9.22

Clustering. Sometimes you want to know which items are similar to others. You might face distinct groups of customers, or your products might be perceived in certain clusters. This example compares cars in terms of three dimensions (acceleration time, weight, and gas mileage) to reveal two distinct groups.

chase a different model. These concepts are relatively easy to understand in two or three dimensions. The tools can search in higher dimensions, looking for groupings across dozens of attributes. The results are more difficult to interpret and act on, but can provide useful information in dealing with complex datasets. Figure 9.22 shows an example that compares cars on three dimensions: acceleration time, weight, and miles per gallon. The results show that there were two distinct groups of cars at that point in time. Heavy cars with low acceleration times (faster), and low miles per gallon versus light, slow cars that achieved more miles per gallon.

Clustering is a useful tool when you are first looking at data or when you need to reduce the number of dimensions. In the car example, you can reduce the analysis to looking at those two groups of cars. Similarly, you might learn that you have three or four types of primary customers, enabling you to focus your analysis and decisions on how representatives from each group will react to changes.

One of the classic data mining tools is market basket analysis (or association rules). **Market basket analysis** was designed to address the question, What items do people tend to purchase together? Figure 9.23 shows the classic example that was purportedly identified by a convenience store. Managers found that on weekends, people often purchased both beer and diapers. This raises the immediate issue of how you can use the results. You might stock the items near each other to encourage people to buy both. Or, you might place the items at opposite ends of an aisle—forcing people to walk past the high-impulse items such as chips. Or, perhaps you discount the diapers to attract more customers and make up the additional profit by not discounting the beer. The tool is generic and is used by

Figure 9.23

Market basket analysis. What items do people buy together? Data mining tools can examine each purchase to identify relationships. You can use this information to increase cross selling.

companies such as Amazon.com and Netflix to show you books that were bought by other customers. If you are interested in the first book, you might also like the books that other people purchased along with that book. Tivo, the television recording system, uses a similar process to identify programs that you might be interested in watching. To illustrate the hazards of data mining, many people have reported interesting twists with Tivo. For instance, if you watch one movie for children, the system will begin to offer more shows geared toward children. Market basket or association analysis is relatively easy to compute—particularly for pairs of items. Several tools exist to perform the computations for you, including a couple of open source tools such as Weka.

Do not let the name fool you. Market basket analysis has many more uses than the example shows. More properly known as association analysis, it works well with categorical data that does not need to be measured numerically. It basically estimates the probability that any two (or more) items will occur together. These items could be purchases, or they could be diseases, events such as Web clicks, votes, and so on.

Several data mining tools have been built in the past few years. They are relatively easy to use, but you might want to review your statistics to help understand the results. One catch with the tools is that some of them can take time to run—even with fast machines, some complex analyses can take hours. For example, market basket analysis is relatively fast with pairs of items, but requires exponentially more computations and time if you attempt to compare more than two items at a time. A bigger problem is that the results might be meaningless or not reproducible. They might simply have arisen because of some random occurrence in the data. The bottom line is that you have to carefully evaluate each piece of information to make sure it is relevant and important.

Expert Systems

Can information technology be more intelligent? Can it analyze data and evaluate rules? Imagine your life as a top-notch manager. Coworkers perceive you as an expert and value your advice and problem-solving skills. You are constantly answering questions and there are always more problems than you can handle. You are using decision support systems and integrated information technology to perform your job better and more efficiently, but it is

Technology Toolbox: Forecasting a Trend

Problem: You need a basic forecast of sales data.

Tools: Many statistical tools and packages exist to build models and estimate trends. The simplest technique is a linear regression. Excel has tools to calculate and display the results.

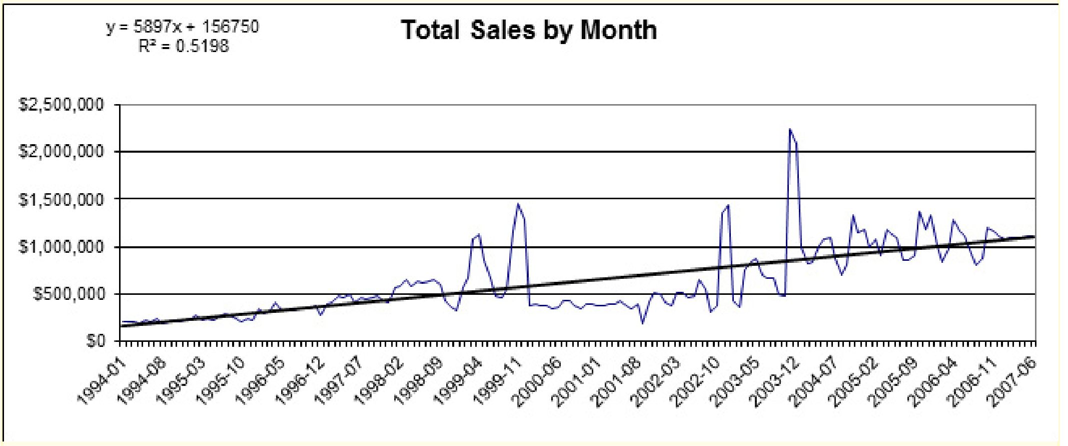

Open the Rolling Thunder database and create a query that computes the total sales by year and month. The Format function converts the date: Format(OrderDate, "yyyy-mm"). Save this query within Access, and close the database. To import the data into Excel, open a new worksheet and choose the Data/Import/New Database Query option. This approach will build a dynamic link to the query.

Create a line chart of the data. Right-click the plotted line and select the option to add a trend line. You should stick with the linear trend, but you might want to test the other choices. Select the options to display the equation on the chart.

To see the actual forecast values, select the entire column of data for SumOfSalePrice, but do not include the title row. Scroll down to the last row in the series and find the small square handle on the lower-right-side corner. Drag this handle to extend the series for six months. Add the appropriate months in the month column. Edit the chart and edit the series selections to include the new rows. Notice that the new points fall exactly on the trend line.

Excel includes more powerful statistical tools, including one that performs multiple regression when you want to examine the effect of many variables. Although this dataset contains only one variable, you can still use it to test the tool. First, insert a new column that numbers the months as 1, 2, 3, and so on. Ensure the Analysis ToolPak is loaded and select Tools/Data Analysis and pick the Regression tool. As the Y range, select the sales column including the title row but excluding the six rows you forecast earlier. Select the month column as the X range. Be sure to check the Labels box, and then run the regression. This tool provides the standard regression coefficients, probabilities, and diagnostics.

Quick Quiz:

1. Why is a linear forecast usually safer than nonlinear?
2. Why do you need to create a new column with month numbers for regression instead of using the formatted year-month column?
3. What happens to the trend line r-squared value on the chart when you add the new forecast rows to the chart?

not enough. Can technology help you with more complex decisions and problem solving? From another perspective, when you encounter new problems with different, complex models, it would be helpful to have an expert assist you with applying and understanding the models. Yet experts or consultants are expensive and not always available. Can you somehow capture the knowledge and methods of experts and use technology to make this knowledge available to workers throughout the company?

Specialized Problems

Expert systems have proven useful for many problems. The goal of an expert system is to enable novices to achieve results similar to those of an expert. The users need to understand the basic problem, learn the terminology, and be able to answer questions. For example, a typical patient would not be able to use a medical expert system because the questions and terms would not make any sense.

Think of an expert system as a consultant in a box. The consultant can solve only certain specific problems. For example, perhaps a retail store manager needs to estimate buying patterns for the next few months. The manager might call a marketing consultant to survey buyers and statistically search for patterns. The consultant will ask questions to determine the basic objectives and identify problems. Similarly, a production manager might be having problems with a certain machine. The manager might call a support line or a repair technician. The advice in this situation will be quite different from the marketing example, because the topics (or domains) of the two problems are different. It would be difficult to create one computer program that could help you with both types of problems. On the other hand, there are similarities in the approach to the two problems. Computerized expert systems are designed to solve narrow, specialized problems. Each problem can be relatively complex, but it must be reasonably well defined. Many business problems fall into this category, and expert systems can be built for each problem.

Diagnostic Problems

Several problems in the world can be classified as diagnostic situations. These problems arise when the decision maker is presented with a set of symptoms and is asked to find the cause of the problem, as well as solutions. Consider a firm that uses a complex machine. If the machine breaks down, production stops until it is fixed. In addition, maintenance tasks have to be performed every day to keep the machine running. The company hires an engineer to perform these tasks. The engineer also knows which adjustments to make if various symptoms appear. This system has been working well, and the company wishes to expand to other locations with a franchise system. The problem is that there is only one engineer, and it would be too expensive to have a highly trained engineer at each location.

One possible solution would be to set up a phone link between the franchises and the engineer. One person at each franchise would be trained in the basics of the machine. If problems arise, the person could call the engineer. The engineer would ask specific questions, such as "What do the gauges show?" The answers will lead the engineer to ask other questions. Eventually, the engineer makes recommendations based on the answers.

Of course, if there are many franchises, the engineer will be too busy to solve all of the problems. Also, if the businesses are located in different countries, the time differences may not allow everyone enough access to the engineer. A bet-

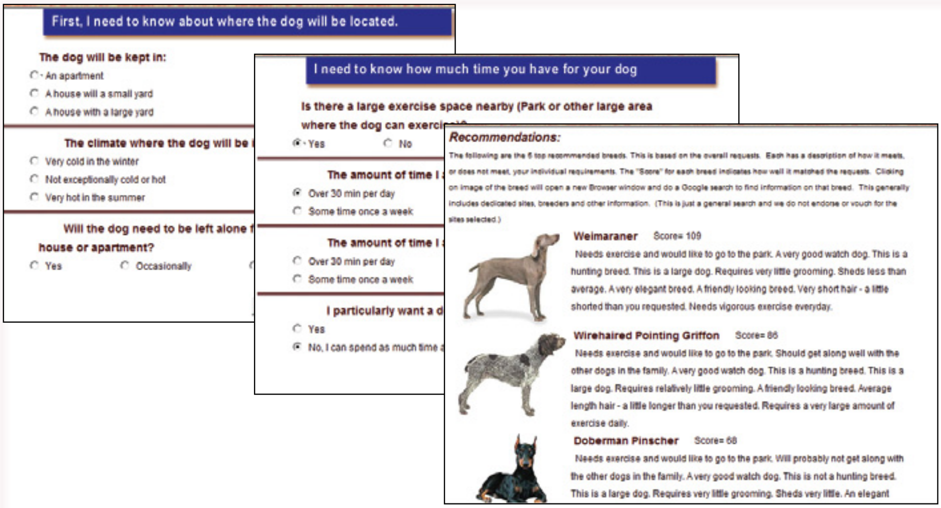

Figure 9.24

Expert System example from ExSys. This sample expert system acts as a knowledgeable dog lover and asks questions about how you view characteristics of dogs. Based on your responses it makes a recommendation of which breed might be good for you. The Web site contains many other examples.

ter solution is to create a computerized expert system. All the expert's questions, recommendations, and rules can be entered into a computer system that is distributed to each franchise. If there is a problem, the on-site person turns to the expert system. The system asks the same questions that the engineer would and arrives at the same recommendations.

As shown in Figure 9.24, expert systems also have the ability to explain their recommendations. In more complex examples, while running the ES, the user can ask it to explain why it asked a particular question or why it arrived at some conclusion. The ES traces through the answers it was given and explains its reasoning. This ability helps the user gain confidence in the decisions, allows mistakes to be corrected, and helps the users remember the answer for future reference.

The business world offers many examples of diagnostic situations, such as identifying causes of defects, finding the source of delays, and keeping complex equipment running. The common characteristic is that you are faced with a set of symptoms, and you need to find the cause.

Speedy Decisions

Other situations can benefit from the use of expert systems. Even if a problem is not exceedingly complex, you could use an expert system to provide faster responses or to provide more consistent recommendations. Several advantages can be gained from making decisions faster than your competitors do. If you can identify a trend in stock prices before anyone else, you can make a higher profit. If you can answer customer questions faster, they will be more likely to shop with you in the future. If you can provide a loan to a customer sooner than anyone else, you will do more business.

Transaction-processing systems keep much of the basic data that you need to make decisions. Decision support systems help you analyze that raw data. Both of these tools enable you to make decisions faster than trying to make the decision without any computers. However, it still takes time for a human to analyze all of the information.

Consider the case of a bank loan. In order to get a loan, you go to the bank and fill out a loan application form. You tell the loan officer why you want the loan and provide basic data on income and expenses. Depending on the amount of money involved, the banker will probably check your credit history, get appraisals on any collateral, and perhaps get approval by a review officer or loan committee. All of these actions take time.

Now, consider the steps involved with a computerized process. First, you need to tell the bank that you want a loan. Instead of driving to the bank, you could use the telephone. With a push-button phone, you enter information directly into the bank's computer. The computer would give you a choice of loan types (car, boat, personal, etc.), and you push a button to select one. You enter the amount of money you want to borrow. The next step is to check your credit history. Your income, expenses, and credit record are available to the bank from national credit reporting agencies. The bank might also have its own database. The bank's computer could be connected to credit agency computers to collect additional data on your credit history.

To make the final decision, the bank needs a set of rules. These rules take into account the size of the loan, the value of the collateral, as well as your income, expenses, credit history, and existing loans. When the bank has determined the proper rules, the computer performs the analyses. If the bankers trust the rules, the computer could make the final decision. For example, there would be no need for a loan officer to be involved in simple decisions, such as making small car loans to customers with large savings accounts. With an expert system, a bank can cut the loan-approval period down to a few minutes on the phone.

Many other decisions need to be made rapidly. The first step in all of these cases is to make sure that the transaction-processing system provides the necessary raw data. The second step is to create a set of rules for making the decision. The difficulty lies in finding these rules. For some problems, there are well-defined rules that can be trusted. For other problems, the rules may not exist. In this case, the company will probably still need a human to make the final decision.

Consistency

The example of the bank loan demonstrates another advantage of expert systems. Business decisions are subject to a wide variety of nondiscrimination laws. An expert system can be used to provide consistent decisions. The rules followed by the ES can be set up to avoid illegal discrimination. Businesses also have credit ratings, which are often determined by Credit Clearing House (CCH). CCH uses an expert system to make the "easy" decisions, which speeds up the process by allowing humans to focus on the more complicated cases. It also leads to consistent application of the rules.

