PART I

CHAPTER 1

CHOOSING THE RIGHT TYPE OF FORECASTING MODEL

INTRODUCTION

Practical business forecasting is both a science and an art. It is a science in the sense that correct use of sophisticated statistical tools will invariably improve forecasting accuracy. It is an art in the sense that empirical data seldom if ever provide an unequivocal answer, so the user must choose between alternative relationships to select those equations that will provide the most accurate forecasts.

There are no perfect forecasts; they always contain some error. While perhaps that is obvious, it is nonetheless important to emphasize this fact at the outset. The point of this book is to show how to minimize forecast error, not to pretend that it can be eliminated completely. To accomplish this goal, a variety of forecasting methods may be used. In many cases, these methods will be complementary, not competitive.

Forecasts can be used for many purposes. Sometimes, predicting the direction of change is sufficient. For example, a model that could accurately predict the direction of the stock market the following day – even without providing any information about how much it would rise or fall – would be extremely valuable and profitable. No such model has ever been successfully constructed, although many have tried, and the goal will presumably remain elusive. At the other extreme, a model that predicted the direction of change in the consumer price index (CPI) the following month without forecasting the magnitude would be virtually useless, since over the past 40 years the monthly changes in the CPI have been negative only about 1 percent of the time.

There are many ways of forecasting, not all of which are based on rigorous statistical techniques. In some cases, informed judgment can provide the best forecasts, such as when "insiders" have company information that is not available to anyone else. Surveys may provide useful information about forecasts for the overall economy, specific sectors, or individual industries and firms. To the extent that these methods improve forecasting accuracy, they should be utilized.

Nonetheless, there is no rigorous way of testing how much informed judgments or survey techniques have boosted forecast accuracy, so they are mentioned only peripherally in this text. Instead, this text concentrates on illustrating how statistical and econometric methods can be used to construct forecasting models and minimize forecast errors. Initially, most economic forecasts were generated with structural equations; more recently, time-series analysis has been utilized more effectively. The benefits and shortcomings of both methods for generating optimal forecasts are identified.

This book is not a theoretical text; the emphasis is placed on practical business forecasting. As a result, theorems and proofs, which can be found in many other texts, will be kept to a minimum, with most of the material related to actual forecasting examples. In particular, this text will illustrate how statistical theory often needs to be adjusted to take into account those problems that recur in actual empirical estimation. Methods of adjusting the models to increase predictive accuracy are not to be denigrated or dismissed; they are an integral part of practical business forecasting.

1.1 Statistics, Econometrics, and Forecasting

Statistics is the application of probability theory and other mathematical methods to improve decision making where uncertainty is involved. Statistical theory and results are used widely in economics, but also apply to a large and diverse number of other disciplines, including sociology, agriculture, astronomy, biology, and physics.

The use of statistics is designed to provide answers where uncertainty exists, whether the uncertainty is due to randomness, or ignorance about the true underlying relationship that is to be tested. To illustrate the first case, we know the underlying probability distribution and hence the proportion of straights that will be dealt in a poker hand over the long run, but not what the next hand will show. To illustrate the second case, we probably do not know the true underlying relationship between capital spending and the rate of interest, or between the rate of inflation and the rate of unemployment, or changes in the value of the dollar and net exports. Cases where the underlying probability distribution is known are rare in economics.

Econometrics is the application of statistical and mathematical methods to the analysis of economic data to verify or refute economic theories. When structural equations are used, a specific theory is being tested for verification or rejection. By comparison, statistical methods are increasingly used with economic data to obtain parameter estimates without specifying any particular theory. Those models are usually known as time-series analysis; one standard technique is integrated autoregressive moving-average (ARIMA) models. Those models consist of correlating a given economic variable with its own lagged values, adjusted for trend and seasonal factors; no attempt is made to postulate an underlying theory.

Economic forecasting often relies on statistical or econometric methods, but even that need not be the case. Some types of forecasts do not involve mathematical techniques at all; for example, surveys or polls may produce valuable forecasts without utilizing any econometric methods. However, these types of forecasts are not featured in this book. Most of the examples will be confined to those types of forecasts that use statistical methods.

1.2 The Concept of Forecasting Accuracy: Compared to What?

No forecast is ever perfect; opinions about what will happen in the future invariably contain errors. Anyone who has ever attempted to predict anything knows that. On the other hand, forecasting can be quite useful if it provides better answers than alternative methods of guessing about the future. The relevant test for any forecast, then, is never whether the results contain errors, but how accurate they are compared to the alternatives. Like that old Henny Youngman one-liner "How's your wife?" the appropriate answer is always "Compared to what?"

Throughout this book, the difference between the *science* of statistics and econometrics and the *art* of forecasting is emphasized. Most of the sophisticated theorems and proofs in those fields are based on highly unlikely assumptions about the distribution of the error terms, and furthermore assume that the data generating process remains the same in the sample and the forecast periods. Adjusting models to generate better forecasts when these assumptions are not satisfied has often been disdainfully called ad hoc adjustment, unworthy of the name of econometrics. Yet it plays a vital role in improving forecast accuracy.

From 1940 through 1970, primary emphasis was placed on theoretical refinements of statistical and econometric procedures, with scant attention paid to systematic methods of adjusting forecasts to improve their accuracy. When macroeconomic models proved unable to predict any of the major changes in the economy in the 1970s, the emphasis gradually shifted to developing methods that produced useful forecasts even if they did not follow the theoretical procedures developed in earlier decades.

In *Forecasting Economic Time Series*, a reference book recommended for those with more advanced mathematical skills, Clements and Hendry[1] have classified the basic issues in forecasting, as opposed to econometrics. They state that "The

[1] Clements, Michael P., and David F. Hendry, *Forecasting Economic Time Series* (Cambridge University Press, Cambridge, UK), 1998.

features of the real-world forecasting venture that give rise to the sources of forecast error . . . induce a fairly radical departure from the literature on 'optimal' forecasting . . . and at the same time help to explain why some apparently ad hoc procedures work in practice" (p. 3).

The approach used here is much less mathematically rigorous than Clements and Hendry's. Also, the discussion of forecasting accuracy begins with structural models and then moves to time-series analysis, contrary to the procedure that they (and others) use. Yet the methodology in which real-world practical forecasting is approached is very much in the spirit of their approach. While the method of least squares is used for the vast majority of the examples, the reader should always keep in mind that the assumptions of the classical linear model seldom hold in practical business forecasting.

In many cases, the underlying data generating function has shifted, the variables are not normally distributed, the residuals are not independent, and the independent variables are not known at the time of the forecast. Even more important, repeated rerunning of regression equations, and the substitution of different empirical data series that measure the same theoretical concept, often help to improve forecasting accuracy but are outside the constructs of the classical linear model. For this reason, the statistical estimates generated under those assumptions must be interpreted carefully, and often with a degree of skepticism.

