

2

Framing the Problem

While there are many different types of quantitative analysis, they all have certain key features and steps in common. As we noted in chapter 1, a quantitative analysis follows the following three stages and six steps:

FRAMING THE PROBLEM

- Problem recognition
- Review of previous findings

SOLVING THE PROBLEM

- Modeling and variable selection
- Data collection
- Data analysis



COMMUNICATING AND ACTING ON RESULTS

- Results presentation and action

In this chapter and chapters 3 and 4, we'll describe each stage and step individually, and provide a couple of examples of quantitative analyses that cover all six steps, but feature the particular stage of analysis being discussed in the chapter. At the end of each of the three chapters we'll lay out two examples—generally one from business and one involving society in general or personal experience—that illustrate how all six steps were employed in an analysis, but again focus in particular on one stage of the analysis. Our three-stage, six-step process isn't the only way to do analytics (for example, there is a Six Sigma methodology for analyzing variation in product quality yielding no more than 3.4 defects per million products produced), but we expect that most analytical experts would agree with it, and it's broad enough to encompass a lot of different types of business problems and analyses.

The Problem Recognition Step

A quantitative analysis starts with recognizing a problem or decision and beginning to solve it. In decision analysis, this step is called *framing*, and it's one of the most critical parts of a good decision process. There are various sources that lead to this first step, including:

- Pure curiosity (common sense, observation of events)



- Experiences on the job
- Need for a decision or action
- Current issues requiring attention (of a person, an organization, or a nation)
- Building on, or contesting, existing theories or past research
- Accepting of project offers or determining availability funding

Note that at this step, the analytics are yet to come. The decision to forge ahead with some sort of analysis may be driven by a hunch or an intuition. The standard of evidence at this point is low. Of course, the whole point of a quantitative analysis is to eventually apply some data and test your hunch. That's the difference between analytical thinkers and others: they test their hunches with data and analysis.

The most important thing in the problem recognition stage is to fully understand what the problem is and why it matters. The answers to these two questions not only make it clear what can be accomplished by solving the problem, but also facilitate the ensuing stages.

Identifying the Stakeholders for the Analysis

Perhaps obviously, the individuals involved at this step are primarily managers and decision makers—the “owners” of the business or organizational problem. However, even at this stage their efforts can be greatly aided by the presence of experienced quantitative analysts who understand the business problem, the decision process, and the likely quantitative approaches to be employed. If all of that knowledge can't be found within one person, you may need a team that jointly possesses it.

It's worth some serious thinking at this step about who the stakeholders are for the analysis you plan to undertake, and how they're

Stakeholder Analysis Worksheet

If you can't answer most of these questions with a "yes," your project may be in trouble from the beginning:

1. Is it clear what executives have a stake in the success of your quantitative analysis project?
 2. Have they been briefed on the problem and the outlines of the solution?
 3. Do they have the ability to provide the necessary resources and to bring about the business changes needed to make the project successful?
 4. Do they generally support the use of analytics and data for decision making?
 5. Does the proposed analytical story and method of communicating it coincide with their typical way of thinking and deciding?
 6. Do you have a plan for providing regular feedback and interim results to them?
-

feeling about the problem (see the "Stakeholder Analysis Worksheet"). Do you have stakeholders who can take action on the results? Are they feeling skeptical that the problem even exists? Are they likely to be persuaded to do something even if the analysis is bulletproof?

The tendency of analysts is often to jump right into the analysis without thinking of stakeholders. The more confident they are in their analytical skills, the less they may worry about who will



ultimately be the audience for the results and the “deciders” about whether to take action.

If you’re persuaded of the need for stakeholder management for your analytical project, some of the common steps involved in that process include:

1. Identifying all stakeholders
2. Documenting stakeholders needs
3. Assessing and analyzing stakeholders interest/influence
4. Managing stakeholders expectations
5. Taking actions
6. Reviewing status and repeating¹

A stakeholder analysis can identify who are the primary decision makers, and how they are most likely to be persuaded by the results from the analysis. Even the most rigorous and bulletproof analysis approach will be of little use if it does not persuade a decision maker to act. In fact, it may even make sense to use a questionable approach from a methodological standpoint if that is the only evidence a decision maker will trust.

For example, Rob Duboff runs a marketing research and strategy firm called HawkPartners. In general, he believes in the value of quantitative research whenever possible. However, he has learned that some executives don’t understand quantitative approaches to learning customer wants and needs, and believe much more in qualitative approaches such as focus groups—convening a small group of customers or potential customers, asking them what they think about a company’s products and services, and observing and recording their responses. Now Duboff knows that focus groups are methodologically suspect. It’s pretty well known in the marketing research field that

customers are likely to tell you what you want to hear, and the fact that they say they like something doesn't mean they would actually buy it. These problems can be mitigated by a skillful discussion leader, but focus group results are not projectable to a wider universe. However, Duboff feels that any research is better than none, and if evidence from a focus group would be trusted and acted on by an executive—and more quantitative results would not—he conducts the focus group.

In a similar sense, the stakeholders for the decision can help to determine the form of output and results presentation. Humans differ in their preferences for seeing quantitative results; some prefer rows and columns of numbers, some prefer graphics, and some prefer text describing the numbers. It's important to elicit those preferences at a relatively early stage. If the results are going to be used not by humans but by a computer—and this is increasingly the case as more and more decisions are automated or partially automated—then it makes little sense to deliberate over the ideal visual format. Just feed it the numbers it thrives on!