Consider the loan example. If each loan officer makes individual decisions, it is hard to determine whether they are consistent with corporate policy. Each individual decision would have to be checked to make sure it was nondiscriminatory. On the other hand, a committee could spend several weeks creating a set of lending rules that can be verified to be sure they are legal and ethical. As long as the bank

Reality Bytes: Don't Bet Against the Machine

Online poker sites are hot. Many people, particularly college students, see it as a simple diversion or a way to pick up a couple of bucks. But have you considered how easy it is for someone to cheat at online poker? Some of the better gaming sites try to detect obvious cheats, but it can be hard to catch a group of players who work together to fleece an outsider. Players also face a new threat: poker bots—automated programs that play poker. Some people are aboveboard and build them to search for new ideas in AI. There is even an international competition for poker bots. Others build poker bots and put them online without telling opponents. But how good are the poker bots? Dr. Jonathan Schaeffer, a professor of computer science at the University of Alberta, has spent over 14 years developing a poker bot, and he competes aboveboard in public tournaments. In 2003, his program went head-to-head against a top poker player. For the first 4,000 hands, the computer was leading, but the poker player eventually analyzed the program's approach and developed a new strategy to beat it. But how many poker games involve top-level players? Dr. Schaeffer observed that "when you look at the low-limit tables, I have not doubt that there are computer programs that can win at that level consistently. And don't forget that these programs can play all the time, without getting tired. And they can play at multiple tables." In 2005, the Golden Palace in Las Vegas sponsored a poker bot competition. Poker ProBot created by Hilton Givens of Lafayette, IN won after nearly 5,000 hands, earning the $100,000 prize. He also had the opportunity to match his program against human pro player Phil Laak. Phil trounced the machine in 399 hands, and also beat the University of Alerta's PokiX bot in 290 hands. But most experts believe that just as in the chess world, the machines will eventually improve to beat humans. In 2011, the U.S. government shut down most of the online poker sites doing business America--primarily because they allegedly violated the U.S.laws against online gambling. Ultimately, online poker is a fool's game.

Adapted from Shawn P. Roarke, "Bots Now Battle Humans for Poker Supremacy," *Fox Sports*, July 20, 2005.

employees follow the recommendations of the ES, the outcome should not be discriminatory. Because there should be few cases where the loan officer overrules the ES, managers will have more time to examine each of these circumstances.

Many business decisions need to be performed consistently to avoid bias and to treat people equally. Loans, pricing, raises, and promotions are some examples. However, there can be problems with using a computer system to enforce standards. The main difficulty lies in creating a set of rules that accurately describe the decisions and standards. For example, it might be useful to have a set of rules regarding raises and promotions, but think about what happens if an employee's job does not fit the basic rules. Organizations continually change, which means the rules have to be monitored and changed regularly.

Training

Training employees is closely associated with problems of consistency. All organizations must train employees. If the tasks are complex and the decisions highly unstructured, it can take years for employees to learn the rules and gain the experience needed to deal with problems. Two features of expert systems help employ-

First, compute the monthly income before taxes.
Next, compute the monthly payment of the loan.
If the payment is greater than 5% of income:
Compute total of other loans payments.
Compute payments as percent of monthly income.
    If this percent is less than 25%:
        If the new loan is less than 10%, make loan.
            Else:
                If total monthly expenses are less than 40% of income, make the loan.
                Else:
                    If less than 50% and has been a customer for more than 5 years or if less than 60% and has been a customer for 10 years and has lived at the same address for 5 years, make the loan.

Figure 9.25

Sample rules for the bank loan. A portion of the business rules that are used to determine whether a person should get a loan.

ees learn. First, employees learn what questions need to be asked. In particular, after using the system for a while, certain groups of questions will occur together. Second, most expert systems have provisions for explaining their answers (and the motivation for each question). At any point, an employee can ask the expert system why it asked a certain question or why it reached a conclusion.

Building Expert Systems

How do you create an expert system? At first glance, you would suspect that expert systems are hard to create. However, except for one step, which is hard, tools exist to make the job easier. The area that causes the most problems when you are creating expert systems is finding a cooperative expert who fully understands and can explain the problem. Some problems are so complex that it is difficult to explain the reasoning process. Sometimes the expert may rely on vague descriptions and minor nuances that cannot be written down. Even though expert systems can deal with these types of problems, it might take too long to determine the entire process. Also, if you transfer the expert's knowledge to a computer, the expert might worry about losing his or her job.

Most expert systems are built as a knowledge base that is processed or analyzed by an inference engine. A **knowledge base** consists of basic data and a set of rules. In most situations, an inference engine applies new observations to the knowledge base and analyzes the rules to reach a conclusion.

The basic steps to create an expert system are (1) analyze the situation and identify needed data and possible outcomes, (2) determine relationships between data and rules that are followed in making the decision, (3) enter the data and rules into an expert system shell, and (4) design questions and responses. A **knowledge engineer** is often hired to organize the data, help devise the rules, and enter the criteria into the expert system shell, or supervise programmers as they create an expert system.

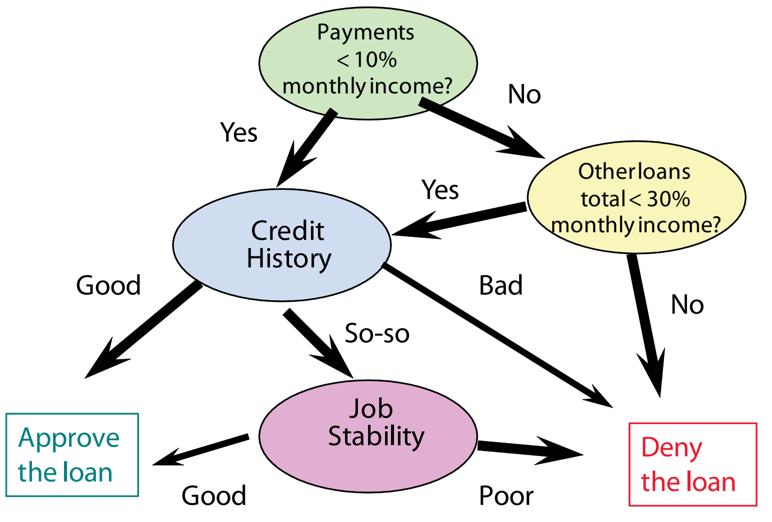

Figure 9.26

Decision tree for sample bank loan expert system. Parts of a knowledge base are often expressed as a decision tree. Each answer to a question leads to additional questions and eventually to a decision. Notice that questions sometimes require numeric answers but can also rely on subjective comments.

Knowledge Base

A knowledge base is more than a simple database. It consists of data but also contains rules, logic, and links among data elements. In most cases, it contains less structured and more descriptive data. For example, an ES for medicine might have a list of symptoms that contains items like "high temperature" and "intense muscle pain." This knowledge base is the reason why the problem must be narrow in scope. Even narrow, well-defined problems can require large amounts of information and thousands of rules or relationships. The real challenge in building expert systems is to devise the knowledge base with its associated rules.

Three basic types of expert systems are in use today. They are defined by how the knowledge base is organized: by rules, frames, or cases.

Rules

The heart of a rule-based ES is a set of logical rules. These **rules** are often complicated. Consider some of the rules that might be needed for an ES to evaluate bank loans, as shown in Figure 9.25. This example has been simplified to keep it short. There will usually be hundreds of rules or conditions to cover a wide variety of situations. Rules are often presented as If… Then… Else … statements. They can include Boolean conjunctions such as AND, OR, NOT. Figure 9.26 presents a portion of a **decision tree** that visually displays the rules.

The difficulty with any ES lies in determining these rules. Some of them will be easy. Others will be complex. Most of them will come from the expert. Unfortunately, most people do not usually express their thoughts in the form of these rules. Although a person might follow rules of this sort, they can be difficult to express. It is even more difficult to remember all the rules at one time. For instance,

Welcome to the Loan Evaluation System.
What is the purpose of the loan? car
How much money will be loaned? 15,000
For how many years? 5

The current interest rate is 10%.
The payment will be $297.02 per month.

What is the annual income? 24,000

What is the total monthly payments of other loans?
Why?

Because the payment is more than 10% of the monthly income.
What is the total monthly payments of other loans?
50.00

The loan should be approved, because there is only a 2% chance of default.

Figure 9.27

Bank loan sample screen. An expert system carries on a dialogue with the user. The ES asks questions and uses the answers to ask additional questions. The user can ask the ES to explain a decision or a question. Hence the ES can be used for training purposes.

say you have lived in the same place for five years and a new person moves into the neighborhood. She asks you to describe the best ways to get to school, the mall, and the grocery store. Then she asks you for the best shortcuts if one of the roads is closed. This problem is relatively simple, but can you sit down right now and provide a complete list of all the rules?

Creating an ES

More commonly today, an ES is built from an expert system shell. This program provides a way to collect data, enter rules, talk to users, present results, and evaluate the rules. To create an ES, you must know what data you need and all of the rules. Once you express this knowledge base in the format used by the shell's inference engine, the shell takes care of the other problems. Many ES shells are available on a wide variety of computers; Jess and Clips are two common systems that are available free or at relatively low cost.

To understand how to create an ES, consider the bank loan example. A typical dialogue with the user (the loan clerk) appears in Figure 9.27. Notice that the ES begins by asking some basic information-gathering questions. The responses of the user are underlined. Once the basic data is collected, the ES performs some computations and follows the built-in rules. Notice that the ES follows the rule that asks for the other loan payments. However, the loan clerk does not know about this rule, so he or she asks for clarification. This ability to ask questions is a powerful feature of expert systems.

Reality Bytes: One Smart Machine is Better than Thousands of People

Many people have touted the potential benefits of crowd-sourcing: breaking a problem into small pieces and using the Internet to assign each piece to thousands of workers. Amazon developed the Mechanical Turk to facilitate these types of tasks. Anyone can sign up to complete tasks and receive wages—often only pennies a chore. The goal is to handle tasks that are easy for humans but difficult for machines. The problem is that people being paid low wages are usually not very good and machines are getting smarter. The Web site Yelp ran a comparison test. First, researchers developed a multiple choice test to try to identify the best Turkers who could correctly categorize a business (restaurant, clothing, etc.), phone number, and address; based on its Web site. The initial test was given to 4,660 applicants, only 79 of them passed. This select group of the "best" were then given the problem of classifying business information where three people saw the same site. If a majority (2-3) agreed, the site was then classified. The researchers then applied a Naïve Bayes classifier data mining algorithm on the same set of problems. The algorithm was trained using 12 million examples submitted earlier by Yelp users. The computer algorithm easily beat the humans—correctly classifying the sites a third more often than the people. Does that mean machines are smarter than humans? Or just better than poorly-paid humans with low motivation?

Adapted from Christopher Mims, "AI Defeats the Hivemind," *Technology Review*, December 20, 2010.

Once you have collected all of the rules involved in the problem, you enter them into the ES shell. The shell lets you type in the questions you want to ask the user. You define the calculations and tell the shell how to look up any other information you need (e.g., the interest rates for auto loans). You then enter the conditions that tell the shell what questions to ask next. If there are many rules with complex interactions, it is more difficult to enter the rules into the shell. One advantage of ES shells is that you generally have to enter only the basic rules and data. As the user enters the data, the shell performs the calculations and follows the rules. The shell also automatically answers the user questions. You do not have to be a computer programmer to create an ES with a shell. With training, many users can create their own expert systems using a shell. However, there are many dangers inherent in ES development, so it helps to have someone evaluate and test the resulting system.

Limitations of Expert Systems

Expert systems are useful tools that can be applied to several specialized problems. However, several important drawbacks arise in their design and use. First, they can be created only for specific, narrowly defined problems. Some complex problems contain too many rules with too many interactions. It quickly becomes impossible to express all of the interrelationships. For example, it is currently impossible to create a medical diagnostic system that covers all possible diseases. However, smaller systems are in use that help determine drug dosages and other treatments such as radiation levels for cancer patients.

Another problem that users and designers have encountered is that it can be difficult to modify the knowledge base in an expert system. As the environment or problem changes, the expert system needs to be updated. The changes are relatively easy to make if they affect only a few rules. However, many expert systems use hundreds of interrelated rules. It is not always clear which rules need to be altered, and changes to one rule can affect many of the others. In essence, as the situation changes, the company is forced to completely redesign the expert system. In fast-changing industries, it would cost too much to continually redesign an expert system. In the lending example, a policy change based on monthly income would be relatively easy to implement. On the other hand, some changes in policy would force a complete redesign of the expert system. For instance, a bank might decide to grant loans to almost everyone but charge riskier people higher interest rates.

Probably the greatest difficulty in creating an expert system is determining the logic rules or frames that will lead to the proper conclusions. It requires finding an expert who understands the process and can express the rules in a form that can be used by the expert system.

Management Issues of Expert Systems

Creating and building an expert system involve many of the same issues encountered in building any other information system. For instance, the problem must be well defined, the designers must communicate with the users, and management and financial controls must be in place to evaluate and control the project.

However, expert systems raise additional management issues. Two issues are particularly important: (1) if an expert transfers knowledge to an expert system, is there still a need for the expert; and (2) what happens when the expert system encounters an exception that it was not designed to solve?

The answer to the first question depends on the individual situation. In cases where the problem is relatively stable over time, it is possible to transfer expert knowledge to software—enabling the firm to reduce the number of experts needed. If this action results in layoffs, the experts will need additional incentives to cooperate with the development of the system. In other cases, the firm will continue to need the services of the experts, to make changes to the ES and to solve new problems. Before starting an ES project, managers need to determine which situation applies and negotiate appropriately with the experts.

The second problem can be more difficult to identify. Consider what happens when workers rely on an expert system to make decisions, and management then cuts costs by hiring less-skilled workers. The new workers do not understand the system or the procedures—they simply follow decisions made by the rules in the ES. If an exception arises, the ES may not know how to respond or it may respond inappropriately. A customer then would be left to deal with an underskilled worker who does not understand the process and cannot resolve the problem.