It is too crude to say that what makes the best forecasts is "whatever works," but that is closer to the spirit of the present approach than the method of choosing rigorous statistical techniques that minimize the root mean square error or other similar measures in the sample period but generate suboptimal forecasts. Sometimes structural econometric models provide better forecasts, and sometimes the results from ARIMA models with no economic structure are better. In certain cases, the *combination* of these methods will generate better forecasts than either method separately. Far from being relegated to the criticism of ad hoc adjustments, changing the model during the forecast period will invariably produce better results than a "pure" unadjusted model, provided it is done properly.

As Newbold and Granger have written,[2] "the evaluation criteria employed should be as demanding as possible since the object ought to be self-criticism rather than self-congratulation" (p. 266). The principal aim should be to build a forecasting model that will generate the smallest forecasting error, not necessarily maximize the goodness-of-fit statistics over the sample period.

The reader should always keep in mind that any forecasting model, no matter how sophisticated the underlying statistical techniques, must perform better than forecasts generated by random variables or naive methods. That means

[2] Granger, C. W. J., and Paul Newbold, *Forecasting Economic Time Series* (Academic Press, San Diego, CA), 1986.

always checking whether the model provides better results than other methods that are available – including naive models, surveys, and qualitative judgments.

A naive model generally assumes that the level or rate of change of the variable to be predicted this period will be the same as last period, or the change this period will be the same as the average change over an extended time period. For a time series without any significant trend, such as the Treasury bill rate, a naive model might state that the bill rate this month will be the same as it was last month. For a time series with a significant trend, the naive model would usually be couched in terms of percentage changes. For example, a naive model might state that the percentage change in the S&P 500 stock prices index next month will equal the percentage change last month, or it might equal the average percentage change over the past 480 months. A more sophisticated type of non-structural model incorporates regression equations using lagged values of the variable that is to be predicted. If more complicated modeling techniques cannot generate forecasts that beat these naive models, the model building attempt is presumably not worthwhile.

For people engaged in industry and finance, where having more accurate forecasts than your competitors will materially improve profitability, forecasts are useful if they provide results that are more accurate than the competition's. A model that accurately predicted the direction of change in the stock market the next day 60 percent of the time would be tremendously valuable – even though it would be wrong almost half the time – regardless of the methodology used to develop those predictions. In a similar vein, calculations by this author have shown, in some semi-annual polls of economists published in the *Wall Street Journal*, over 50 percent of the forecasts incorrectly predicted the *direction* interest rates would change over the next six months. Hence any model that could even predict the direction in which interest rates would move over the next several months would significantly improve the current status of forecasting financial markets.

Yet the decision not to forecast at all means throwing in the towel and claiming that any deviations from past trends can never be predicted. That would be the case only if the variable in question always grew at the same rate and was never subject to exogenous shocks. For even if changes are truly unexpected (an oil shock, a war, a wildcat strike, a plant explosion) forecasting models can still offer useful guidance indicating how to get back on track. Virtually everyone in a management or executive role in business or finance makes guesses about what will happen in the future. While these guesses will never be perfect, they are likely to be much improved if the practitioner combines robust statistical techniques with the ability to adjust the forecasts when actual events do diverge from predicted values.

Forecasting makes practitioners humble. That does not mean people who choose forecasting as a profession are necessarily humble; the opposite is more likely to be true. But unlike economic theories, which can often persist for decades without anyone ever being able to verify whether they are accurate or useful, forecasters generally find out quickly whether or not their opinions are correct.

Since highly visible forecasts of the overall economy or financial markets have compiled a very unimpressive track record over the past 30 years, it is sometimes argued that predicting economic variables is not a useful exercise. Indeed, most consensus forecasts of real growth, inflation, and interest rates have not been much better than from a naive model. In view of these results, some have concluded that forecasting models do not work very well.

Before reaching that conclusion, however, we should try to determine what causes these forecasting errors. For example, suppose the majority of forecasters thought interest rates would rise because inflation was about to increase. The Federal Reserve, also expecting that to happen, tightened policy enough that inflation decreased and, by the time six months had elapsed, interest rates actually fell. I am not suggesting this always occurs, but it is a reasonable hypothesis. Thus before beginning our analysis of how to reduce forecasting errors, it is useful to categorize the major sources of these errors. Some may be intractable, but others can be reduced by a variety of methods that will be explored in this book.

When the econometric model and the mechanism generating the model both coincide in an unchanging world, and when the underlying data are accurate and robust, the theory of economic forecasting is relatively well developed. In such cases, the root mean square forecasting error in the forecast period ought not to be any larger than indicated by the sample period statistics.

This does not happen very often; in the majority of forecasts, the actual error is significantly larger than expected from sample period statistics. In some cases that is because the model builder has used inappropriate statistical techniques or misspecified the model through ignorance. Most of the time, however, unexpectedly large forecasting errors are due to some combination of the following causes:

- structural shifts in parameters
- model misspecification
- missing, smoothed, preliminary, or inaccurate data
- changing expectations by economic agents
- policy shifts
- unexpected changes in exogenous variables
- incorrect assumptions about exogeneity
- error buildup in multi-period forecasts.

1.2.1 Structural Shifts in Parameters

Of the factors listed above, structural shifts in parameters are probably the most common. These may occur either within or outside the sample period. For example, sales at Ace Hardware will drop dramatically when Home Depot opens a store three blocks away. At the macroeconomic level, a recession used to be accompanied by a stronger dollar; now it is accompanied by a weaker dollar. Company profits of American Can were influenced by completely different factors after it became a financial services company.

Perhaps stated in such stark terms, structural shifts are obvious, but most of the time the changes are more subtle. For 1997 through 1999, macroeconomists thought the growth rate of the US economy would slow down from about 4% to the 2–2½% range; yet each year, real growth remained near 4%. Forecasters thought that with the economy at full employment, inflation would increase, causing higher interest rates, lower stock prices, and slower growth, yet it did not happen. At least in retrospect, there were some structural shifts in the economy. For one thing, full employment no longer produced higher inflation. Also, the technological revolution boosted capital spending and productivity growth more rapidly. Yet even after several years, the consensus forecast failed to recognize this shift.

1.2.2 Model Misspecification

Model misspecification could be due to the ignorance of the model builder; but even in the case where the best possible model has been estimated, some terms might be omitted. In many cases these might be expectational variables for which data do not exist. For example, economists agree that bond yields depend on the expected rate of inflation, a variable that cannot be measured directly. A company might find that cutting prices 5% would not invoke any competitive response, but cutting them 10% means competitors would match those lower prices. The missing variable in this case would be the trigger point at which competitors would respond – which itself is likely to change over time.