It may also be the case that certain analytical approaches can help to involve stakeholders throughout the analysis. For example, at Cisco Systems, a forecasting project addressed the possibility that substantially more accurate forecasts were possible through statistical methods (we'll describe the six steps for this example at the end of chapter 7). Some Cisco managers were supportive of the project, but others doubted that better forecasts were possible. Anne Robinson, who managed the project, employed an "agile" methodology for the project, creating new deliverables every few weeks and presenting them to project stakeholders. The more incremental approach to solving the problem helped stakeholders buy into the new approach. Eventually it became clear to even the skeptical managers that the new forecasting approach was far more accurate and could be done more quickly and for more products than the previous nonanalytical approach.



Focusing on Decisions

We have found it helpful in the problem-recognition stage to focus on specific decisions that will be made as a result of the analysis. There are many reasons for this focus. One key reason is that a decision focus makes all participants realize that that is the reason for the quantitative analysis; it's not an idle exercise. Another is that focusing on the decision to be made will help to identify a key stakeholder: the person or group that will make a decision based on the analysis. A third key reason is that if there are no decisions envisioned, it may not be worth doing the analysis.

For example, Mike Thompson, the head of First Analytics, an analytics services firm, describes a meeting he had with a client at the problem-recognition stage. The client, a restaurant chain, believed that the primary focus of the analysis was product profitability. Client executives wanted First Analytics to determine how profitable each menu item was. Mike also subscribes to the idea of focusing on decisions, so he asked the client managers what decisions they would make as a result of the profitability analyses. There was a long silence. One executive suggested that the primary decision would be whether to keep items on the menu or not. Another pointed out, however, that the chain had not eliminated a single menu item over the past twenty years. After some further discussion, the client team decided that perhaps the focus of analysis should be menu item pricing, rather than profitability. "We have changed prices since our founding," one executive observed.

What Type of Analytical Story Are You Telling?

Once you've decided what decisions you are going to make, you can begin to think about how you are going to provide answers or insights for that decision. We'll talk in chapter 4 about *telling a story with data*, which is the best way to communicate results to nonanalytical



people. At this point, you need to begin thinking about what kind of story it is and how you will tell it, although many of the details of the story will come later in the analysis process. Stories are, of course, how numbers talk to people. There are at least six types of quantitative analytical stories; each of them is described below, along with an example or two.

THE CSI STORY. Some quantitative analyses are like police procedural television programs; they attempt to solve a business problem with quantitative analysis. Some operational problem crops up, and data are used to confirm the nature of the issue and find the solution. This situation often does not require deep statistical analysis, just good data and reporting approaches. It is often encountered in online businesses, where customer clickstreams provide plenty of data—often too much—for analysis.

One expert practitioner of the CSI story approach is Joe Megibow, vice president and general manager of online travel company Expedia's US business. Joe was previously a Web analytics maven—and he still is—but his data-based problem-solving approaches have led to a variety of impressive promotions.

Many of the Expedia investigations involve understanding the reasons behind lost online sales. One particular CSI story involved lost revenue on hotel payment transactions. Analysis of data suggested that after a customer had selected a hotel, filled in the travel and billing information, then clicked the "Buy Now" button, a percentage of the sales transactions were not being completed successfully. Megibow's team investigated the reason for the failures, again using Web metrics data and server log files throughout the process.

Apparently, the "Company" field under the customer's name was causing a problem. Some customers interpreted it as the name of the bank that supplied their credit card, and then they also supplied the

bank's address in the billing address fields. This caused the transaction to fail with the credit card processor. Simply removing the "Company" field immediately raised profits for Expedia by \$12 million. Megibow says that Expedia has explored many of these CSI-like stories, and they almost always yield substantial financial or operational benefits.

Sometimes the CSI stories do involve deeper quantitative and statistical analysis. One member of Megibow's team was investigating which customer touchpoints were driving online sales transactions. The analyst used the *Cox regression model*—an approach originally used to determine which patients would die and which would live over certain time periods—of "survival analysis." The analysis discovered that the simpler prior models were not at all correct about what marketing approaches were really leading to a sale. Megibow commented, "We didn't know we were leaving money on the table."²

THE EUREKA STORY. The Eureka story is similar to the CSI story, except that it typically involves a purposeful approach to a particular problem (as opposed to stumbling over the problem) to examine a major change in an organization's strategy or business model. It tends to be a longer story with a greater degree of analysis over time. Sometimes Eureka stories also involve other analytical story types, just because the results are so important to the organizations pursuing them.

At Expedia again, for example, one Eureka story involved eliminating change/cancel fees from online hotel, cruise, and car rental reservations. Until 2009, Expedia and its competitors all charged up to \$30 for a change or cancellation—above and beyond the penalties the hotel imposed. Expedia and other online bookers' rates were typically much lower than booking directly with a hotel, and customers were willing to tolerate change/cancel fees.

However, by 2009 it had become apparent that the fees had become a liability. Expedia's rates were closer to those of the hotels' own rates, so the primary appeal of Expedia had become convenience—and change/cancel fees were not convenient. Analysts looked at customer satisfaction rates, and they were particularly low for customers who had to pay the fees. Expedia's call center representatives were authorized to waive the change/cancel fees for only one reason: a death in the customer's family. A look at the number of waivers showed double-digit growth for the past three years. Either there was a death epidemic, or customers had figured out they could get their money back this way.

Expedia executives realized the market had changed, but change/cancel fees represented a substantial source of revenue. They wondered if the fees were eliminated, would conversion (completed sale) rates go up? In April of 2009, they announced a temporary waiver of fees for the month (a bit of a mad scientist testing story, described below). Conversion rates immediately rose substantially. Executives felt that they had enough evidence to discontinue the fees, and the rest of the industry followed suit.