Specialized Tools

Can machines be made even smarter? What technologies can be used to help managers? Research in artificial intelligence (AI) examined how humans are different from computers. This research led to tools that can be used for certain types of problems. Some of the ideas come from the early days of computers, but it has taken until now for machines to be developed that are fast enough to handle the sophisticated tasks. Ideas in AI have come from many disciplines, from biology to psychology to computer science and engineering.

Humans are noticeably better than computers in six broad areas: pattern recognition, performing multiple tasks at one time, movement, speech recognition, vision, and language comprehension. Some of these concepts are related, but they all represent features that would make machines much more useful. Even with current technological improvements, most observers agree that it will be several years before these features are available.

Pattern Recognition and Neural Networks

One of the early issues in AI research was the question of how human brains worked. Some people suggested that to make intelligent computers, the computers would have to work the same way as the human brain does. An important conclusion from this research is that humans are good at pattern recognition.

Humans use pattern recognition thousands of times a day. It enables people to recognize coworkers, to spot trends in data, to relate today's problems to last year's changes. Many problems in business could benefit from machines that can reliably recognize patterns. For example, what characteristics do "good" borrowers have in common? How will changes in the economy affect next year's sales? How are sales affected by management styles of the sales managers?

Pattern recognition is used by people to solve problems. It is one of the reasons teachers use cases to teach students to solve business problems. If you notice that a problem is similar to a case you have seen before, you can use your prior knowledge to solve the problem. Imagine how useful it would be if an expert system could recognize patterns automatically.

One current technique that is used to spot patterns is the use of neural networks. Initial study indicated that the brain is a collection of cells called *neurons* that have many connections to each other. Each of these cells is relatively simple, but there are approximately 100 million of them. In some respects, a neuron resembles a simple computer. It can be at rest (off), or it can fire a message (on). A neuron responds to other cells (input) to send messages to other neurons (output). A collection of these cells is called a **neural network**. Human neural cells are actually more complicated, but researchers have focused on this simplified form.

A common current example is a bank that uses a neural network to spot credit card fraud. In some cases, Mellon Bank's neural network identified fraudulent patterns even before the human investigators spotted them. It is faster and more accurate than an earlier expert system. The original expert system looked at a limited number of variables and indicated 1,000 suspects a day, which was far more than actually existed and too many for the investigators to keep up with. The new neural network system examines more variables, lists fewer false suspects, and adjusts its methods on its own.

A finance manager might use a form of pattern recognition to search for patterns in the financial markets to forecast future movements. Of course, with thousands of other people searching for patterns, the patterns would not last very long. Similarly, a banker might use pattern recognition to classify loan prospects.

Neural networks can be built with software. Also, computer chips are available today that function as neural networks. Neural networks can be measured in two ways by (1) the number of neurons and (2) the number of interconnections between the individual cells. It is fairly easy to increase the number of cells, but the number of possible interconnections increases very rapidly. For instance, if there are four cells, there are six possible connections. With 10 cells, there are 45 connections. With 1,000 cells, there are half a million connections. In general, if there are N cells, there are $N(N - 1)/2$ possible connections. For many purposes, not every connection is needed, but with millions of cells, a neural network would

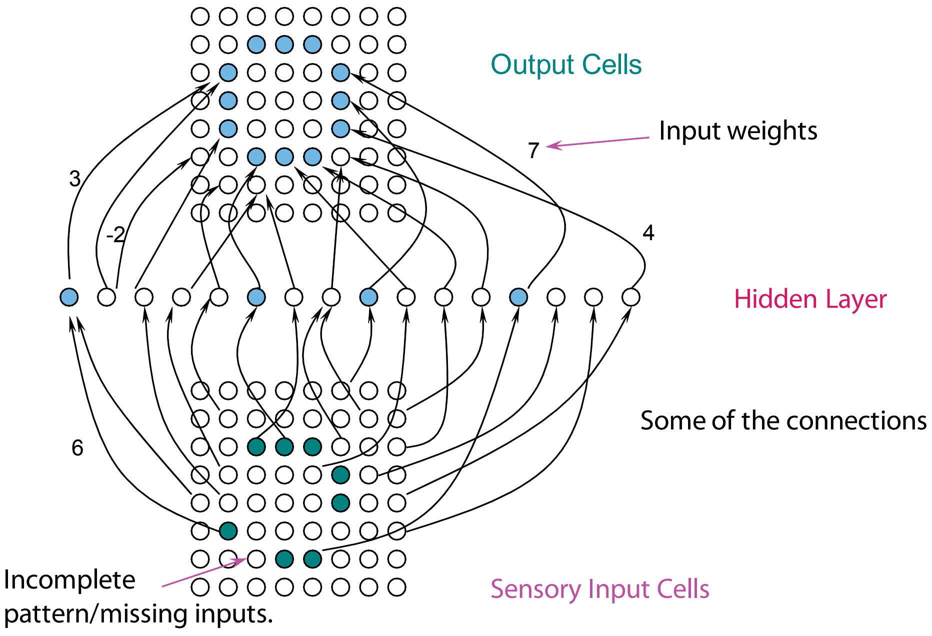

Figure 9.28

Neural net for pattern matching. Input cells convert data to binary form. The required hidden layer recodes the inputs into a new internal representation. The connections represent outputs from the lower layers. When total input levels exceed some value, the receiving cell fires. Any cell can be connected to many other cells. Input weights are determined by training. The output cells are triggered when total input levels from the connections exceed some threshold. Note that a pattern can be recognized even if some input cells are wrong or not firing.

incorporate a large number of connections. Most existing networks use only a few thousand cells.

Figure 9.28 presents a version of how a neural network converts an array of input sensors into a hidden layer and then stores patterns on an output layer. One useful feature of the neural network approach is that it is fairly good at identifying patterns even if some of the inputs are missing.

What can neural networks do that cannot be done with traditional computers? The basic answer is "nothing." However, they provide a new way of thinking about problems. More important, with hardware specifically designed to process neural networks, some difficult problems can be solved faster than with traditional computers. The primary objective of neural networks is the ability to store and recognize patterns. A well-designed network is capable of identifying patterns (such as faces or sounds) even if some of the data is missing or altered.

Another advantage that researchers hope to achieve with neural networks is the ability to simplify training of the computer. The discussion of expert systems noted that changes in the business often mean that knowledge engineers have to redesign the entire expert system. A neural network has a limited ability to "learn" by examining past data. Feeding it proper examples establishes the interconnection weights that enable the network to identify patterns. In theory, neural networks

Reality Bytes: Make the Trains Run on Time

Scheduling trains, planes, or healthcare appointments might appear to be an easy problem. Just figure out the capacity, compute the time needed to get from point A to point B, and then route everything. But, the problem is much more complex. To start with, a complete routing scheme is a difficult problem even in computer science, with no simple solution. But, it gets considerably worse when things go wrong—such as a train delayed by weather. A slowdown at one point can interfere with many other schedules. However, a university research project called Arrival resulted in new scheduling software that efficiently routes trains, even when disruptions happen. The system can handle almost instantaneous rescheduling as disruptions arise. It adjusts timetables, platform allocations, and staff scheduling. The system has been tested in the Netherlands, Germany, and Switzerland. In Berlin, the average wait time was cut in half from four minutes to two. In a test in Italy, the system created a 25 percent reduction in train delays in Palermo and Genoa.

Adapted from *Daily Mail Reporter*, "Brussels: We Can Make the Trains Run on Time," July 16, 2010.

have the ability to learn on their own. In practice, the learning stage is the most difficult component of building a neural network. Most times the designer has to understand the problem and provide hints to the network, along with good sample data. In many ways, training a neural network uses basic properties of statistics related to data sampling and regression.

Machine Vision

Machine vision has many uses in manufacturing environments. Machines are used in optical character recognition, welding and assembly, and quality control. Mechanical sensors have several advantages over humans. They do not suffer from fatigue, they can examine a broader spectrum of light (including ultraviolet and infrared), and they can quickly focus at many different levels (including microscopic).

On the other hand, traditional computer systems are literal in their vision. It is hard for computers to compare objects of different sizes or to match mirror images. It is hard for machines to determine whether differences between objects are minor and should be ignored or if they are major distinguishing features.

Say you are shown a picture of your instructor, and someone adds or subtracts features to it, such as bigger eyebrows, longer hair, or glasses. In most cases, you would still recognize the face. Computers would have difficulty with that problem because they see pictures as a collection of dots (or lines). How does the computer know which changes are important and which are minor?

Machine vision systems are improving rapidly but still have a way to go to become commonplace. For example, companies are working on applications in facial recognition and facial expressions, body tracking (so you can use your hand as a computer pointer), visual tracking of handwriting for use in computer tablets, product inspections for defects, and shape identification.

See what happens when you give a computer the first set of instructions, but it does not hear the commas correctly and thinks you said the second line:

(1) Copy the red, file the blue, delete the yellow mark.

(2) Copy the red file, the blue delete, the yellow mark.

Consider the following sentence, which can be interpreted by humans, but would not make much sense to a computer that tries to interpret it literally.
I saw the Grand Canyon flying to New York.

Or the phrase epitomized in a grammar book:
The panda enters a bar, eats, shoots, and leaves.

Figure 9.29

There are inherent problems with voice recognition. Punctuation and implicit meaning are two difficult areas. Even communication between people has frequent misinterpretations.

Language Comprehension and Translation

Related to voice recognition is the issue of language comprehension, or the ability of the computer to actually understand what we are saying. Technically the two topics are separate, since it might be possible to have a machine understand what we type onto a keyboard. Language comprehension exists when the machine actually understands what we mean. One test of comprehension would be the ability of the computer to carry on a conversation. In fact, Alan Turing, a British pioneer in the computer field, suggested the **Turing test** for computer intelligence. In this test, a human judge communicates with a machine and another person in a separate room. If the judge cannot determine which user is the machine and which is a person, the machine should be considered to be intelligent. Some people have tested this concept (using specific topics). Other people have noted that perhaps you do not have to be intelligent to carry on a conversation.

The Loebner Prize is a Turing test that runs every year to evaluate how close programs are coming to meeting the challenge. Several of the challengers have set up Web sites (e.g., www.alicebot.org, and www.jabberwacky.com) so you can play with these conversational bots. Why do you care? First, you can buy this technology to use it in business applications—such as building Web sites that answer basic questions from customers. Second, as a customer, you might want to learn to recognize when you are dealing with a machine instead of a human.

Language comprehension would be useful because it would make it easier for humans to use computers. Instead of needing to learn a language such as SQL to access data, imagine being able to get answers to questions asked in English (or some other **natural language**). Of course, any natural language has its limitations. The greatest danger with language comprehension is that the machine will interpret your question incorrectly and give you the “right” answer to the “wrong” question. Figure 9.29 provides a simple illustration of the complexities of language comprehension. The first example involves the use of punctuation. A misinterpretation of the command can result in deleting the wrong file. Similarly,

Figure 9.30

A practice Jeopardy match with IBM's Watson and two of the best human players. In three days of televised matches, Watson handily beat the humans.
Source: http://www.youtube.com/watch?v=12rNbGf2Wwo

interpretation of a natural language involves understanding some basic concepts, such as the fact that the Grand Canyon cannot fly.

Ultimately, translating languages requires an understanding of the underlying meaning of a sentence and paragraph. Early translation systems can convert a word from one language into another. For example, check out babel.altavista.com, or any of the other online translators. Some high-end translation systems have been developed that do a better job by at least recognizing common phrases and idioms. Yet even these systems need to be supplemented by human translators.

IBM raised the bar for natural language comprehension in 2011. The company created the Watson system designed to understand language, using statistics to help it understand questions and provide answers. As shown in Figure 9.30, to show off the system, it played several games of Jeopardy against two of the best human opponents. The Watson system won easily. The use of statistical analysis is still debated in computing circles. In language, it provides the benefit of letting the system decide not just an answer but how confident it is in the answer. Just as in real life, if a statement is not clearly understood, the computer can ask for clarification. IBM has suggested that the system can applied to many types of business and research problems, including medical research.

Robotics and Motion

Modern manufacturing relies heavily on robots, and the capabilities of robots continually increase. Most existing robots are specialized machines that perform a limited number of tasks, such as welding or painting. In many firms, there is little need for a general-purpose robot that can "do everything." However, one area that

Reality Bytes: Dealer Services Corp

Dealer Services Corp. provides financing to about 10,000 car dealers. The company essentially writes separate loans for every vehicle, providing individual conditions and payback schedules. The system generated huge amounts of data. Technically-oriented analysts within the company developed their own spreadsheets to analyze data, but they often had inconsistent or bad data. The company purchased WebFocus as a tool to provide centralized access to data and business intelligence. The system largely focused on interactive reporting and charting. When the recession hit in 2008 and 2009, the company managers were able to see results in individual markets and by dealers to see that inventory was not selling. CIO Chris Brady noted that these key insights warned the company early that the recession was going to cause severe problems. So the company tightened lending standards and increased financial reserves. It was also able to provide advice to dealers, such as "stop buying SUVs, they aren't selling." Brady also noted that "We definitely reduced our losses from bad loans and didn't start to see a negative effect until the very end of 2008."

Adapted from Elizabeth Horwitt, "Self-Service BI Catches On," *Computerworld*, December 13, 2010.

remains troublesome is the ability of machines to move. Making a machine that can navigate through an unknown or crowded space is especially difficult. Some work is being done in this area. Liability is a major problem when robots attempt to move among people.

In 2000, Honda built a humanoid robot (Asimo) in Japan. The Asimo robot has two legs and arms. Its most impressive feature is the ability to walk like a human, including up and down stairs. It can also shake hands and hand objects to people. The multi-million-dollar project is the latest step of a 16-year evolution.

In March 2004, DARPA held the first Grand Challenge contest for automated vehicles in the Mojave Desert. The challenge was to create an automated vehicle that could drive itself across 189 miles of desert hitting about 1,000 GPS waypoints in about 10 hours. The Carnegie Mellon $3 million Red Team vehicle (a modified Hummer) made it 7.4 miles. In fall 2005, the contest was run a second time—and five vehicles finished the 132-mile course. A huge gain in a year and a half, and the winning Stanford vehicle was created in less than a year! DARPA announced a new contest for November 2007, where the vehicles will have to drive through city traffic. Congress has mandated that 30 percent of Army vehicles be automated by 2015, and DARPA is attempting to stimulate innovation with these prizes.