It is also possible that the underlying model is nonlinear. In one fairly straightforward and frequently documented case, purchases of capital goods with long lives (as opposed to computers and motor vehicles) generally increase faster when the rate of capacity utilization is high than when it is low. At the beginning of a business cycle upturn, capital spending for long-lived assets is often sluggish even though interest rates are low, credit is easily available, stock prices are rising rapidly, sales are booming, and profits are soaring. Once firms reach full capacity, they are more likely to increase this type of capital spending even if interest rates are higher and growth is slower.

To a certain extent this problem can be finessed by including variables that make the equation nonlinear, and I will discuss just such an example later. For example, investment might grow more rapidly when the rate of capacity utilization is above a certain level (say 85%) than when it is below that level. However, the situation is not that simple because a given level of capacity utilization will affect investment differently depending on the average age of the capital stock, so using a simple rule of thumb will generally result in model misspecification. An attempt to pinpoint the exact rate at which capital spending accelerates is likely to result in data mining and the resultant penalty of relatively large forecast errors.

1.2.3 Missing, Smoothed, Preliminary, or Inaccurate Data

The data used in estimating forecasting models generally comes from one of three major sources. Most macroeconomic data are prepared by agencies of the Federal government, including the Bureau of Economic Analysis (BEA), the Bureau of the Census, and the Federal Reserve Board of Governors. Financial market data on individual company sales and earnings are prepared by individual corporations. In an intermediate category, many industry associations and private sector institutions prepare indexes of consumer and business sentiment, and measures of economic activity for specific industries or sectors; perhaps the best known of these are the Conference Board index of consumer attitudes and the National Association of Purchasing Managers index of business conditions in the manufacturing sector.

Except for specific data based on prices given in financial markets, virtually all macroeconomic or industry data are gathered by sampling, which means only a relatively small percentage of the total transactions is measured. Even when an attempt is made to count all participants, data collection methods are sometimes incomplete. The decennial census is supposed to count every person in the US, but statisticians generally agree the reported number of people in large cities is significantly less than the actual number; many of the uncounted are assumed to be undocumented aliens. Thus even in this most comprehensive data collection effort, which is supposed to count everyone, some errors remain. It is reasonable to assume that errors from smaller samples are relatively larger.

Virtually all macroeconomic and industry data series collected and provided by the government are revised. The issuing agencies named above make an attempt to provide monthly or quarterly data as quickly as possible after the period has ended. These releases are generally known as "advance" or "preliminary" data. In general, these data are then revised over the next several months. They are then revised again using annual benchmarks; these revisions usually incorporate changing seasonal factors. Finally, they are revised again using five-year censuses of the agricultural, manufacturing, and service sectors. In addition, some of the more comprehensive series, such as GDP and the CPI, may be revised because of methodological changes.

The revisions in the data prepared and released by the Federal government are often quite large. Sometimes this is because preliminary data, which appears shortly after the time period in question has ended, are based on a relatively small sample and then revised when more comprehensive data become available. In other cases, seasonal factors shift over time. Data revisions quite properly reflect this additional information.

Most government data are collected from surveys. From time to time, respondents do not send their forms back. What is to be done? The sensible solution is to interpolate the data based on those firms that did return their forms. The problem with this approach is that, in many cases, it is precisely those firms that failed to return their forms that faced unusual circumstances, which would have substantially modified the data. Eventually the problem is solved when more complete numbers are available, but the initial data are seriously flawed.

Sometimes, the methodology is changed. In October 1999, a comprehensive data revision boosted the average growth rate of the past decade by an average of 0.4% because the Bureau of Economic Analysis (BEA) – the agency that prepares GDP and related figures – decided to include software purchased by businesses as part of investment; previously it had been treated as an intermediate good and excluded from GDP. Since software had become an increasingly important part of the overall economy, this change was appropriate and timely.

In another important example, the methodology for computing the rate of inflation was changed in the mid-1990s. As a result, the same changes in all individual components of the CPI would result in an overall inflation rate that was 0.7% lower. These changes reflected the improved quality of many consumer durables, shopping at discount malls instead of department stores, and changes in market baskets that included a higher proportion of less expensive goods. Most economists agreed these changes were warranted, and many thought they were overdue. A commission headed by former Chairman of the Council of Economic Advisers Michael Boskin calculated that the rate of inflation had been overstated by an average of 1.1% per year.[3]

The Federal government statisticians cannot reasonably be criticized for including improved information and methodology in their data releases when they become available. Indeed, failure to include these changes would be a serious error. Nonetheless, the appearance of preliminary data that are later revised substantially raises significant issues in both building and evaluating forecasting models. At least in the past, it has sometimes had a major impact on policy decisions.

For example, one of the major examples of misleading preliminary data occurred in the 1990–1 recession. During that downturn, BEA initially indicated the recession was quite mild, with a dip in real GDP of only about 2%. Subsequent revisions revealed that the drop was much more severe, about 4%.[4] Acting on the data that were originally reported, the Fed assumed the slump was not very severe and hence eased cautiously. If it had known how much real

---

[3] Boskin, Michael J., E. R. Dulberger, R. J. Gordon, and Z. Griliches, "The CPI Commission: Findings and Recommendations," *American Economic Review Papers and Proceedings*, 87 (1997), 78–83.

[4] This result was not entirely a surprise. Joseph Carson, an economist at Chemical Bank who had previously worked at the Commerce Department, stated at the time that he thought real GDP was declining at a 4% annual rate in late 1990.

GDP had really fallen, it probably would have eased much more quickly. Indeed, when the recovery failed to develop on schedule in 1991, the Fed did reduce short-term interest rates to unusually low levels by the end of 1992, and the economy finally did recover. However, that boosted inflationary pressures, causing the Fed to tighten so much in 1994 that real growth plunged to 1% in the first half of 1995. Not until the latter half of that year did the economic effects of those incorrect data completely disappear.

The most accurate forecast would have said the economy is in worse shape than the government reports indicate, so initially the Fed will not ease enough and hence the economy will be slow to rebound, which means the Fed will eventually have to ease more than usual, so two years from now interest rates will be much lower than anyone else expects, in which case inflationary expectations will rise and the Fed will have to tighten again. Of course no one said that, nor could anyone have reasonably been expected to predict such a sequence of events.

This example clearly indicates how inaccurate data can cause poor forecasts. Yet economists were roundly criticized for underpredicting the severity of the recession, overpredicting the initial size of the rebound, and failing to gauge the decline change in interest rates accurately. No forecaster won plaudits following that recession, but it is not unreasonable to suggest that forecast errors would have been smaller with more accurate data.

In May 1974, the wage and price controls imposed by the Nixon Administration ended. As a result, the producer price index (PPI) rose by a record amount that month. For the next few years, the seasonal adjustment program used by the Bureau of Labor Statistics (BLS) assumed the PPI always rose sharply in May, so the seasonally adjusted data for the May PPI showed a big dip, while the unseasonally adjusted data were virtually unchanged. In this case perhaps the obvious solution would have been to ignore those data, but it is not clear what method the forecaster should use. Running regression equations without any data for May? Using seasonally unadjusted data? Treating May 1974 with a dummy variable – e.g., 1 in that period and 0 elsewhere? All these are possible, but none is optimal.