Across town in Seattle lies Zillow, a company that distributes information about residential real estate. Zillow is perhaps best known to quant jocks for its "Zestimates," a proprietary algorithm that generates estimates of home values. But, like Expedia, Zillow's entire culture is based on data and analysis—not surprisingly, since the company was founded by Rich Barton, who also founded Expedia.

One of Zillow's Eureka stories involved a big decision to change how it made its money from relationships with real estate agents. Zillow began to work with agents in 2008, having previously been focused on consumers. One aspect of its agent-related business model was selling advertising by agents and delivering leads to them. Zillow charged the agents for the leads, but the value per lead was not

enough in the view of executives. Chloe Harford, a Zillow executive who heads product management and strategy, was particularly focused on figuring out the right model for increasing lead value and optimizing the pricing of leads.

Harford, who has a PhD in volcanology, or the study of volcanoes, is capable of some pretty sophisticated mathematical analysis. However, she and her colleagues initially relied on what she calls “napkin math” to explore other ways to generate more leads and price them fairly to agents. In April 2010, Zillow created a new feature—immediately copied by competitors—involving selling advertising to agents. It created many more customer contacts than before, and allowed the consumer to contact the agent directly. Zillow also introduced a sophisticated algorithm for pricing leads to agents that attempts to calculate the economic value of the lead, with an estimate of conversion rates. Competitors also do this to some degree, but probably not to the level of sophistication that Zillow does. The leads and pricing of them are so important that Harford and her colleagues frequently test different approaches of them with some of the Mad Scientist testing approaches described below. In short, Zillow’s Eureka stories are intimately tied into its business model and its business success.

THE MAD SCIENTIST STORY. We’re all familiar with the use of scientific testing in science-based industries such as pharmaceuticals. Drug companies test their products on a group of test subjects, while giving a placebo to members of a control group. They pay careful attention to ensure that people are randomly assigned to either the test or control group, so there are no major differences between the groups that might impact the drug’s effectiveness. It’s a powerful analytical tool because it’s usually as close as we can come to causation—the knowledge that what is being tested in the test group is driving the outcome in a causal fashion.

Rigorous testing is no longer just the province of white-coated scientists; it is now an analytical approach that every large organization can employ. There is broadly available software that leads managers or analysts through the testing process. Companies can now base important decisions on real, scientifically valid experiments. In the past, any foray into randomized testing (the random assignment to groups that we mentioned above) meant employing or engaging a PhD in statistics or a “design of experiments” expert. Now, a quantitatively trained MBA can oversee the process, assisted by software that will help determine what sizes of groups are necessary, which sites to use for testing and controls, and whether any changes resulting from experiments are statistically significant.

The mad scientist stories are particularly well suited to organizations like retailers (that have a lot of stores) and banks (that have a lot of branches). That makes it easy to try things out in some locations and use others as controls. It’s also quite easy to do testing on websites, where you can send some customers to one version of a Web page, send other customers to a different version, and see if the results are significantly different (called *A/B testing* in the Web analytics field).

Some examples of mad scientist stories include:³

- Do lobster tanks sell more lobsters at Food Lion supermarkets? The answer is apparently yes if the store was one in which customers already bought lobsters (i.e., they were relatively upscale), and no if the store didn’t attract lobster-buying customers to begin with.
- Does a Sears store inside a Kmart sell more than all-Kmart? Sears Holdings chairman Eddie Lampert is a big fan of randomized testing and has tested a variety of such combinations. We don’t know the answer to this particular

question, but we're guessing that if the answer were a definitive yes, we would have seen a lot more of these blended stores.

- Are the best sales results at the Red Lobster seafood restaurant chain achieved from a low-, medium-, or high-cost remodel of restaurants—and should the exterior or the interior be the primary focus? The result, according to Red Lobster executives, was that the medium-cost interior remodel paid off best. Exterior remodels brought a lot of new customers in, but if they saw that the interiors hadn't been redone as well, they didn't come back.

THE SURVEY STORY. Surveys are a classic method of quantitative research. The survey analyst observes phenomena that have already happened or are happening now. The analyst doesn't try to manipulate the outcome—only to observe, codify, and analyze it. Typically the surveyor seeks to understand what traits or variables observed in the survey are statistically related to other traits. The simplest example would be if we asked a sample of customers of a particular product various things about themselves, including demographic information like gender and age. If we also asked what products they liked, we could then determine whether men like certain products more than women, or whether certain products are more likely to be liked by younger people.

Surveys are popular and relatively easy to carry out. However, we have to remember that the results and stories based on them can vary considerably based on how questions are asked and how they vary (or not) over time. For example, the US Census has worked for literally decades on questions about the race of US citizens. The number of racial categories in census surveys keeps expanding; in the 2010 census there were fifteen choices, including “some other

race.” That was a popular choice for the more than 50 million Latino US citizens, 18 million of whom checked the “Other” box.⁴ If there is that much confusion about race, imagine what difficulties survey researchers can have with slippery topics such as politics, religion, social attitudes, and sexual behavior.

We also have to remember that just because two variables in a survey analysis are related, they may not be causally related. We’ll have more to say about this issue in chapter 6, but for now we’ll just point out that there may well be other variables that you’re not looking at that might be the causal factor driving the phenomena you care about.