As computer processors decline in size and price, it becomes easier to build intelligent mobile systems, making it possible to build robots and automated vehicles. Major automobile manufacturers are working on slightly less automated systems to provide assistance to drivers. For example, crash-warning sensors can automatically apply brakes for an inattentive driver. Slide-control systems to prevent spinouts on sharp turns have been installed in luxury vehicles for several years.

Reality Bytes: Guessing for Holiday Sales

Holiday shopping is critical for many U.S. businesses. At the high end, hobby, toy and game stores make 34 percent of their sales in November and December. Even for other retailers, the season is critical. Almost 20 percent of all retail sales value arises in the last two months of the year. Department stores hit 25 percent. The one sector that is lower than then one-sixth time frame is car dealers. The problem is that retailers need to order products and have them on the shelves well before the season begins. Consequently, retailers need good forecasts of both the expected sales level and details on which products are going to be hot.

Adapted from Phil Izzo, "Number of the Week: Outsize Importance of Holidays for Retailers," *The Wall Street Journal*, December 24, 2010.

Machine Intelligence

What would it take to convince you that a machine is intelligent? The Turing test has been proposed as one method. Many other tests have been proposed in the past. At one time, people suggested that a machine that could win at chess would be intelligent. Today's chess-playing computers have beaten even the top human players. Another test proposed was the ability to solve mathematical problems, in particular, the ability to write mathematical proofs. An early AI program created in the 1950s could do that. Today, for a few hundred dollars, you can buy programs or even small calculators that manipulate mathematical symbols to solve equations.

Some people have suggested that intelligence involves creativity. Creativity is probably as hard to measure as intelligence. Even so, examples of computer creativity abound. A few years ago, a programmer developed a system that created music. The interesting feature of the program was that it allowed people to call on the phone and vote on the music. The computer used this feedback to change its next composition. Not only was the computer creative, but it was learning and adapting, albeit in a limited context. Today, you can buy software that creates music or plays background to your solo.

Although business applications to much of this current research is somewhat limited, there are two main reasons for staying abreast of the capabilities. First, anything that makes the computer easier to use will make it more useful, and these techniques continue to improve. Second, you need to understand the current limitations to avoid costly mistakes.

DSS, ES, and AI

What are the differences between DSS, ES, and AI systems? The differences among decision support systems, expert systems, and artificial intelligence can be confusing at first. Take a simple problem and see how a computer system based on each method might operate. A common financial problem is to determine how much money to lend to customers. Any firm that grants terms to customers—not just financial institutions—must make this decision. Figure 9.31 discusses the differences among a DSS, ES, and AI approach to the inventory problem.

Decision Support System	Expert System	AI/Nerual Network
Loan Officer Data: Income Existing loans Credit report Model: Lend in all but the worst cases. Monitor for late and missing payments. Output: Name Loan #Late Amt Brown 25,000 5 1,250 Jones 62,500 1 135 Smith 83,000 3 2,435 …	ES Rules What is the monthly income? 3,000 What are the total monthly payments on other loans? 450 How long have they had the current job? 5 years … Should grant the loan since there is a 5% chance of default.	Determine Rules Data/Training Cases Loan 1 data…paid Loan 2 data… 5 late Loan 3 data… bankrupt Loan 4 data… 1 late Neural Network Weights Evaluate new data, make recommendations.

Figure 9.31

Comparison of techniques for a loan. A DSS can display background data for a loan officer and can also monitor customer payments. An ES could help managers decide if they should make the loan by evaluating more complex rules. An AI such as a neural network can analyze past loans and determine the rules that should be used to grant or deny future loans.

In a relatively simple system, the computer would retrieve data about the customer and the prior loans to that customer. Historically, loan officers used basic data and personal factors to make the lending decision. In some instances, these rules of thumb led to problems—with bad decisions and sometimes discrimination. The DSS could also be used to monitor existing loans and payments. As part of a transaction-processing system, it can notify managers when customers continually make late payments and help identify problem loans.

To improve consistency and reduce the decision time, many firms have moved to expert systems to help evaluate loans. Statistical analysis of prior loans is used to establish a set of rules that are coded into the ES. In some cases, the ES can then be operated with push-button phones or over the Internet. In straightforward cases, the ES can make the final decision and approve the loan. In more difficult situations, the preliminary results and data can be forwarded to a human loan officer to factor in personal judgment and factors not considered by the ES.

Of course, the value of the ES depends heavily on the accuracy of the underlying rules (and the supplied data). These rules might change over time or as economic conditions change. A neural network can be used to examine the prior loans automatically to identify the factors that predict successful and unsuccessful loans. Once these factors are identified, they can be coded into the ES to automate the decision process. In this situation, the AI/neural network takes the place of (or supplements) the decisions of the human expert.

The Importance of Intelligent Systems in e-Business

How can more intelligent systems benefit e-business? Disintermediation is a primary aspect of e-business. Businesses can interact directly with customers, with less need for middle levels such as retail stores. However, these middle levels often existed because they provided more explanations and support to customers. If you remove that level, how are you going to deal with thousands or millions of customers? If you have to hire hundreds of workers to answer customer questions, you will lose most of the potential benefits of disintermediation. One of the solutions to this problem is to implement more intelligent systems that can provide automated support to customers.

In many ways, the Internet adds complexity to the daily lives of customers and managers. The Internet provides access to huge amounts of data—and it is growing constantly. The growth adds more data, but it also means that the availability and use of information is constantly changing

Agents

A recent application of AI techniques has arisen in the context of the Internet. A key issue of the Internet is searching for data. Although the Internet dramatically improves communication, there are problems with maintaining the "interpretation" of the information from various systems. Originally, most data on the Web was stored as standard pages of text using HTML. Search engines would simply scan these pages and build searchable indexes.

Increasingly the Internet is being used to store and transmit objects composed of data, pictures, spreadsheets, sounds, and video. From a pure transmission standpoint, any object can be decomposed into raw data bits and sent between computers. Where we run into problems is searching for the objects. Consider a simple example where you want to find a new printer, so you search the Internet for prices. Today, many vendors store the product descriptions and prices in a database, and then build the HTML page on demand when you go to the site. Since the page is not static, the search engines do not index it.

One solution to this problem is to create software agents. **Agents** are object-oriented programs designed for networks that are written to perform specific tasks in response to user requests. The concept of object orientation is important because it means that agents know how to exchange object attributes, and they have the ability to activate object functions in other agents. The tasks could be simple, such as finding all files on a network that refer to a specific topic. One key feature of agents is that they are designed to communicate with each other. As long as your agent knows the abilities or functions of another agent, they can exchange messages and commands. General Magic was a pioneering company that created a standard programming language for agents. With this language, agents can transfer themselves and run on other computers. Agents also have a degree of "intelligence." They can be given relatively general commands, which the agents reinterpret and apply to each situation they encounter.

Consider an example illustrated by Figure 9.32. You have been working hard and decide to take a vacation. You want to go to a beach but do not have much money to spend. You are looking for a place where you can swim, scuba dive, and meet people at night. But you also want the place to have some secluded beaches where you can get away from the crowds and relax. You could call a travel agent and buy a package deal, but every agent you call just laughs and says that next

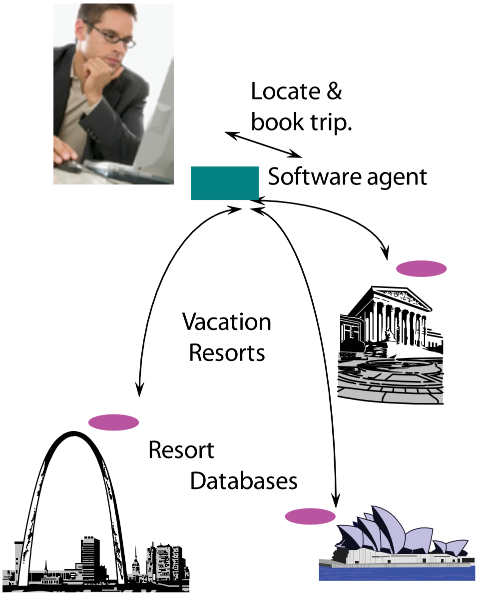

Figure 9.32

Software agents. A personal software agent might be used to book a vacation. It would take your initial preferences and communicate with other agents to find sites that matched your preferences. It might also be able to negotiate prices with competing resorts.

time you should call three months ahead of time instead of only three days ahead. You suspect that a beach resort probably has last-minute cancellations and you could get in, but how do you find out? Thousands of possibilities exit. If all of the resort computers had automatic reservation agents, the task would be fairly easy. You would start an agent on your computer and tell it the features you want. Your agent sends messages to all of the automated resort agents looking for open spots at places that matched your features. When your agent finds something close, it brings back details and pictures to display on your screen. When you decide on a resort, the agent automatically makes the reservations.

You might argue that some of the existing travel search sites come close to this example. On the surface, yes, sites such as PriceLine offer a few of the elements. However, the technology and methods beneath the surface are completely different and largely rely on humans to set prices. The travel industry is likely to be one of the first to incorporate these technologies—largely because many of the reservation systems are already linked together at some level. At some point, the negotiating agent needs access to the data. Today, these decisions are made at central search engine sites that have access to the data on millions of sites. In a true agent-based system, the search engine would not need direct access to the data, but would communicate with the agents located at each site.

Notice three important features of software agents. First, the agents need to know how to communicate. It is not as simple as transmitting raw data. They must understand the data and respond to questions. Second, imagine the amount of network traffic involved. In the vacation search example, your agent might have to contact thousands of other computers. Now picture what happens when a thousand other people do the same thing! Third, all of the agents are independent. You, as well as other computer owners, are free to create or modify your own agent. As long as there are standard methods for agents to exchange attributes and activate functions, they can be modified and improved. For instance, you might program your agent to weight the vacation spots according to some system, or you might teach it to begin its search in specific locations.

Programmers have begun to incorporate expert system and other AI capabilities into these agents. By adding a set of rules, the agent becomes more than just a simple search mechanism. The more complex the rules, the more "intelligent" it becomes, which means you have to do less work. In fact, software agents have the potential to dramatically increase the research in AI. Currently, because of limited standards and the difficulty of creating them, there are few examples of useful agents. As increasing numbers of people use agents and begin demanding more intelligence, it will become profitable for researchers to work harder at building reliable, intelligent software.

Support and Problem-Solving Applications

Increasingly, your customers want personalized attention to help in both selecting products and solving problems. Yet it is expensive to provide individual personal support to every customer. Instead, firms are developing expert systems and other intelligent applications to help customers with a more personalized touch. For example, look at Amazon.com's recommendation system. It began with books but has been expanded to most of their products. As you purchase items at Amazon.com, the system gives you a list of similar products that you might be interested in. For instance, if you purchase several science fiction books, it will suggest new releases of similar books. The system can increase sales because it helps show customers items that they might not have found otherwise.

More complex products can benefit from more sophisticated expert systems that help analyze customer needs and help configure the correct components. For example, a computer vendor could build a system that asks questions to help identify the applications that a customer will run. It could then suggest specific enhancements such as adding RAM or a second disk drive to improve performance.

Similarly, many firms are building expert systems to help customers with problems. If a customer has a problem installing a new product, he or she can turn to the Web site. The system asks questions to identify the problem and then make suggestions. The advantage of the expert system is that it is available 24 hours a day, can solve most of the easy problems, and is less embarrassing to customers who might think their questions are too "silly" to ask a human troubleshooter.

Intelligent systems can also be useful for B2B and other forms of e-business. The systems might analyze past purchases and suggest new products, or automatically analyze sales patterns and help managers develop new products and close out unprofitable lines. They can be used to develop automated ordering systems that predict customer demands, schedule production, and generate automated sales orders and payments.

It can be difficult to develop these applications, but firms that build powerful systems will attract customers and increase the level of sales to each customer. Ultimately, these systems could be the primary reason people switch to buying items over the Internet.

Cloud Computing

How can cloud computing be used to analyze data? Cloud computing offers some options for analytical tools—but the process is only beginning. Some decision tools require relatively large amounts of computation and storage capabilities. In some cases, these tools could be run on large parallel processing supercomputers. Some research institutions provide online access to supercomputers—but they are traditionally used for scientific research. For more typical business research, it should be possible to use multiple processor systems leased from cloud computing providers. Companies such as Amazon provide access to multiple processors based solely on the amount of time used. Consequently, researchers can develop algorithms, test procedures, and fine-tune the analysis. Then, when the main data sets are loaded, additional processors can be leased for a short time to process the main results. With minimal fixed costs, cloud computing makes it possible to tackle large problems at relatively low costs.

It is likely that companies will offer data analysis as package services through online systems. Currently, several market research firms will conduct studies, format data, and run analyses for other companies. As more data becomes accessible through online systems, it will become possible for cloud-based companies to offer direct access to analytical services. For instance, if you want to offer services similar to Amazon or Netflix to show customers similar items purchased by customers, a company could automatically read your sales data, perform the forecasts, and display the results using links across the Web. You would not have to develop the algorithms or purchase high-speed computers to handle the computations. On the other hand, you might have less control over customization or modification of the algorithms. Leased systems might not provide a competitive advantage, but they might make it easy for your company to add desirable features to your Web site at relatively low cost.

Summary

Managers make many different decisions. Every business discipline builds models to help people analyze problems and make decisions. Without models and tools, people rarely make good decisions. Many tools have been created to help you analyze data and make decisions. In many cases, you will want to build a data warehouse to retrieve and organize the data. You can build a DSS, building a model in a spreadsheet or more advanced statistical analysis tool. A good DSS extracts the needed data from the database, evaluates the model, and displays the results in a form that helps managers visualize the problem and quickly choose a solution. The biggest difficulty with DSS tools is that you generally need to be trained to create the models and understand the results.

More sophisticated data mining tools can be employed to semiautomatically search for correlations and other relationships within the data. Common regression, clustering, and classification tools use statistical measures to identify the importance of various attributes and to build forecast equations. Market basket analysis identifies items that customers purchase together. The information provided by this analysis is a powerful tool for cross selling related products. Data mining tools are powerful techniques, but the results are not always useful, and you must carefully evaluate the implications of the results to ensure that they are realistic and repeatable.