Of course, it is not only the Federal government that revises their data. Companies often restate their earnings for a variety of reasons. They book sales when they ship the goods, but if they aren't paid for, writeoffs must be taken. Sometimes reorganizations, or sales of divisions, result in huge one-quarter writeoffs. Other times, accounting errors are at fault. Analysts try to take these anomalies into account, but most if not all attempts to predict stock prices based on reported company earnings suffer from the changes and inconsistencies in these data.

There will never be any perfect solution to the issue of data revisions. Nor does it make any sense to castigate government statisticians for providing the most accurate estimates possible based on incomplete data and the changing nature of the economy. Nonetheless, a few observations relating to data revisions are appropriate at this point.

1. Evaluation of forecast accuracy should take into account the data at the time when the forecasts were issued, rather than the most recently revised data. This means, for example, that an attempt to evaluate forecasting accuracy of macroeconomic forecasts made many years ago provides far different results depending on which set of "actual" data are used.
2. Some, although certainly not all, forecast error stems from the assumptions of changes in fiscal and monetary policy that are based on the preliminary data issued by the government. Later revisions of these data sometimes make it appear that those assumptions were unwarranted.
3. When estimating a structural model over an extended period of time, it is useful and appropriate to use dummy or truncated variables in the regression equation. For example, the methodological changes in the CPI that began in 1994 can be entered explicitly as an additional variable; before 1994, any such variable would have the value of zero.

1.2.4 Changing Expectations by Economic Agents

This is often cited as one of the major reasons given for the failure of macroeconomic modeling in the 1970s and the 1980s. It has been argued that economic forecasts based on past historical evidence cannot be accurate because people adjust their behavior based on previous events, and thus react differently to the same phenomena in the future. This concept is generally known as the Lucas Critique;[5] however, it was formulated by Oskar Morgenstern in 1928,[6] so it is hardly a recent idea. Formally, we can express this concept by saying that the data generation process underlying the model has changed during the sample period, or between the sample and the forecast periods. I mention the roots of this concept to emphasize that it far predates the idea that mismanaged monetary policy in the 1950s and 1960s was the primary factor that caused the short-term tradeoff between inflation and unemployment.

Indeed, the Lucas Critique is just a special case, although an extremely well-known one, of changing expectations. Economic agents often change their behavior patterns based on what has happened in the past. That is not only true at the macro level. Growth in individual company sales will be significantly affected as competitors enter and exit the industry. Firms will raise or lower prices depending on how their competitors react. Borrowers may have a higher or lower rate of default on loans depending on recent changes in the bankruptcy laws.

Lucas and others, and Morgenstern before them, claimed that econometric models would not work whenever economic agents learned from previous

[5] Lucas, Robert E., "Some International Evidence on Output–Inflation Tradeoffs," *American Economic Review*, 63 (1973), 326–34.

[6] Morgenstern, Oskar, *Wirtschaftsprognose: eine Ubtersuchung ihrer Vorussetzungen und Moglichkeiten* (Julius Springer, Vienna), 1928.

experience and adjusted their behavior accordingly in the future. Yet many economic links continue to hold over a wide variety of different experiences. On a ceteris paribus basis, consumers will spend more if income rises, although admittedly their increase in consumption will be greater if they think the change is permanent rather than temporary. If interest rates rise, capital spending will decline. If the value of the currency increases, the volume of net exports will decline. If the growth rate for profits of an individual firm accelerates, the stock price will rise. There are many similar examples where structural relationships continue to hold in spite of changing expectations.

Sometimes, a change in expectations in one area of the economy will generate changes in other sectors that are consistent with past experience. One major example of this occurred in the US economy in the second half of the 1990s. Expectations about future profit growth shifted significantly, so that the price/earnings ratio of the stock market doubled even though bond yields were at just about the same level in 1995 and 2000. Few forecasters were able to predict that change. On the other hand, the rise in stock prices and the decline in the cost of equity capital impacted consumer and capital spending in a manner consistent with previous historical experience. In addition, the more rapid growth in capital stock stemming from an increase in the ratio of capital spending to GDP boosted productivity growth, which reduced the rate of inflation and lowered interest rates further. That in turn boosted real growth enough that the Federal budget position moved from a deficit to a surplus, which further boosted equity prices. Predicting the change in the stock market was difficult; but given that change, predicting more robust growth in the overall economy was more straightforward. Conversely, when the stock market plunged, all of the reverse factors occurred – lower capital spending, a slowdown in productivity, and a return to deficit financing.

1.2.5 Policy Shifts

Anyone who tries to estimate an equation to predict short-term interest rates will soon find that, during the mid-1970s, Fed Chairman Arthur Burns used monetary policy to offset the recessionary impact of higher oil prices, leading to unusually low real interest rates; whereas in the early 1980s, Chairman Paul Volcker refused to accommodate the further increase in oil prices, leading to unusually high real interest rates. The real Federal funds rate equals the nominal rate minus the change in inflation over the past year. Its pattern is shown in figure 1.1.

No model estimated on data through 1979 would have predicted the massive increase in real interest rates that started in late 1980. With hindsight, of course, one can include a well-chosen set of economic variables that track this pattern, but that is not the point. In July 1980, the Blue Chip consensus forecast of the six-month commercial paper rate for 1981 was 8.7%; the actual figure was 14.8%. This is one of the clearest policy shifts that ever occurred in the US economy.

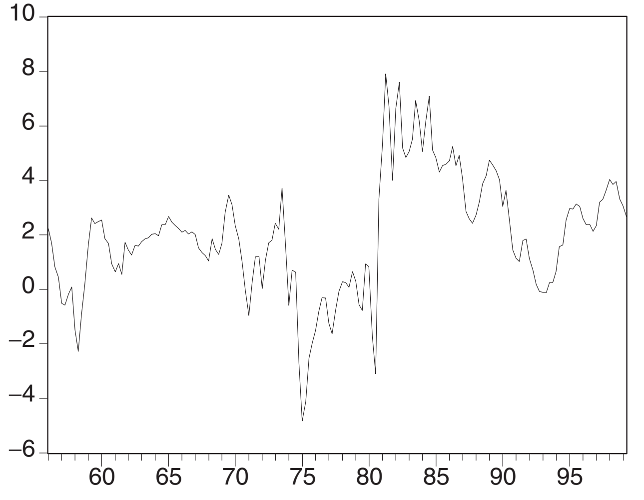

Figure 1.1 The real Federal funds rate.