Survey stories often involve asking people about their beliefs and attitudes, but they don’t have to involve people. Take, for example, this survey of airplanes conducted during World War II, related in a classic statistics textbook:

During the Second World War it was necessary to keep planes in action as much as possible, so it was decided to see if the number of time-consuming engine overhauls could be reduced without risk. A retrospective survey was made of planes that were lost, and contrary to all expectations, it was found that the number of planes lost as a result of engine troubles was greatest right after overhaul, and actually decreased as the time since overhaul grew longer. This result led to a considerable increase in the intervals between overhauls, and needless to say, to important revisions in the manner of overhauling to make sure that all those nuts and bolts were really tightened up properly.⁵

If you’re planning to do or analyze a survey, make sure that you’ve thought very carefully about the meanings of your survey questions or variables. A variable is any measured characteristic, with two or more levels or values, of properties of people, situations, and

behaviors. Gender, test scores, room temperature, love, happiness, and team cohesiveness are good examples of variables.

Also, it's important to ensure that your survey sample is representative of the population you want to study. How you perform the survey can affect the sample. For example, if you want to survey young people's attitudes or behaviors, don't hire a survey firm that only contacts the members of the sample through landline telephones. That's a very typical approach, but we all know that many young people don't have, and don't ever intend to have, a landline. So they would be underrepresented in a sample that employs only landlines.⁶

THE PREDICTION STORY. Prediction stories are all about anticipating what will happen in the future. While it's pretty difficult to get good data about the future, taking data about the past and understanding the factors that drive past events is pretty straightforward for quantitative analytics. Typically this is referred to as *predictive analytics* or *predictive modeling*.

There are a variety of prediction stories that an analyst can construct. Below is a sample of possibilities; note how specific they are:

- *Offer response:* Which customers will respond to an e-mail of a free shipping offer within two business days with a purchase of \$50 or more?
- *Cross-sell/upsell:* Which checking account customers with account balances over \$2,000 will purchase a one-year CD with an interest rate of 1.5 percent, responding within one month, given a mail solicitation?
- *Employee attrition:* Which employees of more than six months who haven't yet signed up for the 401(k) program will resign from their jobs within the next three months?

There are many other predictive analytics possibilities. In business, a common approach to prediction is to determine what offer the customer is most likely to accept. The most sophisticated versions of this “next best offer” analytics are increasingly automated; no human needs to see the offer before it is made available to the customer, and there can be hundreds or thousands of different offers.

Microsoft, for example, has an incredible ability to dynamically tailor “offers” for its Bing search engine (the product is free, so Microsoft is just trying to get you to use it). The offers tempt you to try out Bing, to create a Bing search bar on your browser, to try a particular Bing feature, and so forth. The customization of the offer is based on a variety of factors—including your location, age, gender, and recent online activity—that it can determine from your cookies and other sources. If you have signed up for Microsoft Passport, the company has even more information about you that allows for targeting the offers even more effectively. Microsoft is able (facilitated by the Infor Epiphany Interaction Advisor software they use) to instantly compose a targeted e-mail the moment you click on an offer in your inbox; it all takes about 200 milliseconds. Microsoft says it works extremely well to lift conversion rates.

Often, prediction stories can be a bit of a fishing expedition. We don’t know exactly what factors will allow us to predict something, so we try a lot of them and see what works. Sometimes the results are unexpected. For example, in the Microsoft Bing offers we’ve just described, the number of Microsoft Messenger buddies you have turns out to be a good predictor of whether you’ll try out Bing.

At Google, the company wanted to predict what employee traits predicted high performance. Some analysis determined that the factors Google was originally using—grades in college and interview ratings—were poor predictors of performance. Since they weren’t sure what factors would be important, they asked employees to

answer a three-hundred-question survey. As Laszlo Bock, the head of People Operations at Google, noted: “We wanted to cast a very wide net. It is not unusual to walk the halls here and bump into dogs. Maybe people who own dogs have some personality trait that is useful.”⁷

Bringing pets to work didn’t prove to predict much of anything, but Google did find some unexpected predictors. For example, whether a job applicant had set a world or national record or had started a nonprofit organization or club were both associated with high performance. Google now asks questions about experiences like these on its online job interviews.

Of course, if the factors that predict something make no sense at all, it’s a good idea to go back and recheck your data and your analysis. But actually looking at some data can outperform a human futurist’s predictions much of the time. As a caution, remember that predictive stories use data from the past to tell stories about the future. If something in the world has changed since you did your analysis, the predictions may no longer hold.

THE “HERE’S WHAT HAPPENED” STORY. Stories that simply tell what happened using data are perhaps the most common of all. They provide the facts—how many products were sold when and where, what were the financials that were achieved last quarter, how many people did we hire last year. Since they are reporting-oriented stories that often don’t use sophisticated math, it might seem that they would be easy to tell. However, the great rise in data within today’s organizations has been mirrored by a similar rise in reports based on data. Therefore, it’s sometimes difficult to get the attention of the intended audience for the reports you create or distribute.

This type of story is particularly well suited to visual displays of information. Suffice it to say that if you are providing reports in rows and columns of numbers, you aren’t likely to get the attention you

need. Many of us even tire today of colorful graphs and charts, but most people would say they are more worthy of attention than numbers on a page. Since chapter 4 is about communicating results, we'll say more about how to make this kind of report more interesting and attention-getting there.

The Scope of the Problem

By definition, a data-driven story and the quantitative analysis behind it can be somewhat narrow in scope, simply because it requires gathering data and applying it to a testable hypothesis (see “Examples of Testable Hypotheses”). It's difficult to gather data on very broad problems. However, it's important at this step not to prematurely limit the scope of the problem or decision. Thinking about the issue should be expansive, and you should have a number of alternative directions in mind. For example, if an organization recognizes a performance problem within a particular business unit or region, it should be open to a variety of causes of the problem—from customer dissatisfaction to operational issues to problems with products or services.

In the example of Transitions Optical at the end of this chapter, the problem recognition and framing step was prompted by a vague sense that marketing spending was too high, but the decision frame was expanded into one involving an overall optimization of marketing spending levels and media used.