Expert systems provide a different level of decision automation. They analyze data based on rules defined by an expert. They can handle complex and missing data. They are particularly useful at helping novices reach a better decision. They are excellent tools for specialized, narrowly defined problems. The biggest concern arises when someone tries to build an ES for a problem that is too large, too variable, or unstructured. The ES is not likely to work, and it will cost money and time to build. Worse, the proposed decisions might be nonsensical.

Scientists are continually working on ways to automate even more decisions. Neural networks are incredible tools for analyzing pattern data. They have been successfully applied to problems ranging from speech recognition to lending analysis. You can purchase software that will quickly build a neural network to analyze data. One drawback to neural networks is that the relationships tend to be highly nonlinear and difficult to interpret. So, you might reach an answer, but you might not be able to explain the relationships. Other tools that are increasingly useful in business applications are machine vision, robotics, and language comprehension.

Automated tools can help you provide better service to customers and suppliers, particularly when the intelligent systems can be reached online. If you can help a customer solve a problem 24 hours a day without having to pay hundreds of humans, you can gain happier customers with little additional cost. Several companies are working on building intelligent agents that will enable customers and vendors to interact automatically by following rules that you specify.

A Manager's View

It is hard to make good decisions. You need fast access to huge amounts of data, the ability to evaluate various models, and a way to visualize the problem and the solution. Various levels of tools are available to help. The tools provide different types of intelligence and support. As a manager, you need to understand the context of the problem and know which tools can be applied to solve a problem. You also need to be enough of an expert to recognize when a system provides useless or bad answers so that you can avoid disasters.

Key Words

agent
artificial intelligence (AI)
data marts
data mining
data warehouse
decision support system (DSS)
decision tree
expert system
extraction, transformation, and loading (ETL)
geographic information system (GIS)
global positioning system (GPS)
knowledge base
knowledge engineer
market basket analysis
metadata
models
natural language
neural network
online analytical processing (OLAP)
optimization
parameters
prediction
rules
simulation
Turing test

Web Site References

Expert Systems and AI Tools	
A.L.I.C.E. conversation	Alice.pandorabots.com
CLIPS (Started by NASA)	clipsrules.sourceforge.net
ExSys (Commercial)	www.exsys.com
International NeuralNetwork Society	www.inns.org
Jess (Java)	www.jessrules.com
Mathworks neural network toolbox	www.mathworks.com/products/neuralnet
Machine Vision	
CalTech	www.vision.caltech.edu
Carnegie Red Team	www.cs.cmu.edu/~red/Red
British Machine Vision	www.bmva.org

Review Questions

1. Why is it so important to build models and analyze data using a scientific process?
2. What is the purpose of a data warehouse?
3. How is an OLAP cube browser better than using queries?
4. What is the role of a manager in a DSS?
5. How is a GIS used to answer business questions?
6. How can data mining help you make better decisions and what are some of the primary techniques?
7. What is an expert system and what are the characteristics of the problems that it is designed to solve?
8. What tools are available to help businesses construct expert systems and what features do they provide?

9. What types of problems are best suited for a neural network?
10. How would a fully autonomous vehicle alter business?
11. What are the differences between a DSS, ES, and AI features?
12. Why are intelligent systems so important for e-business?
13. How could cloud-based computing and the Internet be used to add intelligence to business applications and decisions?

Exercises

1. Work through at least two of the examples on the Exsys Web site. What features do the two examples have in common? Give an example (unrelated to any on the Web site) of a business problem or something from your life that could benefit from an expert system.
2. An HR manager wants to develop an expert system to evaluate potential employees applying to work for a job in your department. Assuming the job can be performed by a business-school intern, list some of the questions you would ask potential employees. Create a decision tree to evaluate the basic questions. Try to generate three possible outcomes: acceptable, unacceptable, and personal interview to decide. Your goal is to reduce the number of people needing a personal interview.
3. Interview an expert in some area and create an initial set of rules that you could use for an expert system. If you cannot find a cooperative expert, try researching one of the following topics in your library: fruit tree propagation and pruning (what trees are needed for cross-pollination, what varieties grow best in each region, what fertilizers are needed, when they should be pruned); requirements or qualifications for public assistance or some other governmental program (check government documents); legal requirements to determine whether a contract is in effect (check books on business law).
4. Describe how you could use data mining tools to help you find a new vehicle to buy. How well do the car-buying sites such as Edmunds perform these tasks for you?
5. Obtain an expert system (e.g., Jess and CLIPS are free). Create a set of rules to evaluate a simple request for a car loan. If you do not have access to the tool, at least build the decision tree.
6. Identify a problem that would be well suited for a neural network. Explain how the system would be trained (e.g., what existing data can be used?). Explain why you think the problem needs a neural network and what benefits can be gained.
7. For the following problems identify those that would be best suited for an expert system, decision support system, or a more advanced AI system. Explain why.
   a. A financial manager in a bank wants to find the best investments for her clients.
   b. A venture capital investor wants a system to do a first-pass evaluation of new business plans that are submitted via a Web site.

c. A plant manager wants to be able to forecast when equipment will need special maintenance. Currently, the engineers use a fixed schedule, but some products and operations seem to cause more problems than others.

d. A fast-food restaurant chain wants to help managers determine the appropriate number of employees to schedule for each day in the year.

e. A cell-phone provider wants an application on its Web site to help customers select a cell phone that best meets their needs.

f. A company has to select a city to host a new factory which happens about once a decade.

g. A marketing department wants an application to help them set the best prices for new products in different markets.

8. Examine existing travel sites, including aggregators such as Kayak.com. What features could be added to increase the intelligence level to assist consumers?

9. Use a spreadsheet to create the example from the Human Resources Management example. Fill in the market adjustment column so that raises match the performance appraisals. Remember, total raises cannot exceed $10,000.

DeptStor.mdb

Department	Customers/month
Clothing—Children	180
Clothing—Men	150
Clothing—Women	180
Electronics	200
Furniture	150
Household	250
Linen	300
Shoes	300
Sports	400
Tools	340

You are a midlevel manager for a small department store. You have collected a large amount of data on sales for 2012. Your transaction system kept track of every sale (order) by customer. Most customers paid by credit card or check, so you have complete customer data. Walk-in customers who paid cash are given a separate customer number, so you still have the sales data.

You are trying to determine staffing levels for each department. You know that the store becomes much busier during the end-of-the-year holiday season. For summer months, you have thought about combining staff from the departments. From conversations with experienced workers, you have determined that there is a maximum number of customers that can be handled by one person in a department. These numbers are expressed as monthly averages in the table.

You are thinking about combining workers from some of the departments to save on staffing—especially over the spring and summer months. However, working multiple departments makes the sales staff less efficient. There are two considerations in combining staff members. First, if any of the departments are reduced to a staff of zero, sales in that department will drop by 10 percent for that month. Second, total staffing should be kept at the level defined by the monthly averages. If average staffing (total across all departments) falls below the total suggested, then sales in all departments will fall by 2 percent for each tenth of a percentage point below the suggested average.

10. Using the database and a spreadsheet, determine how many workers we need in each department for each month. Present a plan for combining departments if it can save the company money. Assume that sales members cost an average of $1,000 a month. Two queries have already been created by the MIS department and are stored in the database: SalesbyMonth and SalesCountbyMonth. The first totals the dollar value; the second counts the number of transactions.

11. Write a report to upper management designating the appropriate sales staff levels for each department by month. Include data and graphs to support your position. (*Hint*: Use a spreadsheet that lets you enter various staffing levels in each department in each month, and then calculate any sales declines.)

12. Create a PivotTable for the company that enables managers to evaluate sales by employee by department by month.

Technology Toolbox

13. Create the PivotTable report for Rolling Thunder Bicycles. Briefly summarize any patterns or problems you identify.

14. Using the Rolling Thunder Bicycles query, create a PivotChart and compare sales of the different models over time. Identify any patterns that you see.

15. Research an alternate cube browser (such as SQL Server or Oracle) and explain how it is different from the Excel PivotTable. If you have access to the tool, build a small example.

16. Compute the average number of days it takes to build a bicycle (ShipDate – OrderDate) for each month. Import the data into Excel and forecast the trend. First, forecast it based on all of the data. Second, forecast it for three time periods: (a) the early years, (b) the middle years, and (c) the most recent years. Look at the initial chart to estimate the breaks between these three sets, or just divide it into three equal-size groups if you do not see any good break points. Comment on any differences or problems.

17. Using federal data (start at www.fedstats.gov), compute a regression analysis of Rolling Thunder sales by state by year compared with at least population and income.

18. Choose a company and identify the primary dimensions that you would create in an OLAP cube if you were a manager at that company. Hint: Focus on your major area.

Teamwork

19. Have each person find and describe a problem that could benefit from a GIS. Make sure it needs a GIS, not just a mapping system. Combine the results and compare the types of problems to identify similarities.

20. Build a decision tree to test whether a specific student next year can take this class.

21. Many classes fill up early and the dean keeps a wait list for adding students at the last minute. Build a decision tree to evaluate each student on the list to determine priority for adding that specific person to the class.

22. Each person should identify a problem that could be solved with a rules-based expert system. Combine the results and compare the types of problems to identify similarities.

23. Using Rolling Thunder Bicycles, have each person forecast the sales by one model type for six months. Combine the individual model results and compare this value to the forecast based only on total sales.

24. Select an economic data series such as personal income (check www.fedstats.gov). Place members into one of three subgroups. Have each group forecast the series using a different methodology. Compare the results. If you have sufficient data, leave out the most recent data, and then forecast those values and compare the forecasts to the actual.

25. Choose a publicly-traded company and collect basic quarterly financial data for the company for at least 10 years. It is easiest if you assign specific years to each person. Put the data into a spreadsheet or simple database. Create a PivotTable and Pivot Chart to examine the data. Create a report containing some of the charts and tables and any conclusions you can make.

26. Have each person find a problem that could benefit from a neural network. Describe how the system would be trained. Combine the individual comments and identify any commonalities.

Rolling Thunder Database

27. Identify shipments where receipts do not match the original order. Provide a count and value (and percentages) by supplier/manufacturer.

28. Analyze sales and discounts by employee and by model type. Are some employees providing higher discounts than others? Are we discounting some models too much or not enough?

29. The company wants to create an online ordering system. Create a decision tree to help novices select the appropriate bicycle and components. If necessary, consult with a friend or relative who can be considered a bicycle expert.

30. Use queries to extract sales data by model type and month. Use a spreadsheet to forecast the sales of each quantity of model type by week for the next year. *Hint: Use Format([OrderDate],"yyyy-mm") to get the month.*

31. What pattern-matching types of decisions arise at Rolling Thunder that could benefit from the use of neural networks?

32. What aspects of customer service might be automated with expert systems? What are the potential advantages and disadvantages?

Additional Reading

Barabba, Vincent and Gerald Zaltman, *Hearing the Voice of the Market*, Harvard Business Press: Cambridge, MA, 1991. [Overcoming design biases at GM.]

Franses , Philip Hans, *Time Series Models for Business and Economic Forecasting*, 1998, Cambridge University Press: Port Chester, NY. [An introduction to forecasting.]

Dangermond, Jack and Adena Schutzberg, "Engineering, geographic information systems, and databases: a new frontier," *Journal of Computing in Civil Engineering*, July 1998 pp. 121-122. [Uses of GIS.]

Krill, Paul, "Analytics redraw CRM lines," *Computerworld*, December 3, 2001. [A summary of vendors developing products to analyze customer data.]

Neil, Stephanie, "Blue Cross dissects data to improve care," *PC Week*, February 8, 1999. [BCBS uses data warehouse and online application processing DSS to improve service.]

Wilkinson, Stephanie, "PC apps help to take a byte out of crime," *PC Week*, February 23, 1998. [GIS and imaging tools help police predict and solve crimes.]

Cases: Financial Services Industry

The Industry

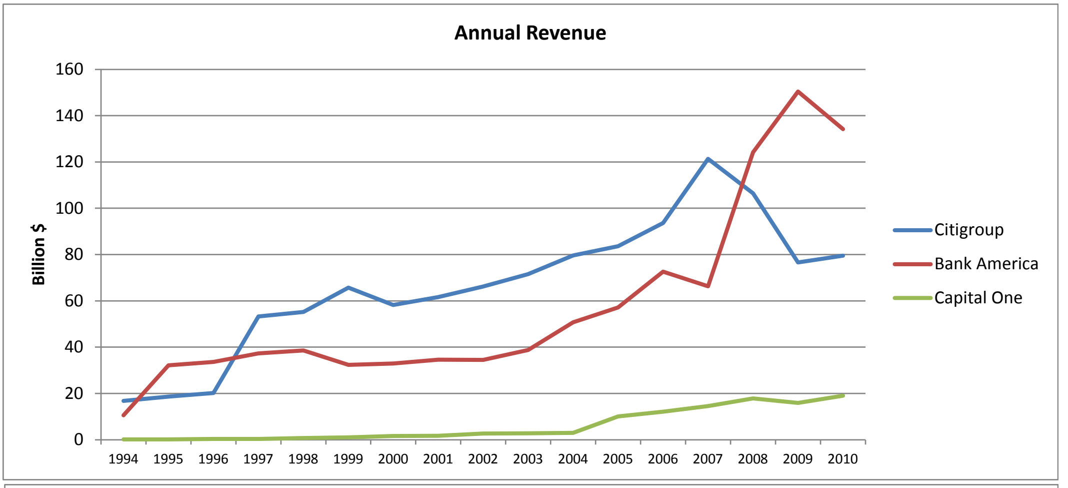

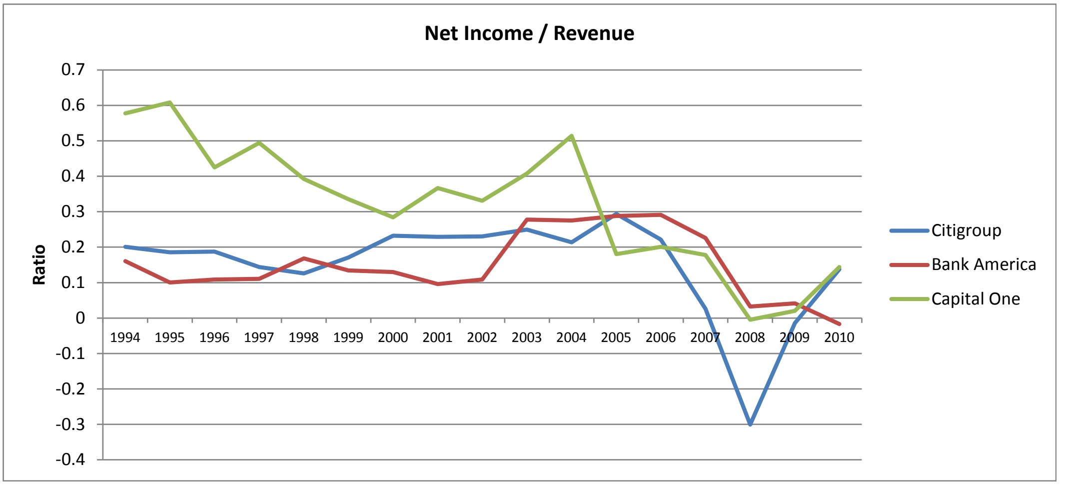

The financial industry is interesting. It is huge and everyone interacts with the trillions of dollars a day handled by the global financial system. It is also complex, with many different types of firms involved, and constantly creating new financial tools. At one point, the financial industry was driven by banks. Today, banks still play important roles, but where do people put their money? In the stock market or other investments. So brokerage firms have a strong role. How do consumers pay for things? With debit cards and credit cards. So the card-processing sector, led by Visa and MasterCard, plays a huge role in the industry. As the national laws have changed, mostly relaxed, over the past few years, the industry has become even more complex. Banks can once again sell stocks and investments, as well as insurance. And brokerage firms and insurance companies can perform banking functions.