What lessons can forecasters learn from this experience? In the short run, fluctuations in short-term interest rates are determined primarily if not exclusively by the action of the Federal Open Market Committee. That is why short-term interest rate forecasting today is reduced to a series of guesses about what the Fed will do next. In the long run, however, we learn another lesson. If the Fed holds short-term rates at below equilibrium for an extended period, eventually both inflation and interest rates will rise; whereas if it holds short-term rates above equilibrium, eventually both inflation and interest rates will decline. In this case, a model that captured this underlying relationship would provide very little guidance in predicting interest rates in the short run, but would be useful in the long run. In particular, a forecast that interest rates and inflation would start to decline in 1982, hence setting in motion the biggest bull market in history, would have been particularly valuable. Yet hardly anyone believed the few forecasters who accurately predicted that development.

Even the best econometric model is not designed to predict the impact of unexpected policy or exogenous changes in the short run. However, once these changes have occurred, correctly structured models should be able to offer valuable insights into what will happen in the longer run.

1.2.6 Unexpected Changes in Exogenous Variables

The change in Fed policy under Paul Volcker is a classic example of a policy change initiated by the government. As seen by forecasts made at the time, it was a major surprise. Another major example of an unexpected policy change, although it occurred over several years, was the decision by senior officials in the Clinton Administration to reduce the level of real per capita government spending during his tenure as President.[7] That changed the budget deficit to a budget surplus, which (as already noted above) was one of the factors causing an almost unprecedented increase in the price/earnings ratio of the stock market.

Changes of this sort are undertaken by government officials. However, other shocks that affect the economy have nothing to do with policy, such as energy shocks, wars, and natural disasters. Unless foreseen, they will not be incorporated in any forecasts. Yet if they were predicted, vigorous action would be taken to offset or eliminate these developments.

I have already noted how the Fed acted quite differently to the first and second energy shocks in 1973–4 and 1979–80 respectively. However, that was not the only change; private sector economic agents also reacted differently. The first energy shock was viewed by most consumers and businesses as a once-in-a-lifetime event, so they did not alter their behavior patterns very much. As a result, oil imports continued to increase, and eventually oil prices rose again. After the second energy shock, attitudes changed significantly. Most people now expected that massive price increases would continue on a regular basis, and forecasts were common that oil prices would rise to $100/bbl by the end of the twentieth century. As a result, both consumers and businesses started using less energy, buying more fuel-efficient motor vehicles, and constructing more fuel-efficient buildings. Those plans were successful enough to reduce oil imports, so in 1986 energy prices plunged by more than half. In 1998 they were lower in real terms than in 1972, before the first energy shock occurred.

Any forecast of the economy in the 1980s – whether right or wrong – was influenced by the assumption about energy prices. However, this example indicates the value of some of the alternative types of forecasts discussed in section 1.3: conditional vs unconditional, point vs interval, and alternative scenarios weighted by probabilities. An appropriate way for many businesses to proceed would have been to generate alternative forecasts based on different scenarios about oil prices: higher, steady, or lower. When prices gradually started to decline in the mid-1980s as the worldwide energy glut increased, more weight would have been given to the lower-price scenario, so businesses would have been better prepared when crude oil prices suddenly fell by more than half in 1986.

There is little to be gained by pointing out that forecasts are inaccurate when they fail to predict unexpected exogenous shocks, many of which would never have occurred if they had been accurately predicted. However, models that correctly analyze the impact of these shocks when they do happen can still be quite useful in indicating what lies ahead.

---

7 Most of these changes were suggested by Treasury Undersecretary, and then Secretary, Robert Rubin.

1.2.7 Incorrect Assumptions About Exogeneity

In some cases, models designed for forecasting generate much larger errors than would be indicated by the sample period statistics because some of the independent variables are assumed to be exogenous when they really are not. Technically, an exogenous variable is one whose value is not determined within an economic model, but which plays a role in the determination of the endogenous variables. However, as a practical matter, there are degrees of exogeneity. Only a relative handful of variables, such as weather conditions and defense expenditures, are exogenous in all circumstances. Most of the time, policy variables have some endogenous components as well.

For example, foreign currency values are often considered to be exogenous. After the collapse of the Thai baht, Korean won, Indonesian rupiah, and Malaysian ringgit in the latter half of 1997, US net exports declined dramatically in 1998 and the first half of 1999. As a result, manufacturing production rose much more slowly than total GDP; whereas during boom years, production usually rises faster than overall GDP. North Carolina has the highest proportion of workers in manufacturing, so its growth rate fell sharply after the collapse of those currencies.

A model that linked growth in North Carolina employment to the value of the dollar (among other variables) would show a high correlation. However, a forecast made in 1997 would have been inaccurate if it had assumed the values of those currencies would remain stable. In such a case, the model would appear to work well, but forecasts of the North Carolina economy would be far off the mark. In this case, the equations might have continued to work well in the sense of high correlations and low standard errors, but the forecasts would have been poor because of the inability to predict the exogenous variables.

In the past, monetary policy used to be treated as exogenous, although this error is made far less often today. Even in the days before Paul Volcker, the Fed routinely tightened monetary policy as inflation increased. Thus assuming that monetary policy variables were exogenous and would not change invariably led to forecast errors that were much larger than expected from the sample period statistics.

1.2.8 Error Buildup in Multi-period Forecasts

Analyses of macroeconomic models undertaken many years ago by this author showed that the single biggest source of error in multi-period forecasting was caused by using the lagged dependent variable on the right-hand side of the equation. If current consumption were estimated as a function of lagged consumption, for example, an error made one quarter could distort all the forecasts from that point forward. I will discuss a variety of methods to overcome that difficulty; now that this error has been well documented, it does not occur so much in multi-period forecasting models. Nonetheless, it is an error that beginning modelers often commit.

1.3 Alternative Types of Forecasts

When most people think of forecasts, they think of point estimates. For example, sales will rise 12% next year, the Dow will climb to 12,000 a year from now, the Federal Open Market Committee will vote to boost the Federal funds rate 25 basis points at its next meeting, the price of oil will climb 20% over the next six months, and so on.

While it is true that point estimates are the most common type of forecasts, there are many other ways in which forecast results can be presented. Sometimes a range for the predicted variable is more appropriate; other times, different probabilities are assigned to alternative scenarios. Sometimes the penalties associated with being too high or too low are equal; at other times, the loss function is asymmetric. In this section, I discuss some of the more common types of alternative forecasts.

1.3.1 Point or Interval

Suppose a company has a limited amount of excess manufacturing capacity. If sales grow less than 5% per year, the company will be better off using its existing facilities. If sales grow more than 5% per year, it will be better off building a new plant. In this case, the point estimate for sales growth is not as important as the probability that sales growth will exceed 5%.