We've referred to this first step in quantitative analysis as problem recognition, but it can also be an identification of opportunities. Joseph Jagger (1830–1892), a British engineer, realized that there was an opportunity to “break the bank” at the Monte Carlo casino.⁸ Jagger gained his practical experience of mechanics working in Yorkshire's cotton manufacturing industry. He extended his experience to the behavior of a roulette wheel, speculating that its outcomes

Examples of Testable Hypotheses

- The type of products that a customer has bought from us in the past year is the best guide to what e-mailed offers he or she will respond positively to in the future.
- Years of education is a good predictor of the level of performance rating an employee will receive in knowledge work jobs.
- Price markdowns of 10 percent made in the week before a holiday are less effective than those made at other periods.
- An end-cap display is the most effective placement of our product in a retail store for lifting weekly sales.
- Our customers can be grouped into four distinct segments with regard to the products they buy.
- Our ability to raise prices on a class of consumer staple products without hurting demand is significantly lower during economic recessions.
- Our business units that have centralized inventory management facilities tend to maintain lower average days of inventory for their production processes.

were not purely random sequences but that mechanical imbalances might result in biases toward particular outcomes. What if there were imperfections in the roulette wheel that he could exploit to his advantage? He went to Monaco to test this concept.

There are thirty-seven numbers in a French/European roulette wheel: 1–36 and 0. When a wheel is spun once, the theoretical probability that each number will come out is equal to $1/37$. Therefore the

proportion of each resultant number in a large number of spins should be roughly $1/37$. Jagger speculated that mechanical imbalances, if any, in wheels would cause specific numbers to appear more often than the probability of $1/37$.

With these thoughts in mind, Jagger hired six clerks to observe the six roulette wheels at the legendary Beaux-Arts Casino in Monte Carlo, each covering a different wheel. Each had specific instructions to record all of the results that came from each spin. When he analyzed the results, Jagger found that five of the roulette wheels produced the random results that one would expect. On the sixth wheel, however, he found that nine particular numbers (7, 8, 9, 17, 18, 19, 22, 28, and 29) appeared more often than mere chance could account for. Jagger concluded that the wheel was biased—that is, imperfectly balanced. He accordingly placed his first bets on July 7, 1875, and quickly won a considerable amount of money (£14,000—equivalent to around sixty times that amount in 2012, or over \$1.3 million, adjusted for inflation). The casino caught on to Jagger's betting strategy, and eventually neutralized it—but not before he had won the current equivalent of over \$6 million. Quite an analytical opportunity!

Getting Specific About What You Want to Find Out

While it's important to think expansively early in the problem recognition step, by the end of it you'll need to have created a clear statement of the problem, with concrete definitions of the key items or variables you want to study. Here's why: it makes a big difference how things are defined in quantitative research. For example, let's say you were a television executive interested in learning what channels consumers watched. Two different analytical consultants have approached you with proposals to learn the answer. Just for fun you decide to hire both of them to see how their results compare.



One consultant proposes to ask consumers to record (using either an online or a paper form) the actual channels and programs watched each day for a week. The other suggests asking the survey respondents to rank the channels they generally watch on television over the last several months. Both have well-designed survey samples that represent the desired population.

While these two consultants are trying to solve very similar problems, they are likely to come back with very different results. The one who proposes that consumers record actual programs and channels watched each day is likely to get more accurate results, but the extra burden of recording is likely to mean a lower level of participation from the survey sample. (Nielsen Media Research, which does channel and program monitoring on an ongoing basis, has about a 50 percent dropout level, and its recording is automated.) The other problem with this consultant is that viewing patterns might be overly influenced by the particular season or programming offered during the particular week of the study.

The other study is likely to be less accurate, but since it covers a broader time period, it is less likely to be influenced by seasonal factors. Most importantly, the results of the two surveys will probably be so different as to be difficult to reconcile. That's why it's important to finish the problem recognition step with a clear idea about what you want to study.

Review of Previous Findings Step

Once the problem is recognized, all the previous findings connected to it should be investigated. This is still a step within the first stage of analysis (framing the problem) because investigating previous findings can help analysts and decision makers think about how the



2. Review of previous findings



problem has been structured thus far, and how it might be conceptualized in different ways. Quite often, analysts will discover something in the review of previous findings that will lead to a substantial revision of the problem recognition step. That in turn can lead to a different set of previous findings.

Basically at this step we are asking, “Has a story similar to this been told before?” If so, we can get ideas for our own analysis. The review of previous findings can suggest any of the following:

- What kind of story could we tell? Does it involve prediction, reporting, an experiment, a survey?
- What kind of data are we likely to want to look for?
- How have variables been defined before?
- What types of analyses are we likely to perform?
- How could we tell the story in an interesting way that is likely to get results, and different from past stories?

One of the key attributes of quantitative analysis (and of the scientific method more broadly) is that it draws on previous research and findings. For example, searching thorough the problem-related knowledge appearing in books, reports, and articles is very important in getting to the bottom of the problem. It may help to identify relevant variables and any association among the identified variables.

A complete review of any of the previous findings is a must in any given quantitative analysis. You cannot make something out of nothing in analytics. You may only begin to solve the problem once you have a total grasp of the previous findings. Just remember one thing: your problem is not as unique as you think, and it's likely that many people have already done just what you are trying to do. Do not reinvent the wheel; what you need to do is search, search, and search again. These days, by using a search engine like Google, you can easily muster up most of the material related to your issue. By just arranging and evaluating the material, you can identify a potential model or approach to solve the problem.