Making Money

How do brokerage firms make money? They used to make a profit on transaction commissions, as much as several percentage points per trade—amounting to hundreds or thousands of dollars every time a client bought or sold a stock. With the advent of online trading, pushing the discount brokerage firms, commissions fell to fixed rates of $10 to $20 per trade, a tiny fraction of what they were and under 10 cents a share. The big firms tried to convince customers that their higher fees were worthwhile because they also provided investment advice. But with tainted advice, and a huge amount of data available free to customers online, the full-price firms suffered.

How do banks make money? Banks make money in two ways: interest rate spreads and fees. Interest rate spreads are the difference between the rates the banks charge borrowers and the amount the bank has to pay to obtain the money. Banks select different lending markets (consumer versus business) and have diverse ways of obtaining funds (deposits, loans from other banks or the Federal Reserve). Likewise, fees vary depending on the type of bank. Consumer-oriented banks receive fees from customers (such as checking account fees). Business or investment banks charge fees for more sophisticated services.

How do credit card companies make money? Here the answer is trickier, because several types of companies are involved. The transaction-processing organizations (Visa and MasterCard) receive a fee for every transaction handled through the clearing system. However, banks ultimately issue the cards and are responsible for the money. A consumer bank issues a card to customers and is responsible for paying vendors for legitimate transactions. The consumer bank makes money on consumer fees (usually low today) and interest charges to customers on outstanding balances. It also uses the cards as a marketing tool to attract customers and encourage them to purchase other services. The merchant bank is responsible for ensuring that the merchant is legitimate. It makes money by taking a percentage cut of the merchant's credit card (or debit card) sales. In 2011, the U.S. government joined Europe and imposed caps on merchant fees for debit cards (not credit cards). So banks began cancelling debit card promotional programs and scrambled to find other ways to charge fees.

Brokerage Firms

More than three-fourths of American's liquid assets ($12 trillion) are in stocks, bonds, and money market funds [Revell 2002). How safe are these assets? Do you trust your broker? Do you trust the advice from brokers? In April 2003, forced by a lawsuit led by New York attorney general Eliot Spitzer, 10 of the largest Wall Street investment banks paid fines of $1.4 billion (Nocera 2004). The money was a settlement for providing misleading investment analysis. The brokerage firms usually wrote glowing reports of companies—urging people to buy stock in specific companies. Many of the reports were written because the firms were vying for investment banker contracts with the same companies. In exchange for glowing reports, the companies provided more business (and fees) to the brokerage firms. James Freeman, a former research director, notes that analysts "went from having investment banking deals pushed on them in the 1980s to becoming the greedy pigs at the center of it" (Nocera 2004).

Credit Cards

At one time, the credit card industry was neatly divided into three main segments: the upper end controlled by American Express, the middle tier targeted by Citicorp, and the subprime market pursued by Providian. As a fourth group, MBNA targeted a wide swath of customers with its affiliate cards. Most card companies gained market share by blindly sending 5 billion solicitations a year to U.S. consumers (Gross 2002). Capital One grew by using targeted marketing and thousands of marketing tests with different groups of potential customers. Nonetheless, with competition at a peak for high-end customers, profits were squeezed. Companies were increasingly tempted to go after the subprime market—where they could charge higher interest rates in exchange for greater risk. But as the economy turned down and interest rates increased, this strategy results in greater losses when customers do not pay their bills or even file for bankruptcy (Dugas 2002). Providian lost 90 percent of its stock market value from mid-2001 to mid-2002. Capital One also faced greater losses, but not as large as many other companies. The 2003 settlement with Visa and MasterCard ultimately affected profits for the industry. Wal-Mart had sued the two credit card companies to break up their power to force merchants to accept higher costs for debit cards. The win by the merchants reduces the power of the card companies and banks, and should ultimately reduce fees paid by merchants (www.cardweb.com).

Banks

Citibank (or the parent company Citigroup) is one of the largest banks in the world. In 2003, it earned $17.9 billion in profits, the most ever by a single bank (Economist 2004). In 2010, net income was $10.6 billion (2010 Annual Report). Over the years, it has suffered several problems along with other banks, including from the huge losses on international loans to South America in the 1980s to the investment banking scandals in the 1990s. The company agreed to pay $400 million as part of the $1.4 billion investment banking settlement in 2003, the largest payment by any of the companies involved. Also like other banks, the company grew through acquisitions of other banks (Primerica in 1988), insurance companies (Travelers in 1993), brokerage firms (Shearson and Salomon in 1998), and credit cards (Sears in 2004) (Rosenberg 2003). In fact, the last decade of banking in the United States can be best characterized by the mergers. A handful of banks control a huge number of deposits (Citibank, Bank of America, J.P. Morgan Chase, Wells Fargo, and Wachovia) along with a few large regional banks (Stone 2003). Ultimately, one limiting factor to mergers is that federal law currently prohibits a single U.S. bank from holding more than 10 percent of the consumer deposits. Banks have also been one of the big beneficiaries of the loose federal monetary policy of the early 2000s—making it easy to borrow money at almost no cost.

The entire banking industry was rattled in the 2008 recession and housing market crash. All of the major banks in the U.S. and many banks in other countries were bailed out by governments. A few banks were allowed to fail, and a few others merged with larger banks. The banks most involved in shaky mortgage loans were the hardest hit and the first to go. Brokerage firms were also forced to merge in order to survive. For example, WaMu (savings bank) was eventually sold to JP Morgan Chase, and Merrill-Lynch (brokerage) was picked up by Bank of America. The crisis revamped banking rules and brokerage oversight, but the process is ongoing and will take many more years to stabilize (Sorkin 2009).

Decisions

All aspects of the financial industry heavily utilize information technology. All firms face difficult daily questions on how to balance risks and profits. Overall, banks have profited because they have become more adept at identifying potential bad loans and diversifying risk. The entire industry has become more interconnected through new financial devices to share risk and create new products. The use of credit-default swaps and derivatives has been credited with protecting banks in the downturn of the early 2000s (Economist 2004). Derivatives spread risks across the entire industry and multiple firms, which is good. But, it means that all companies, investors, and regulators have to carefully monitor a wide range of financial indicators and evaluate large-scale effects. How can anyone possibly handle this much data to make informed decisions?

Additional Reading

Dugas, Christine, "Bankruptcy Filings Set Reocrd in 2002," *USA Today*, February 14, 2003.

Gross, Daniel, "Capital One," What's in Your Wallet?" *Slate*, October 25, 2002.

"Just Deserts?" *The Economist*, April 15, 2004.

Nocera, Joseph, "Wall Street On the Run," *Fortune*, June 1, 2004.

Revell, Janice, "How Much Do Brokers Have to Hide?" *Fortune*, April 28, 2002.

Rosenberg, Yuval, "Sandy Hands Over the Keys to Citi," *Fortune*, July 16, 2003.

Stone, Amey, "Will Citi Upstage BofA's Big Deal?" *BusinessWeek*, October 29, 2003.

Sorkin, Andrew Ross, *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves*, Viking, 2009.

http://www.cardweb.com

Case: Citigroup

Charles Prince III, general counsel of Citigroup explained the overall challenges facing management: "You've got five or six or seven businesses—credit cards to mortgages to personal loans to investment banking to commercial insurance. They relate in important ways, but they're different. And they're all over the place. No one has ever had a company as broad in geographic scope [101 countries], as broad in product set, and as deep in size" (Loomis 2001). Sandy Weill gained control of Citigroup (ticker: C), the parent company, in 2000 after a merger between his company (Travelers) and Citibank. In 2001, most of the top-level management was reorganized, but the structure is a little non-traditional. Robert Willumstad is head of consumer business. Michael Carpenter runs the corporate division including the Salomon Smith Barney brokerage and Citi's commercial bank. Thomas Jones runs the investment management and private bank. Bictor Menezes heads the emerging markets section. Charles Prince and Jay Fishman, technically co-chief operating officers, also report directly to Weill. Prince handles operations and administration, Fishman risk and finance. Neither has any direct authority over the other business leaders. One of Fishman's jobs is to synthesize the overall risk picture of the business (into a 54-page monthly book), so that everyone can see the current issues and risks. The lines in many cases are somewhat blurry. Prince, for example, has a responsibility to promote cross selling of products. Because most of the company growth has come from acquisitions, few people down in the hierarchy know much about the other divisions and options. Chief financial officer Todd Thomson also has oversight of cross selling. Yet, neither Prince nor Thomson has direct authority over the businesses that do the cross selling. Weill points out that "this company is too big to micromanage, but it's not so big that you can't know what's going on" (Loomis 2001). Overall, the combination of a huge bank and brokerage firm creates significant advantages over the competitors. When necessary, Citigroup can swing an investment underwriting deal by offering a Citibank loan—something that is much harder to do at brokerage firms like Merrill Lynch.

One controversial step Weill has taken is to minimize the use of the Internet. Instead of trying to run it separately, he moved the projects into the subdivision units and asked them to merge them into their own operations—while trying to reduce costs. Similarly, the Salomon Smith Barney retail brokerage is a full-service company and does not offer low-cost Internet trading. On the other hand, Deryck Maughan, CEO of E-Citi worked hard to establish an electronic foreign-exchange market called Atriax in partnership with three other large banks. "We are saying that we would rather disintermediate ourselves in partnership with others and grab a large share of the new market than sit around in some pre-Information Age factory" (Loomis 2001). In October 2003, Weill remained chairman but made Prince CEO and Willumstad president (Stone 2003).

American financial institutions were hammered by the recession of the early 1990s, even Citibank came close to going under. Yet with the recession of the early 2000s, most banks, including Citigroup, increased profits. Average return on bank assets jumped from 8 percent in the 1990s to 16 percent in the 2000s. One reason is that the 2001 recession was not as deep, and banks are not as tied to corporations as before, supplying 40 percent of the funds today versus 50 percent a decade earlier. Banks also profited from the drop in interest rates engineered by the Fed. But banks are still the biggest source of funds for companies. The big change: a substantial improvement in risk management at American banks. Doug Woodham of Moody's observes "there has been a step change in risk management" (The Economist 2004). For example, mortgages have been turned into securities that are sold on a national market—spreading the risk away from the issuing bank. Credit-default derivatives have accomplished the same task for corporate loans, with the market reaching $350 billion in 2003.

The Hispanic population represents a growing market—particularly for Citibank with its international reach and reputation. In 2004, Citibank created a checkless checking account specifically for Hispanic customers. Many workers from South America come to the United States to make money, which they send back to their families. Yet many of them do not realize the importance of banks and ATMs in the U.S. economy. Most have difficulty opening checking accounts. The Citibank account relies on a debit card and Citibank's international collection of ATMs. Workers can deposit money in the account, and family members can withdraw it almost anywhere in the world through a local ATM (Wentz 2004).

Citibank is one of the largest issuers of credit cards in the United States. Yet that market is stagnant with intense competition and little opportunity for profit.

Some firms have tried to expand profits by going after the subprime market, but Citigroup has resisted. Instead, it is aiming to expand its reach globally. Leveraging its experience in emerging markets, it began offering Visa cards to Chinese consumers in February 2004. The charges can be paid using either dollars or Chinese yuan. Charles Prince and other dignitaries opened the service with parades through the streets of Shanghai. In 2003, only 25 million of the 1.3 billion Chinese citizens held credit cards. Some experts expect the number to easily triple in 10 years. Although the untapped market seems to offer incredible benefits, the risks are enormous. There is only one credit bureau in the entire country, minimal market research, and high bank fraud. Similar expansions of credit in Hong Kong and Korea resulted in huge losses from personal bankruptcies. Citigroup is mitigating the risk by performing their own background checks, only issuing cards to clients older than 21, and requiring an annual income of at least $6,000 (restricting the market to only 10 percent of Shanghai's 16 million residents). Citibank has also devised a mathematical scoring system to asses each individual's credit risk. The McKinsey consulting firms notes that the current default rate in China is 1.5 percent compared with 5 percent in the United States. An interesting twist to the risk problem is that the customers tend to pay off their card in full each month, so the bank has to make its money on fees instead of interest charges. (Baglole 2004).