A similar case might be made for advertising budgets. If a firm thinks a $1 million expenditure on advertising will boost sales by at least $5 million, it will decide to go ahead and spend the money. It doesn't matter so much whether the increase in sales is $6 or $10 million, but if it is $4 million, the expenditure will not be made.

At the macro level, suppose the Fed decides that 3% is the highest level of inflation that is tolerable. If inflation rises 1%, 1½%, or 2%, there will be no change in monetary policy. If it exceeds 3% – or if it appears likely it will soon exceed 3% if policy is not changed – the Fed will boost short-term interest rates.

A company may have a loan covenant with the bank stating that if cash reserves drop below a certain level, the loan will be called. That level might be correlated with the assumption of increased profitability, so a decline in profits would trigger the loan call. In that case, the key forecast is whether company profits have risen or not, rather than the precise amount they would increase.

1.3.2 Absolute or Conditional

Forecasts can be either absolute or conditional. Some examples of absolute, or unconditional forecasts are: real GDP will grow 4% next year, the Republicans will retain (or regain) majority control of Congress, and company sales will rise at least 15% per year over the next decade. However, many forecasts are issued on a conditional basis: real GDP will grow 4% next year if the Fed does not tighten, the Republicans will be the majority party in Congress if they also capture the Presidency, and sales will grow if competitors do not double their capital spending and advertising budgets.

The choice of which type of forecast is appropriate will depend largely on how the results are to be used. A speculator in financial markets wants to know whether prices will rise or fall, not whether they will rise or fall under certain circumstances. An automobile dealer wants to know what lines of vehicles will sell most quickly, so he can optimize his ordering procedure. A pharmaceutical company wants to know how rapidly a new drug will be adopted.

Conversely, conditional forecasts can often be quite useful. Firms might want to determine how fast sales are likely to grow under normal business conditions, using those results as guidelines for rewarding superior performance. If sales are then affected by some exogenous event, guidelines can be adjusted accordingly. Forecasts of production planning might be determined based on the assumption that materials are delivered on time, compared with what might happen if a strike occurred. The most common way of delivering conditional forecasts is by using alternative scenarios, as discussed next.

1.3.3 Alternative Scenarios Weighed by Probabilities

A forecast that sales will rise 8% if the economy booms, rise 6% if real growth remains sluggish, and fall 2% if there is a recession may appear to be an excuse to avoid offering a firm forecast at all. However, that is not always true. In many cases, firms need to be prepared to take appropriate action if the economy falters even if the probability of that occurring is relatively low.

Based on the historical forecasting record of macroeconomists, it would appear that recessions were not predictable. Consider the case of a lending institution involved in sub-prime auto loans. As long as the economy remains healthy, the vast majority of these loans will be repaid; if a recession strikes, the loss rate will rise enough to put the company out of business. Prudence might dictate less risky loans; but if the company is too picky, it will lose business to competitors and won't make enough loans to stay in business.

In this case the most appropriate procedure would be to assess the probability of a recession occurring next year. If it were only 5%, then the lending institution would continue to expand its sub-prime loan portfolio. On the other hand, if it were to rise to 25%, some trimming would be in order. Note that in this case the probability of an actual downturn the following year is well below 50%, yet some adjustment in corporate strategy is warranted.

The alternative-scenario method of forecasting can also be used for long-range planning, since long-term economic forecasts are generally little more than trend extrapolations in any case. The company might discover that the probability of meeting its stated goal of a 15% annual gain in sales and earnings would occur only if the most optimistic macroeconomic forecast, with a probability of only 10%, were met. The company could then make plans to move into faster-growing areas of the economy or, alternatively, trim its ambitious long-term goals to more realistic levels.

1.3.4 Asymmetric Gains and Losses

So far we have been assuming that a forecast error of +8% carries the same penalty as an error of –8%. Often, however, that is not the case. For many companies, if sales increase faster than expected, that is fine; but if they don't, disaster strikes. I have already described such a situation for a sub-prime auto lending company. The same general type of argument could be applied to municipal bonds; as long as the community tax base grows above a certain rate, the interest and principal will be repaid, but if it dips below that rate, the issuing authority will default on the bonds.

In many companies, the rewards for exceeding the plan are substantial: bonuses, promotions, and larger budgets for next year. Similarly, the penalties for failing to meet planned targets are severe, including loss of employment. In a situation of that sort, many planners will set targets below their predicted level, so they will appear to have exceeded their goals. Eventually, management may catch on to this trick and fire all the planners, which is another risk. Nonetheless, the percentage of plans that are exceeded compared with the percentage that are not met strongly suggests that corporate planners are well aware of the asymmetric loss function.

Money managers may face a similar dilemma. If they beat the benchmark averages – Dow Jones Industrials, S&P 500, or Nasdaq composite index – they are handsomely rewarded; investors will switch their assets into those funds, and salaries and bonuses rise. If their performance falls short of the gains posted by the major averages, they will lose customers and possibly their own jobs.

This is not just a hypothetical example. The so-called January effect occurs because many money managers aggressively buy growth stocks early in the year (or the previous December) and, if they can show substantial gains, lock in those gains and buy the equivalent of index funds for the rest of the year. In the same vein, very few money managers who are already ahead of the average for the first three quarters of the year would take risks in the fourth quarter that would jeopardize their hefty year-end bonuses.

1.3.5 Single-period or Multi-period

So far we have not specified how many time periods in the future are being predicted. That can make a great deal of difference in the way a model is formulated. In models used to forecast only one period ahead, it might well be appropriate to use the lagged value of the variable that is being predicted. Interest rates in the next period might very well depend on rates this period, as well as on other variables such as the inflation rate, growth rate, unemployment rate, value of the currency, budget surplus or deficit, and other relevant variables.

However, suppose the model is used to predict interest rates on a monthly basis for the next 12 months. In this case, the forecasts for interest rates later in the year would depend on "lagged" values of interest rates that were not known at the time of forecast. For example, suppose the forecast made at the beginning of March for interest rates depends on the level of interest rates in January and February. As the year progresses, the forecast for interest rates in June would depend on their level in April and May, which are not yet known.

For this reason, using the lagged dependent variable for multi-period forecasts causes serious difficulties that do not exist for the single-period forecast. That does not rule out the use of lagged dependent variables on an a-priori basis, but it does raise a red flag. One of the tenets of the classical linear model, as will be shown in the next chapter, is that the values of all the independent variables are known at the time of forecast. Obviously that is not the case when the lagged dependent variable is used in multi-period forecasting. Hence it is advisable to use a different approach when multi-period forecasts are required.

1.3.6 Short Run or Long Range

To a certain extent, the difference between short- and long-run forecasts can be viewed as the difference between single- and multi-period forecasting. However, whereas short-term forecasts are more generally concerned with deviations from trends, long-run forecasts are often designed to predict the trend itself. As a result, different methods should be used.