An example of a successful review of previous findings took place during World War II. Adolf Hitler had ordered the production of a powerful new rocket bomb called the V-2, and in 1944 the Luftwaffe began to terrify the citizens of London. Over the next few months, 1,358 V-2s, out of at least 3,172 rockets distributed over the various Allied targets, flew out of the sky and landed in London, resulting in the death of an estimated 7,250 military personnel and civilians.

During the attack on London, many observers asserted that the points of impact of the bombs were grouped in clusters. The British were interested in knowing whether the Germans could actually target their bomb hits or were merely limited to random hits. If the Germans could only randomly hit targets, then deployment throughout the countryside of various security installations would serve quite well to protect the nation. But if the Germans could actually target their bombs, then the British were faced with a more potent opponent; the deployment of security installations would do little to protect them. The British government engaged statistician R. D. Clarke to solve this question. Clarke applied a simple statistical test based on his review—or existing knowledge—of previous findings.

Clarke was aware that the *Poisson distribution* could be used to analyze the distribution of bombs. The Poisson distribution expresses the probability of a number of events occurring in a fixed period of time, area, or volume if these events occur with a known average rate. The only thing we have to know to specify the Poisson distribution is the mean number of occurrences. If the bombs are falling randomly, the number of bombs that hit any particular small area follows a Poisson distribution. For example, if the average number of bombs that hit is 1 bomb per area, we can easily calculate the probabilities that no bomb will hit, exactly 1 bomb will hit, exactly 2 bombs will hit, exactly 3 bombs will hit, and exactly 4 or more bombs will hit, just by plugging these numbers in the Poisson formula.

To measure the number of bombs that may hit any specifically defined small area, Clarke divided South London into 576 squares of one-quarter square kilometer each, and counted the numbers of squares containing 0, 1, 2, 3, etc., flying bombs. If the targeting was completely random, then the probability that a square is hit with 0, 1, 2, 3, etc., hits would be governed by a Poisson distribution. The actual fit of the Poisson pattern for the data was surprisingly good, which lent no support to the clustering hypothesis (see the website for this book). The British were relieved by Clarke's conclusion. Fortunately, the Germans surrendered in 1945 before the V-2 could do much more damage. (*Note:* Despite its inability to be guided effectively, that rocket became the technical basis of the US space program.)

Just as Clarke did when he realized that the problem of the falling bombs could be described by a Poisson distribution, you can go back and review the problem recognition step after you have reviewed previous findings (see "Some Methods for Reviewing Previous Findings"). You may find that you need to modify your story, your problem scope, your decision, or even your stakeholders. If you have revised those a bit, or if you're still happy with the original problem

Some Methods for Reviewing Previous Findings

- Do an Internet search for key terms related to your analysis.
- Consult a statistics textbook for analyses similar to the one you're proposing.
- Talk to analysts around your company to see if they've done something similar.
- Check your company's knowledge management system if it has one.
- Talk about the problem with analysts at other (but noncompetitive) companies.
- Attend a conference (or at least look online at conference agendas) on analytics to see if anyone else is presenting on related topics.

definition, you can consider your problem framed and move along to actually solving it using quantitative analysis.

Reframing the Problem

Although we've laid out the analytical problem-solving process as a linear one of six steps in three stages, it is nothing if not iterative. Every step sheds new light on the problem, and it's always a good idea to think about how the new knowledge might shed light on previous steps. Although you can't spend forever reexamining each step, it's worth some time thinking about what the review of previous findings suggests about framing the problem (the "Worksheet for Framing the Problem" can help).

Worksheet for Framing the Problem

Have you framed the problem well? If so, you should be able to answer all or most of these questions positively:

1. Have you defined a clear problem or opportunity to address what is important to your business or organization?
2. Have you considered multiple alternative ways to solve the problem?
3. Have you identified the stakeholders for the problem, and communicated with them extensively about it?
4. Are you confident that the way you plan to solve the problem will resonate with the stakeholders, and that they will use the results to make a decision?
5. Are you clear on what decision is to be made—and who will make it—on the basis of the results from your analysis once the problem is solved?
6. Have you started with a broad definition of the problem, but then narrowed it down to a very specific problem with clear phrasing on the question to be addressed, the data to be applied to it, and the possible outcomes?
7. Are you able to describe the type of analytical story that you want to tell in solving this particular problem?
8. Do you have someone who can help you in solving that particular type of analytical story?
9. Have you looked systematically to see whether there are previous findings or experience related to this problem either within or outside your organization?

10. Have you revised your problem definition based on what you have learned from your review of previous findings?
-

For a good example, Rama Ramakrishnan, a retail analytics expert who is now CEO of the start-up CQuotient, describes a situation suitable for reframing in one of his blog posts:⁹

Take the “customer targeting” problem that arises in direct marketing. Customer targeting is about deciding which customers should be mailed (since mailing every customer is expensive). This is an old problem that has been studied by numerous researchers and practitioners. The most commonly used approach is as follows:

1. send a test mailing to a sample of customers
2. use the results of the test mailing to build a “response model” that predicts each customer’s propensity to respond to the mailing as a function of their attributes, past history etc.
3. use this model to score each customer in the database and mail to the top scorers.

This looks reasonable and may well be what the business cares about. But perhaps not.

The words “response model” suggest that the mailing *caused* the customer to respond. In reality, the customer may have come into the store and made a purchase anyway (I am thinking of multichannel retailers and not pure-play catalog retailers.



For the latter, without the catalog, it may be impossible for customers to make a purchase so the word “response” may be appropriate).