To train brokers and agents, Citigroup needed a better way to communicate. In 2004 and 2005, the company installed IBM's Digital Media for Banking technology that enables it to stream live and archived content to 300,000 PC desktops across the company's office and branches (Mearian 2004). Communication with external organizations is more complicated. A key feature in the modern banking system is communications from the banks to the credit bureaus. On a regular basis, Citigroup sends a tape of customer data and payment histories to various credit bureaus. On May 2, 2005, Citigroup sent a tape to an Experian facility in Texas via UPS. However, the tape never made it to the company and UPS was unable to track the package. The unencrypted tape contained data on 3.9 million customers including Social Security numbers and payment history. After the incident, Citigroup decided it would be safer to begin encrypting the data on the tapes (McMillan 2005).

In early 2007, Citigroup laid off 17,000 workers in an effort to cut costs by $10 billion. The company also reorganized IT operations, largely by consolidating data centers. The company said that "simplification and standardization of Citi's information technology platform will be critical to increase efficiency and drive lower costs as well as decrease time to market." The company also planned to move 9,500 back-office positions to lower-cost locations, including off shore (Vijayan 2007). Analysts observed that Citi had little choice because between 2000 and 2005 its revenue increased by 8 percent while costs soared by 15 percent. Citigroup planned to close half of its 42 data centers. Much of the savings will come by replacing server clusters with grid computing based on thousands of smaller computers networked together (Crossman 2007). Another part of the plan involved reducing its credit card platforms from 12 down to 2, and reducing mortgage-origination software systems from 5 to 1 (McDougall 2007). However, at the same time, Citigroup announced that it was upgrading 325,000 employee desktop PCs to Windows Vista (Lai 2007).

The housing market collapse of 2007-2008-ongoing caused huge problems for most banks. The main problem was that banks had made loans to weak customers who could not pay them back. When housing prices collapsed (many cut in half), borrowers decided to default on their loans and walk away. Banks who made the loans had resold the loans as packages to everyone else. The entire banking industry was hit with a liquidity crisis. Citigroup alone posted losses of $27.7 billion in 2008 and $1.6 billion in 2009. With a huge infusion of cash by the federal government, Citigroup returned to profitability ($10.6 billion) in 2010. Homeowners did not really gain anything. Banks were also required to increase their capital holdings to reduce the risk of future meltdowns.

To reduce costs, Citigroup cut another 52,000 jobs at the end of 2008. Many of the job cuts came through selling subsidiaries, including the sale of Citigroup's India-based computer management center to Wipro (Thibodeau 2008).

Largely due to the financial crisis, banks largely made few changes in operations or technology. On the other hand, security issues continued. Citigroup was hit by hackers who stole card numbers from over 360,000 accounts (McMillan 2011). Hackers did not get into the main database, but were able to guess account numbers and log into the customer Web site. Sensitive data such as Social Security Numbers, expiration dates, and CVV codes were not compromised. Still, the thieves racked up $2.7 million in charges on 3,400 accounts. The bank covered those losses, and issued new cards to all of the affected customers.

Many banks have jumped on the smartphone platform and offer apps to help customers monitor their accounts and transfer money. Citigroup encountered a problem with their early version. The software stored account information on a filed on the iPhone, and apparently this data was not encrypted. If someone lost a phone, it would have been relatively easy for someone to read the file to get the account information. Citi releaed an upgrade that removed the file (Keizer 2010).

In another interesting security problem, Citigroup reported that a former vice president (Gary Foster) allegedly stole $19.2 million from corporate accounts that were in his responsibility (Bray and Rothfeld 2011). Basically, he is alleged to have set up false contracts and just wired the money to his own account. The transactions were uncovered in an audit. Shannon Bell, a Citi spokeswoman stated that "We are outraged by the actions of this former employee. Citi informed law enforcement immediately upon discovery of the suspicious transactions and we are cooperating fully to ensure Mr. Foster is prosecuted to the full extent of the law."

Questions

1. How does Citigroup use models and information systems to make decisions in the credit card market?
2. How is the Chinese and Southeast Asian market for credit (especially in Korea) different from that in the United States?
3. How could Citigroup make better use of the Internet?
4. How can models, expert systems, and other information tools help Citigroup manage such a large organization?

Additional Reading

Baglole, Joel, "Citibank Takes Risk by Issuing Cards in China," *The Wall Street Journal*, March 10, 2004.

Bray, Chad and Michael Rothfeld, "Former Citi Executive Charged in $19 Million Theft," *The Wall Street Journal*, June 27, 2011.

Crossman, Penny, "Wall Street-Style Power," *Information Week*, April 14, 2007.

The Economist, "Just Deserts?" April 15, 2004.

Keizer, Gregg, "Citi Confirms Critical Bug in iPhone Mobile Banking App," *Computerworld*, July 26, 2010.

Lai, Eric, "Ballmer: Citigroup to Upgrade 500,000 PCs to Vista in Next Year," *Computerworld*, April 19, 2007.

Loomis, Carol J., "Sandy Weill's Monster," *Fortune*, April 1, 2001.

McDougall, Paul, 'Citigroup Cost Cutting Puts a Bull's-Eye On IT," *Information Week*, April 14, 2007.

McMillan, Robert, "Citigroup to Encrypt Data sent to Credit Bureaus," *Computerworld*, June 7, 2005.

McMillan, Robert, "Citigroup Hackers Made $2.7 Million," *Computerworld*, June 24, 2011.

Mearian, Lucas, "Citigroup Rolls out Streaming Media to Corporate Desktops," *Computerworld*, May 28, 2004.

Stone, Amey, "Will Citi Upstage BofA's Big Deal?" *BusinessWeek*, October 29, 2003.

Thibodeau, Patrick, "Wipro Says Buying Citigroup Unit will Boost IT Infrastructure Outsourcing in India," *Computerworld*, December 23, 2008.

Vijayan, Jaikumar, "Citigroup to Lay Off 17,000, Overhaul IT Operations," *Computerworld*, April 11, 2007.

Wentz, Laurel, "Banks Tailor Efforts to Homesick Hispanics," *Advertising Age*, April 5, 2004.

Case: Capital One

As a newcomer to the credit card industry, Capital One (ticker: COF) needed an edge. The company created the edge by building the Information Based Strategy (IBS). The system is highly focused on testing and evaluating various options before offering them to the public. The three-step approach: (1) create a new product and find a target population, (2) create a test by changing the variables and seeing how the group members react, and (3) use the test results to further divide the segment, and then test specific campaigns against those segments. By specifically targeting rates, fees, and options to each market, Capital One was able to obtain high response rates, reducing the costs of acquiring a customer. The company conducted 45,000 tests in 2000 alone (120 per day). The company extended the process to its Web site, tracking visitor activity, tailoring options to each specific customer, and using the background data to buy ads on other sites with the appropriate demographics. For example, if sports enthusiasts responded well to certain features, those features would be advertised on sporting sites. With this process, the company opened 2 million new accounts online (Cohen 2001).

The company uses the test results and marketing to tailor interest rates and fees to individual customers. In effect, it has created 100,000 different segments or product combinations. The process extends to the customer call center. Capital One uses Cisco Systems Global Service Logistics (GSL) system to route calls. The system retrieves information about the customer and routes the call to the most appropriate representative. For example, a customer who routinely pays off the monthly charges might be routed to a representative to sell a platinum card with a higher balance.

In a different twist, Capital One uses a similar process to hire and promote employees. It records specific data on every hire, including scores on a timed math test and a behavioral test. During promotions, the characteristics of successful employees are listed for each job. These characteristics, and the initial test scores, are used to refine the hiring process for each type of job. The process fits employees to jobs and reduces hiring costs (Cohen 2001).

For several years, Capital One had enviable growth rates with revenue growing from $95 million in 1995 to $4.97 billion in 2001, from 6 million customers to 33 million (Cohen 2001). Yet because of the extremely competitive nature of the industry, Capital One had to extend into the subprime market to capture more customers. With the economic downturn, the subprime market crashed in 2002. Capital One suffered along with the others in the industry. For the first time, managers also revealed how dependent Capital One was on the subprime market. In 2002, 40 percent of its cards were in the hands of subprime customers (Albergotti 2003), far above what most investors had believed. The market punished Capital One's stock. Yet, ultimately, Capital One's system worked. The company's risk management techniques gave it one of the lowest levels of bad loans—until 2003. The company wrote off bad debt, but earnings eventually rose and the company survived. The company has reduced its exposure in the subprime market. Federal oversight also forced companies to reevaluate the subprime market—regulators have been attempting to limit fees and penalty charges (Smith 2002). But, the strategic shift leaves Capital One facing more competition in the higher quality markets. Its profit margins have shrunk accordingly—down to 16 percent in 2004, half their earlier levels (Byrnes 2004).

Searching for new markets, Capital One is widening its search into more traditional banking areas. Richard D. Fairbank, the CEO, knows that growth in the card market will be slow. "A lot of the different financial markets are evolving, but the most evolved is the credit card business. This is pretty close to the endgame. We've really got to work for a living these days in the credit card market." His new target: auto loans, installment loans, and even international loans. In particular, he wants to target small businesses. Capital One's lending to small businesses has grown from $400 million in 1999 to $3.3 billion in 2003. At the end of 2003, the company managed $46.3 billion in credit card loans, $8.5 billion in autos, $5.4 billion in installment loans, and $7.6 billion in international loans (Kuykendall 2004).

In the meantime, the company is cutting expenses, eliminating as many as 2,500 of its 9,000 positions in its Richmond, Virginia, headquarters. Fairbanks told employees he needed to cut expenses by 20 percent. "Our businesses need to improve their cost positions to compete in the future against leading players in the financial-services industry" (Hazard 2004). Capital One is also interested in using the Internet to reduce costs. Processing payments online instead of paper checks reduces costs by $1 a year per customer. Persuading customers to accept electronic statements instead of mail cuts costs an additional $5 per customer per year. Rick Long, director of U.S. card operations at Capital One, notes that "our ROI [return on investment] models are built on lowering costs" [Wade 2004). He was emphasizing that Capital One is primarily interested in using the Internet only where it is cost effective. The company is also concerned that customers have to be enticed to use the Internet instead of coerced. "If we drive them to the site, they may not stay" [Wade 2004). Most customers who are interested in using the Internet, enroll within the first six months.

To broaden its reach and financial offerings, Capital One embarked on a campaign of buying retail banks. The company purchased North Fork Bancoprt as well as Hibernia Corp. (Wei 2007) Capital One IT managers were so impressed with Hibernia's deposit-transaction system that they migrated their banking system onto its mainframe package. In its other 9 recent acquisitions, Capital One migrated most operations onto its existing platforms. In particular, HRM, finance, and other operations are transferred to Capital One's PeopleSoft ERP system (McGee 2006).

Like other institutions, Capital One is trying to reduce expenses. In 2007, the company announced job cuts of 2,000 employees in an effort to reduce costs by $700 million (Wei 2007). The IT department is working to cut expenses by reducing the number of servers it runs. Using server virtualization, a hardware-based server can be split into multiple virtual servers, with each one running different operating systems and different applications. One goal is to reduce costs by moving from 1,600 down to 1,100 servers. Le Congdon, VP of corporate technology also said that the technology "buys us time in circumstances where a specific application is ready to move over and we need to shut down that physical location, or where we need to move of aging hardware or software. It also gives us better management tools and controls across the environment. If we can go from 50 percent to 85 percent utilization on the servers, that's free money" (Thibodeau 2006).

Capital One was also hurt by the financial crash. But, because the firm focused on higher-end customers (compared to Providian), it was able to survive the liquidity crash. Rob Alexander, CIO, pointed out that the company was still investing in technology to handle customer transactions. He also emphasized the importance of business analytics at Capital One to ensure the company was making the best decisions possible (Fonseca 2008). Alexander also noted that the IT department maintains an advanced technology group to look at cutting-edge technologies and test them to see how they might improve operations or deliver new business opportunities in the near future.

In 2011, Capital One purchased the U.S. online-banking business of ING—a Dutch bank that pioneered some online banking. The deal will make Capital One the fifth largest U.S. bank measured by the amount of deposits (Grocer 2011).

Questions

1. How does Capital One use decision support systems to reduce risk and increase sales?
2. With its advanced systems, how did Capital One lose so much money in 2003?
3. How can Capital One use technology to reduce its costs?

Additional Reading

Albergotti, Reed, "Fewer Credit-Card Solicitations Are in the Mail," *The Wall Street Journal*, November 5, 2003.

Byrnes, Nanette, "Coming of Age At Capital One," *BusinessWeek*, February 16, 2004.

Cohen, Jackie, "Growth Formula," *Computerworld*, July 2, 2001.

Fonseca, Brian, "The Grill: Capital One's CIO Discusses Business Strategies for Tough Financial Times," *Computerworld*, November 3, 2008.

Grocer, Stephen, "Deal Profile:" Capital One Wins ING Direct USA," *The Wall Street Journal*, June 17, 2011.

Hazard, Carol, "Capital One Shifts Focus; Firm to Cut Richmond, Va., Jobs to Reduce Costs," *Richmond Times*, April 29, 2004.

Kuykendall, Lavonne, "Amex, Capital One Talk About the Future," *The American Banker*, June 7, 2004.

McGee, Marianne Kolbasuk, "Capital One: A Serial Acquirer," *Information Week*, June 26, 2006.

Smith, Geoffrey, "The Bill Comes Due For Capital One," *BusinessWeek*, November 4, 2002.

Thibodeau, Patrick, "Capital One Cuts Servers With Virtualization Push," *Computerworld*, February 13, 2006.

Wade, Will, "Coming Full Circle on E-Strategies," *The American Banker*, June 15, 2004.

Wei, Linling, "Capital One to Cut About 2,000 Jobs," *The Wall Street Journal*, June 28, 2007.

Case: Providian/Washington Mutual

The credit card industry has changed in the decades since the 1960s when Bank of America flooded households with millions of cards. Today, trillions of dollars of transactions are paid with credit cards, with 1.2 billion in use in the United States alone. The average cardholder has 2.7 bank cards, 3.8 retail cards, and 1.1 debit cards for a total of 7.6 cards per person. Despite the huge variety, the top 10 issuers handled about 78 percent of the total value at the end of 2002. About 24 percent of all consumer retail transactions are paid with credit and debit cards (www.cardweb.com).