One of the principal goals of short-term forecasting, and one that has been emphasized by time-series analysis, is to remove the trend from time-series variables so the underlying properties of the series may be properly examined. If company sales have been growing an average of 12% per year, the challenge in short-term forecasting is to indicate how much sales next year will deviate from that trend. Long-range forecasters, on the other hand, might want to determine how many years it will take for the trend growth in sales to diverge from that 12% average gain. The difference is analogous to the split responsibilities of the COO, who asks "How are we doing?", and the CEO, who asks "Where are we heading?"

In large part, then, the method of building forecasting models will be different depending on whether the primary goal is short-term or long-range forecasting. In general, the same model will not be optimal for attempting both goals.

1.3.7 Forecasting Single or Multiple Variables

In the models discussed above, it has been implicitly assumed that the independent variables – the variables on the right-hand side of the equation – are either known in advance or are truly exogenous. In the case of financial decision or qualitative choice models, actual information is entered for economic and demographic data. In the case of sales forecasting models, the variables are either exogenous to the firm or are determined by management decisions.

In the case of macroeconomic and financial forecasting models, however, that assumption is not generally valid. Interest rates depend on expected inflation, which is generally not known. Net exports depend on the value of the currency, which also is not known. In cases of this sort, it is necessary to build multi-equation models in order to explain all the endogenous variables in the system. In the case of macro models, some variables are generally treated as exogenous, such as changes in fiscal and monetary policy, but even these are often related to the state of the economy. Only variables such as wartime expenditures, energy shocks, or weather conditions are truly exogenous.

1.4 Some Common Pitfalls in Building Forecasting Equations

Before turning to a brief review of statistics, I will illustrate some of the most common pitfalls that occur in estimating regression models for forecasting. These topics will be treated in a more rigorous fashion after the statistical groundwork has been prepared, but it is useful to introduce them initially so they can be kept in mind as the statistical and econometric exposition unfolds.

I have already noted that there is no such thing as a perfect forecast. Even if all of the statistical methods are applied correctly, some random error will occur. This error can be quantified and measured for any existing data set, and can be used as an estimate of the forecast error that can be expected. In the vast majority of cases, though, the actual forecasting error is larger than is indicated by the regression equation or econometric model. Some of the major reasons for unexpectedly large forecast error are discussed next.

The residuals in any stochastic equation, which are supposed to be independent, may be correlated with each other. As a result, there are far fewer independent observations than indicated by the statistical program. Hence the goodness-of-fit statistics are overstated, and the forecasting errors are understated. Structural relationships estimated with time-series data – consumption as a function of income, prices as a function of unit labor costs, or interest rates as a function of the inflation – are all likely to have serially correlated residuals. Because consumer spending patterns, for example, change slowly over time, the number of independent observations is probably far less than the sample period data would indicate. Consequently, the standard errors are significantly understated.

Virtually all statistical and econometric tests are based on the underlying assumption that the residuals are normally distributed. Often, however, that is not the case. That is another reason why the calculated goodness-of-fit statistics overstate the robustness of the equation.

The "law of large numbers" indicates that as the sample size increases, all distributions with a finite variance tend to approach the normal distribution. However, that is scant comfort to those who must deal with relatively small samples. Furthermore, some financial market data do not have bounded data; in particular, percentage changes in daily stock prices are not normally distributed. Every once in a while, an unexpected event will cause a much larger change than could be expected from past history – especially in financial markets. Such distributions are colloquially referred to as "fat tails." Estimates based on the assumption of a normal distribution when that is not the case are likely to generate disappointing forecasts.

Spurious correlation may destroy the usefulness of any model for forecasting, even if the sample period statistics appear to provide a remarkably accurate fit. Many studies have shown that series that actually have no correlation – because they were generated from random number series – can provide highly significant goodness-of-fit statistics if enough alternative regressions are calculated. This problem has become particularly virulent in the PC era, where it is a simple matter to run hundreds if not thousands of regression equations very quickly.

The problem of "data mining" has also run rampant because of quick and inexpensive computing power. This issue always represents somewhat of a dilemma. One does not want to test only one or two versions of any given equation. After all, the theory may not be precisely specified; and even if the long-run determinants are well determined, the lag structure and adjustment process may not be known. Empirical approximations of theoretical concepts may not be precise, so it is logical to try several different measures of the concept in question. Also, research results are often improved when alternative specifications were tried because the first attempt did not produce reasonable results. Yet having provided all these reasons for diligent research, it is much more likely that econometricians and statisticians will "torture the data until they confess" instead of failing to calculate the necessary minimum number of regressions. Such attempts at curve fitting seldom produce useful forecasting equations.

Sometimes the equation fits very well during the sample period, and the goodness-of-fit statistics hold even in the forecast period, yet the equation generates very poor forecasts because the values of the independent variables are not known. For example, sales growth for a particular company or individual product line is likely to change if competitors react to an erosion of their market share. At the macroeconomic level, financial markets certainly will react differently to anticipated and unanticipated changes in policy. Consumers are likely to alter their spending patterns based on what they think will happen in the future as well as changes in current and lagged income and monetary conditions.

It is not very helpful to develop theories that produce optimal forecasts under severely stylized sets of assumptions that are rarely encountered in the real world. Practical business forecasting invariably consists of two interrelated steps: use of standard statistical theory that has been developed based on restrictive assumptions, followed by modification of that theory to improve actual forecasting accuracy. These two steps cannot be considered in isolation. Thus even in this introductory chapter, I have pointed out some of the major pitfalls that commonly occur in forecasting models. Further details will be provided throughout the text.

The following examples are indicative of many cases where robust economic theories, which have been verified by sophisticated econometric methods, do not generate accurate forecasts unless they are further modified.

- *Example 1.* Economic theory says that the riskless long-term interest rate is related to the underlying growth rate of the economy, the Federal budget deficit ratio, and the expected rate of inflation. Econometrics can be used to test this theory. However, it cannot be used for forecasting unless, in addition, we can find an accurate way to predict the expected rate of inflation. Essentially the same comments could be made for forecasting the stock market, foreign exchange rates, or commodity prices. Since inflationary expectations are not formed in a vacuum, they could presumably be tied to changes in economic and political variables that have already occurred. So far, no one has been very successful at this.
- *Example 2.* The price of oil is tied to the world demand and supply for oil, which can perhaps be predicted accurately by econometric methods, using the geopolitical situation of Saudi Arabia *vis-à-vis* the US and other major powers as a major factor in the forecast. However, world economic hegemony cannot be predicted econometrically – and probably cannot be predicted very well with any method – so this is not a useful forecasting model. Certainly no one in the early 1980s publicly predicted the fall of the Berlin Wall by the end of the decade.
- *Example 3.* Historically, the growth rate for PCs, modems, and other high-tech equipment can be accurately tracked over the sample period by identifying the time when major innovations were introduced and matching their performance to various growth curves. In the future, since the timing of such innovations is unknown, such a set of regression equations would not serve as a useful forecasting model.
- *Example 4.* Economic theory says that the value of the dollar depends on relative real interest rate differentials; the *higher* the real rate in the US, the more likely it is that the dollar will appreciate. However, economic theory also says

that a stronger dollar will attract capital from abroad, hence resulting in a *lower* interest rate than would otherwise occur. Both of these theories can be verified separately, but unless further adjustments are made they are useless for predicting either the value of the dollar or interest rates, since they lead to opposite conclusions. This is indicative of a larger problem in forecasting, where an individual theory may provide robust empirical results *in isolation* but may be useless for forecasting because the factors that are being held constant in the theory are in fact always changing.