What these response models really do is identify customers who are likely to shop rather than customers likely to shop as a result of the mailing. But maybe what management really wants is the latter. For those customers who are either going to shop anyway or not going to shop regardless of what is mailed to them, mailing is a waste of money and potentially costs customer goodwill too. What the business may really want is to identify those customers who will shop if mailed, but won't if not mailed.

This re-framing of the customer targeting problem and approaches for solving it are relatively recent. It goes by many names—uplift modeling, net lift modeling—and the academic work on it is quite minimal compared to traditional response modeling. Yet, for many retailers, this is a more relevant and useful way to frame and solve the customer targeting problem than doing it the old way.

In this example, a thorough review of previous findings might have revealed the recent work on uplift and net lift modeling, and that might occasion a reframing of the problem. Ramakrishnan suggests that in such situations with relatively new modeling approaches, “Since the new problem hasn't received enough attention (by definition), simple algorithms may yield benefits quickly.”

We'll conclude this chapter on framing the problem with a couple of examples, one from business and one from law, in which the framing process was critical to the outcome. One is a good example of framing, and one is an example of incorrect framing. You haven't learned



much yet about the steps beyond the framing stage, but we're confident that you can make sense of them in these examples.

Analytical Thinking Example: Transitions Optical

One of the most common analytical problems in business is deciding how much to spend on a specific activity. And that's a particularly difficult decision for marketing spending. The department store founder John Wanamaker—and some European retailers before him—are renowned for saying, "Half the money I spend on advertising is wasted; the trouble is I don't know which half." Today, however, companies can use quantitative analysis to find out which marketing expenditures are effective, and which are not—and what the most effective combination of marketing expenditures is. This is typically called *marketing mix analysis*, and it's increasingly popular for firms that sell to consumers.

PROBLEM RECOGNITION AND FRAMING. Transitions Optical, which offers photochromic lenses for glasses, was getting some pressure from its corporate parents (Transitions is jointly owned by PPG and Essilor) with regard to its level of marketing spending. PPG, in particular, isn't in the business of consumer marketing, so that parent was especially skeptical about the cost and value of advertising and promotion. There were specific questions about whether particular advertising and marketing campaigns were effective or not. The overall intuitive feeling was that spending was too high, but there was no empirical data to answer the question of what level of marketing spend was optimal. Transitions executives decided to frame the problem as one of optimizing marketing expenditures and approaches in a way that maximized sales lift for the dollars invested. According to Grady Lenski, who headed Marketing at the time, "We were relying heavily on art to make marketing decisions; we needed more science."

REVIEW OF PREVIOUS FINDINGS. No previous findings on this topic existed; Transitions had customer data that would make such an analysis possible, but it was fragmented across the organization. Lenski and some of his colleagues were aware that it was possible to analyze the effectiveness of different marketing approaches, but didn't know the details.

MODELING (VARIABLE SELECTION). Marketing mix optimization models, which are increasingly employed by large organizations to optimize marketing spending, involve variables of marketing response, marketing costs, and product margins. The optimization models, using linear and nonlinear programming methods, find the weekly or monthly advertising, promotion, and pricing levels that maximize revenue, profit margin, or both. They also determine which particular advertising media are most effective for maximizing these outcomes. They also typically contain a series of "control" variables that might affect consumer spending and purchase behavior, such as weather and macroeconomic data.

DATA COLLECTION. This was one of the most difficult aspects of the analysis for Transitions, since the company works with intermediaries (optical labs, for example) and historically had little contact with or data about end customers. Hence, it couldn't accurately measure whether advertisements were seen by customers or whether they provided any sales lift. Transitions embarked upon a multiyear effort to gather end customer data from its channel partners (some of whom were competitors of its parent companies). Lenski had previously been head of the retail channel, so that facilitated gathering the information. The customer information came into Transitions in thirty different formats, but the company was eventually able to get it into an integrated data warehouse for analysis. Lenski commented

that the Marketing organization also needed to persuade different parts of the Transitions organization to provide data. The first time Transitions did the analysis, it did so without a data warehouse.

DATA ANALYSIS. Transitions hired an external consultant to do the data analysis, since it had no one in-house who was familiar with marketing mix optimization models. The analysis initially took several months, since the data had to be gathered and the model involves ruling out a wide variety of other explanatory factors for any marketing response (including weather, competitor marketing, etc.). Now that the models have been developed and refined, they can be finished in a few days.

RESULTS PRESENTATION AND ACTION. Transitions felt that interpreting and presenting the results was important enough to require in-house capabilities, so internal staff were hired to do it. The in-house experts take the model from the consultants and discuss it with executives to determine its implications and combine them with their intuitions about the market. Overall, the results have led to higher spending on marketing for Transitions, particularly for television advertising.

Analytical Thinking Example: *People v. Collins*

People v. Collins was a jury trial in California that made notorious forensic use of mathematics and probability, and it's a good example of how framing the problem incorrectly can lead to a bad outcome.¹⁰ The jury found defendant Malcolm Collins and his wife, Janet Collins, guilty of second-degree robbery. Malcolm appealed the judgment, and the Supreme Court of California eventually set aside the conviction, criticizing the statistical reasoning and disallowing the way the decision was put to the jury. We will examine this case within the six-step framework.

PROBLEM RECOGNITION. Mrs. Juanita Brooks, who had been shopping, was walking home along an alley in the San Pedro area. She was suddenly pushed to the ground by a person whom she couldn't see. She was stunned by the fall and felt some pain. Immediately after the incident, Mrs. Brooks discovered that her purse, containing between \$35 and \$40, was missing. A witness to the robbery testified that the perpetrators were a black male with a beard and moustache, and a Caucasian female with blonde hair tied in a ponytail. They had escaped in a yellow car. At the seven-day trial, the prosecution experienced some difficulty in establishing the identities of the perpetrators of the crime. The victim could not identify Janet Collins and had never seen her assailant; identification by the witness was incomplete. The prosecutor—perhaps desperate to save the case—decided to help the jury determine the probability that the accused pair fit the description of the witnesses.