A huge percentage of the U.S. population already holds several credit cards, and banks are constantly competing to give more cards to high-income customers. To most observers, the market appears to be saturated. Customers will switch banks if one offers a better rate or different benefits. Providian Financial (PVN), a San Francisco-based company, looked at this market and decided to find a new niche: the subprime market. According to cardweb, approximately 25 percent of the population has weak credit ratings. The objective was to find low-income people with poor or no credit, offer them a credit card, and then charge huge fees and high interest rates. Shailesh Mehta, who eventually became CEO, developed a complex mathematical model that allowed the company to identify the subprime customers who would be most likely to use the cards, but not default on the loans. One executive characterizes the process as "we found the best of the bad." Since other banks refused to serve this market segment, Providian had solid growth rates and, despite the 24 percent and higher interest rates, was able to claim that it was providing a useful service to the customers. Eventually, other banks (notably Capital One and Household International) jumped into the market. By 2000, 20 percent of the cardholders were classified as subprime (cardweb), and Providian was the fifth largest card issuer in the nation (Koudsi 2002). In 2000, the company was ordered to pay $300 million in restitution for misleading, unfair, and deceptive business practices. To keep growing, the company had to find a new market. It failed in an experiment to go after the platinum market. So, it went after increasingly risky customers. By 2001, when the economy faltered, the subprime market felt the impact first, and the company saw its default rates jump—to a huge 12.7 percent. Amid claims of misleading accounting, the stock price plummeted 90 percent. Cardweb notes that 1.3 million cardholders declared bankruptcy in 2001—which generally erases unsecured credit card debt. Other banks focusing on the subprime market faced similar problems. Bernhard Nann of Fair Isaac, which evaluates consumer creditworthiness, notes that "data used in traditional behavioral models is not quite as powerful as we would like. We'd like to go and expand the universe of data that can be looked at to produce better predictions. That's particularly important in the subprime area." In particular, the company wants to include utility and rent payments in its calculations (Punch 2003).

With the economy improving in 2003 and 2004, Providian was able to sell off some of its loans and write off the really bad ones. The company also refocused its efforts to go after the middle market, people with a FICO score between 600 and 720 (Albergotti, 2003). Although weak, the company was able to report a profit in late 2003 (Business Week 2003). By 2004, some experts saw an improvement in the credit card industry profits. Many homeowners had already refinanced their mortgages—and paid down credit card debt. But, that leaves them free to increase the borrowing—hence increasing profits (Stovall 2004).

Yet, with increasing competition for the platinum and midmarket segments, banks continue to search for ways to make money on credit cards. With low balances and low interest rates, and few new customers, banks turned to fees. Income from late fees and penalties was predicted to reach $13 billion for 2004, with total fees likely to account for 39 percent of revenue (Simon 2004). Providian joined in the parade and increased its fees annually. However, the company also created a "Real Rewards" program, where cardholders accumulate points for cash rebates that can also be used to offset late fees.

In 2004, Providian settled the accounting and insider-trading lawsuits for $65 million (Kuykendall 2004).

In 2005, Providian was purchased by Washington Mutual Bank WaMu (Hallinan 2005). WaMu was formed in 1889 in Washington State. It runs consumer and small-business retail banks from the West Coast to the Midwest. It also has an aggressive home-lending program. In the past few years, it has become known as a bank for taking risks and targeting subprime loans—both in credit cards and mortgage loans. WaMu was pressured by the crashing housing market in 2006 and 2007. Because of its aggressive subprime loans, the company was one of the first to offer new loans to customers to try and move them into safer loans as rates increased, in an effort to keep large number of customers from defaulting on housing loans (Carrns 2007 and Hagerty and Carrns 2007).

WaMu aggressively uses technology to reduce labor costs in its banks. In most of its branches, tellers assist customers with transactions that are handled electronically. Effectively, tellers work at advanced ATM stations that enable them to handle tasks only slightly more complex than a standard ATM machine. If a customer processes a transaction involving cash, the money is dispensed by the machine.

The real estate crash and financial liquidity crunch were tough on WaMu. Essentially, the government forced the sale of WaMu to JP Morgan Chase, who picked up the assets for relatively low cost. Shareholders of WaMu stock pretty much lost all of their value. For a couple of years after the crash, top-level managers tried to argue that the government should not have forced the sale of WaMu, and that government actions might have contributed to the failure of the company. But, in the end, most observers simply blamed the over-aggressive lending practices of the company. When the real estate market crashed, an increasing number of loans turned bad, and even moderately risky customers walked away from their debt. These were the customers targeted by WaMu, and it lacked the capital and stable profits to balance these risks. JP Morgan had to write off at least $31 billion in bad loans (Sidel, Enrich, and Fitzpatrick 2008).

Questions

1. Why is the subprime market so risky and how does Providian use information technology to minimize the risks?
2. How are the platinum and midmarket accounts different from the subprime market, and can Providian use the same models?
3. How can Providian use information technology to increase its revenue from fees?

Additional Reading

Albergotti, Reed, "Fewer Credit-Card Solicitations Are in the Mail," *The Wall Street Journal*, November 5, 2003.

BusinessWeek, "UBS Financial Cuts Providian to 'Neutral'," October 30, 2003.

Carrns, Ann, 'Washington Mutual Earnings Slide on Home-Loan Losses," *The Wall Street Journal*, April 18, 2007.

http://www.cardweb.com/cardlearn/stat.html

Hagerty, James R. and Ann Carrns, "WaMu Leads in Risky Type of Lending," *The Wall Street Journal*, April 17, 2007.

Hallinan, Joseph T, "WaMu Has Deal To Buy Providian For $6.45 Billion," *The Wall Street Journal*, June 7, 2005.

Koudsi, Suzanne, "Sleazy Credit," *Fortune*, February 19, 2002.

Kuykendall, Lavonne, "Providian Says It Can Move On After Settling Class Action for $65M," *The American Banker*, June 9, 2004.

Punch, Linda, "Subprime Cards, Act II," *Credit Card Management*, November 2003.

Sidel, Robin, David Enrich, and Dan Fitzpatrick, "WaMu Is Seized, Sold Off to J.P. morgan, In Largest Failure in U.S. Banking History," *The Wall Street Journal*, September 26, 2008.

Simon, Ruth, "Credit-Card Fees Are Rising Faster; As Economics of Industry Change, Issuers Seek To Boost Income From Penalties; $35 Late Charge," *The Wall Street Journal*, May 31, 2004.

Stovall, Sam, "A Rising Rating for Consumer Finance," *BusinessWeek*, February 24, 2004.

Case: Merrill Lynch

For decades, Merrill Lynch (ticker: MER) was the premier brokerage firm. It led the industry in sales and innovations. The firm helped create hundreds of modern companies with its access to capital and the ability to sell stocks. In the 1970s and 1980s, its innovations in money market funds rocked the financial world and resulted in substantial changes to the federal banking laws. It accomplished many of these innovations through technology—and the recognition that money is really data. It expanded its reach overseas, and became the largest retail brokerage and the largest underwriter of stocks and bonds in the world. Through the stock market bubble of the 1990s, Merrill prospered—at least in terms of market share. They helped finance several start-ups and retail stock brokerage revenues grew by $3 billion from 1996 to 1998. But the growth carried equally huge costs: only $100 million of that money made it to the bottom line. By 1998, Merrill's profit margins were 10 percentage points lower than its competitors. At the same time, the growth of online account management and discount brokerage firms had chipped away at the retail end of Merrill's operations. By 2001, even the directors knew that Merrill was in trouble and might not survive. Being the largest was not very useful if the company was not making money (Rynecki 2004). In 2001, the board made the surprising move of appointing Stan O'Neal, the CFO to become the new CEO of Merrill. Recognizing the importance of profits, and seeing limits on the ability to increase revenue, he began focusing on cutting costs. Ultimately, he eliminated 24,000 jobs, including 20 percent of investment banking and analyst positions. He closed 300 field offices and completely pulled out of Australia, Canada, New Zealand, and South Africa. He reduced the number of stocks traded directly by the firm by 75 percent. He replaced almost all of the management from the top down—largely replacing them with younger staffers looking to make their marks. He reportedly often states that "ruthless isn't always that bad" (Rynecki 2004). The result: in 2003, the firm earned a record $4 billion in profit, and pretax margins reached 28 percent—vastly exceeding the high-growth years of the 1990s.

Despite the profitability, the firm still has to worry about revenue. In 2004, with the slow economy and shaky financial markets, Merrill ranked last in sales growth among the Newsday Top 100—with a three-year sales decline of 14.8 percent (Murray 2004). Merrill Lynch is refocusing its retail brokerage operations—on wealthy customers. The company pays higher interest rates to customers with balances over $10 million (McGeehan 2004). Merrill also jumped into the credit card market in 2004. The company realized that customers withdrew $3 billion from their accounts in 2003 to pay off credit card debt. Co-branded with MBNA, the card is designed to provide one-stop shopping for its customers. Merrill is targeting its wealthier customers—in an effort to provide more services, but also to keep control over a larger percentage of their money. The company is offering a variety of awards to entice customers to use them for large transactions (Lieber 2004).

Despite the renewed emphasis on cost, Merrill still has to handle business, and that requires technology. In 2003, the company began the first steps of a $1 billion upgrade to their broker workstation systems. The new Client 360 system is designed to be a total wealth-management tool—largely focused on customer relationship management. It is designed to provide a complete view of customer data to the broker. John Killeen, the chief technology officer, notes that “the tool suite around wealth management is pretty well established across the industry, still what will make this unique is that we are facilitating the relationships for our financial advisers with their clients. … It’s a single-screen representation of the most important and most prevalent questions that a client may ask a financial adviser when they have them on the phone. It will talk to balances; it will talk to progress towards plans, and any important notices that affect that client. And within one or two clicks, we can drill down to greater and greater detail. It represents a huge productivity gain for our financial advisers” (Pallay 2004). Most of the front-end system being developed is installed by Thomson Financial, which is overseeing the integration of work from 400 vendors. The back-end system remains Merrill’s proprietary system, largely running on Microsoft software. The company built a framework that separates the back-end and middle tier systems from the front end. This approach enables the company to alter various components on either end without having to rebuild the entire system. Despite the complexity of the project, Killeen’s focus has been on the end-use applications: “I think the lynchpin around the tools will be client data—who has the greater understanding of the total client relationship, what is important to the client at the different phases of their financial life cycle, and how do we put all of that together so that our advisers are best positioned to work with their client?”

Technology has been driving several changes in market trading for several years. For a while, the stock exchanges frowned on automated trading schemes, but by 2007, they had adopted them wholeheartedly. Brokers, including Merrill Lynch, have developed complex models of trading strategies. Some look for arbitrage opportunities across markets, others look for timing gaps. These systems can generate hundreds of trade orders per second, and generally hold positions for only a few seconds or minutes (Martin 2007). Economically, these systems serve to remove bumps and reduce imperfections in the market. In the process, the large brokers make profits not available to the average investors, but the profits pay for the technology.

In terms of customer-based technology, Merrill Lynch is reportedly the largest customer of Salesforce.com with licenses for 25,000 users (Weier 2007). Salesforce provides customer relationship management software via an online system. The company offers its Wealth Management Edition for Merrill Lynch and other financial institutions. In terms of processing transactions, federal regulations began changing the rules in July 2007 with the implementation of Reg NMS. The regulation requires brokers to automatically and instantly route orders to the market that displays the best price (Horowitz 2007). Only the largest brokerage firms have been able to afford the technology to maintain constant connections and price searches with multiple markets. With many of the exchanges automated, the brokerage computers need to be fast and secure to handle the number of searches required for each transaction. Tim Cox, VP of market structure strategy at Merrill Lynch observed that "The big are going to get bigger. Reg NMS over the course of a couple of months will make technology a real differentiator" (Horowitz 2007).

Merrill Lynch only barely survived the financial crisis of 2008. The company had made huge bets on the housing market—and held those bets on its balance sheet. Its competitor Lehman Brothers was "allowed" to fail by federal regulators—partly to send a message to the rest of the industry. Merrill Lynch would probably have been close behind, but Bank of America took a $50 billion gamble and bought out the company (Farrell 2010). Ultimately, the question of whether Merrill Lynch was worth billions of dollars to Bank of America remains to be seen. Can the division generate enough revenue to justify its continued existence? After all, it is still a high-priced brokerage firm. Perhaps it can help Bank of America in high-end deals for IPOs and other investment banking tasks.

Questions

1. How can Merrill Lynch survive as a full-service broker, particularly with the regulatory changes in investment research?
2. How can Merrill Lynch use information technology to attract and keep the high-end investors that it wants?
3. How can the company use technology to reduce its costs and improve its profit margins?

Additional Reading

Farrell, Greg, *Crash of the Titans: Greed, Hubris, The Fall of Merrill Lynch and the Near-Collapse of Bank of America*, Crown Business 2010.

Horowitz, Jed, "How New Rule May Tip Brokerage Scales," *The Wall Street Journal*, July 6, 2007.

Lieber, Ron, "A New Contender for Your Wallet; Merrill Jumps into Card War With No-Fee Rewards Offer, But Does It Beat Platinum?" *The Wall Street Journal*, April 20, 2004.

Martin, Richard, "Wall Street's Quest To Process Data At The Speed Of Light," *Information Week*, April 21, 2007.

McGeehan, Patrick, "Schwab Tells Less-Affluent Customers They Have To Pay More Than Richer Clients To Make Trades," *The New York Times*, May 26, 2004.

Murray, Christian, "Top 100; After a Low, Merrill Looking Up," *Newsday*, June 14, 2004.

Pallay, Jessica, "An Interview with Merrill Lynch's John Killeen," *Wall Street & Technology*, May 2004.

Rynecki, David, "Putting the Muscle Back in the Bull," *Fortune*, April 5, 2004.

Weier, Mary Hayes, "New Deals Show Software-As-A-Service More Than A Niche," *Information Week*, February 28, 2007.

Summary Industry Questions

1. What information technologies have helped this industry?
2. Did the technologies provide a competitive advantage or were they quickly adopted by rivals?
3. Which technologies could this industry use that were developed in other sectors?
4. Is the level of competition increasing or decreasing in this industry? Is it dominated by a few firms, or are they fairly balanced?
5. What problems have been created from the use of information technology and how did the firms solve the problems?