These examples provide a flavor of the problems of building a practical forecasting model. Many of the examples involve interrelationships between several variables that must be predicted simultaneously. However, even in the cases where the independent variables are actually known ahead of time, and in that sense are truly exogenous, model builders often go astray by failing to realize the spurious correlation introduced by common trends in several of the time series.

Using econometrics to build forecasting models is deceptively difficult. As Clive Granger has put it, "econometric modeling is an activity that should not be attempted for the first time."[8] It takes practice to develop useful forecasting models.

Problems and Questions

1. As an economist, you are asked to prepare quarterly forecasts for the next two years for shipments of oil-drilling equipment. Data on company and industry shipments are available back to 1959. Figure 1.2 shows the relationship between constant-dollar shipments of oil-drilling equipment and the relative price of crude oil.

(a) Would you prepare an unconditional or conditional forecast? If the latter, for what variables would you prepare alternative scenarios?
(b) How would you generate forecasts of oil prices?
(c) In general, would you predict that the next time oil prices rise sharply, shipments of oil-drilling equipment would rise rapidly as they did in the 1970s and the 1980s?

2. The loan portfolio of a bank has been growing at an average of 10% per year. The bank officers would like to expand growth to 15% per year,

continued

[8] Granger, C.W. J., *Forecasting in Business and Economics*, 2nd edn (Academic Press, San Diego, CA), 1989.

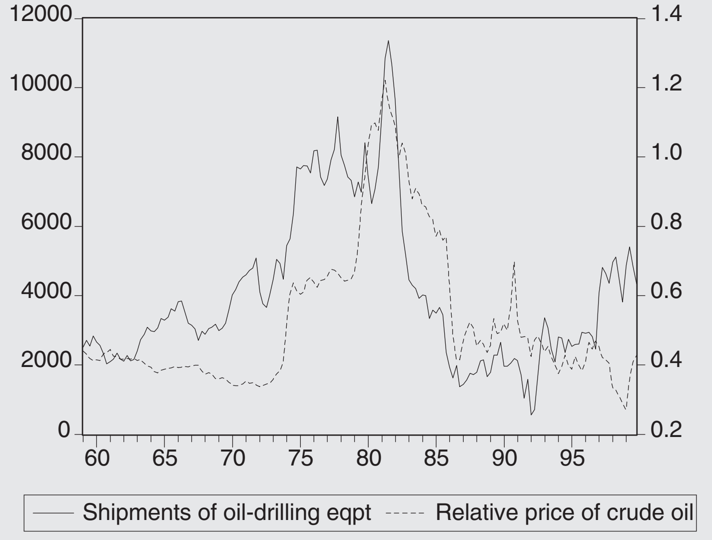

Figure 1.2 Problem 1.

and have asked you to develop a model that would evaluate credit risks on personal loans more accurately. The CEO also points out that the default rate on loans has fallen from 11% in 1990 to 2% in 1999.

(a) How would your advice to the CEO differ if (i) the consensus forecast called for a continuation of 5% growth and a further decline in the unemployment rate, (ii) a decline in the growth rate to $2\frac{1}{2}$% and a gradual rise in the unemployment rate, or (iii) the reemergence of recession next year?

(b) Suppose your result showed that credit card loans could be doubled by reducing the APR from 13.9% to 6.9%. Under what circumstances would you recommend that move, and under what circumstances would you advise against it?

(c) The two largest banks in the metropolitan area have merged and have significantly increased the average monthly charge on checking accounts. What information would you need to determine whether it would be advisable to offer "free" checking accounts in an attempt to obtain more customers who would then borrow money from the bank? To what extent would the likely macroeconomic outlook influence this decision?

3. A hotel chain would like to determine whether to build a major new hotel in Las Vegas or Orlando; essentially its decision will be based on whether the "gambling" market or the "family entertainment" market is

continued

expected to grow faster. Historical data are available on the number of trips to each city, the occupancy rate for hotel rooms, the average amount spent per traveler, the proportion arriving by automobile and airplane, and total available rooms.

(a) What other variables would be useful in making this determination? (Hint: what proportion of the visitors are foreign?)
(b) Suppose one presidential candidate said he thought gambling was sinful, whereas the other said consumers should have freedom of choice. How would that influence your forecast?
(c) In recent years, Florida voters have decided against permitting on-shore casinos in the state, but the issue will be voted on again in the future. How would that affect your forecast?
(d) The recent building boom in Las Vegas could affect the decision either way. On the negative side, there may be excess hotel rooms. On the positive side, more hotel rooms may attract more business conventions and stay for vacations. How would you determine whether the building boom should have a positive or negative impact on the decision?

4. As a financial manager, your client wants you to recommend the purchase of Mutual Fund X, Y, or Z. From 1997 through 1999, Fund X earned a total rate of return of 32% per year, Fund Y has earned 19% per year, and Fund Z has earned 8% per year.

(a) What additional data would you need to make an informed choice of which fund to purchase?
(b) Further analysis shows, not surprisingly, Fund X outperforms the market when it is rising at above-average rates, but underperforms it otherwise, whereas the opposite is true for Fund Z. However, your client has made it clear he does not want a conditional forecast. How would you proceed?
(c) The day before your client calls, the Federal Reserve Board has just voted to raise the funds rate by 50 basis points. How does this influence your decision? (Hint: was the change expected or unexpected?)

5. Your client is a state government, and legislators are debating whether to raise the tax on cigarettes by an additional $1/pack. Proponents of the bill claim that (i) needed revenue will be raised without boosting other taxes, and (ii) higher prices will reduce smoking, which will benefit the general health of society. However, as an economist, you know there is a tradeoff: the more people who quit, the less additional revenue will be raised.

continued

(a) How would you estimate the price elasticity of demand?
(b) Would you expect the price elasticity to be larger for younger or older smokers? How would this affect your overall conclusion?
(c) How would your answer differ if (i) Nevada were the client, and it was planning to raise its tax rate but California was not, or (ii) California were the client, and it was planning to raise its tax rate but Nevada was not?
(d) How would your estimates of tax revenue in the short run and the long run vary?