REVIEW OF PREVIOUS FINDINGS. It is recognized that the court generally discerns no inherent incompatibility between the disciplines of law and mathematics and intends no disapproval or disparagement of mathematics as a fact-finding process of the law. There have been some criminal cases in which the prosecution used mathematical probability as evidence.

MODELING (VARIABLE SELECTION). The model suggested by the prosecutor is the probability that the accused pair fits the description of the witnesses.

DATA COLLECTION (MEASUREMENT). The prosecutor called to the stand an instructor of mathematics at a state college. Through this witness, he suggested that the jury would be safe in estimating the following probabilities of encountering the attributes of the criminals and crime:

Black man with beard	1 in 10
Man with moustache	1 in 4
White woman with pony tail	1 in 10
White woman with blonde hair	1 in 3
Yellow motor car	1 in 10
Interracial couple in car	1 in 1,000

DATA ANALYSIS. The mathematics instructor suggested that when events are independent, the probabilities of their happening together can be computed by multiplying each probability.

$P(A)$ = the probability that the accused pair fits
the description of the witness

$$= \frac{1}{10} \times \frac{1}{4} \times \frac{1}{10} \times \frac{1}{3} \times \frac{1}{10} \times \frac{1}{1000}$$

$$= \frac{1}{12,000,000}, \text{ or one in 12 million}$$

RESULTS PRESENTATION AND ACTION. The prosecutor arrived at a probability that there was only one chance in 12 million that any couple possessed the distinctive characteristics of the defendants. Accordingly, under this theory, it was to be inferred that there could be but one chance in 12 million that the defendants were innocent. The jury returned a verdict of guilty.

The Collinses appealed this judgment. The California Supreme Court thought that undoubtedly the jurors were unduly impressed by the mystique of the mathematical demonstration but were unable to assess its relevancy or value. The court set aside the conviction, criticizing the statistical reasoning and disallowing the way in which the decision was put to the jury. The Supreme Court pointed out that

the specific technique presented through the mathematician's testimony suffered from two important defects. First, the prosecution produced no evidence whatsoever showing the validity of the odds, nor evidence from which such odds could be in any way inferred. Second, there was another glaring defect in the prosecution's technique: an inadequate proof of the statistical independence of the six factors which were brought as evidence by the prosecution (e.g., bearded men commonly sport moustaches).

More importantly, the case and evidence had been framed incorrectly by the prosecutor. Even if the prosecution's conclusion was arithmetically accurate, it could not be concluded that the Collinses were the guilty couple. There was absolutely no guidance on a crucial issue: of the admittedly few such couples that might be encountered in the world, which one, if any, was guilty of committing this robbery?

The relevant variable in this case was not the probability that the accused pair fits the description of the witnesses, but the probability that there are other couples fitting the description of the witnesses, since the accused pair already fit the description. Depending on exactly how many couples there are in the Los Angeles area, the probability of at least one other couple fitting the description might be as high as 40 percent (see the website for this book). Thus the prosecution's computations, far from establishing beyond a reasonable doubt that the Collinses were the couple described by the prosecution's witnesses, imply a very substantial likelihood that the area contained more than one such couple, and that a couple other than the Collinses was the one observed at the scene of the robbery.

After an examination of the entire case, including the evidence, the Supreme Court determined that the judgment against the defendants must therefore be reversed. Bad framing of problems can clearly lead to bad decisions.

Notes

Chapter 2

1. S. Babou, "What Is Stakeholder Analysis?" The Project Management Hut, <http://www.pmhut.com/what-is-stakeholder-analysis>.
2. "Expedia Travels to New Heights," *SASCom Magazine*, Third Quarter 2011, 14.
3. All of these testing stories are examples from customers of Applied Predictive Technologies, a software company, though we approached the companies independently. For more on the methods behind this approach, see Thomas H. Davenport, "How to Design Smart Business Experiments," *Harvard Business Review*, November 2009.
4. Mireya Navarro, "For Many Latinos, Race Is More Culture Than Color," *New York Times*, January 13, 2012.
5. W.A. Wallis and H.V. Roberts, *Statistics: A New Approach* (New York: Free Press, 1960).
6. Lisa Carley-Baxter et al., "Comparison of Cell Phone and Landline Surveys: A Design Perspective," *Field Methods* 22, no. 1 (February 2010): 3–15.
7. Saul Hansell, "Google Answer to Filling Jobs Is an Algorithm," *New York Times*, January 3, 2007, www.nytimes.com/2007/01/03/technology/03google.html.
8. "Joseph Jagger," Wikipedia, http://en.wikipedia.org/wiki/Joseph_Jagger; "Joseph Jagger: The Man Who Broke the Bank," www.wildjackcasino.com/joseph-jagger.html; "Joseph Jagger," www.realmoneycasinos.net/joseph-jagger.html; "Roulette—The Men Who Broke the Bank at Monte Carlo—Joseph Jagger," www.wiseguyroulette.com/roulette-history/joseph-jagger/.
9. Rama Ramakrishnan, "Three Ways to Analytic Impact," *The Analytic Age* blog, July 26, 2011, <http://blog.ramakrishnan.com/>.
10. *People v. Collins*, 68 Cal. 2d 319 (1968); http://scholar.google.com/scholar_case?case=2393563144534950884; "People v. Collins," http://en.wikipedia.org/wiki/People_v._Collins.