Contents

1
Introduction

By a forecast will be meant any statement about 'the future', where the future is relative to the analyst's viewpoint. So as well as the common sense notion of a forecast of what will happen tomorrow, or next Saturday, the term will equally apply to the outcome of the 1997 General Election made *now* but based on what was known at the end of 1996, for example. Forecasts are often constructed *ex post* as a way of evaluating a particular forecasting model or forecasting device, presumably with the hope that the past forecast performance of the model will serve as a useful guide to how well it might forecast in the future. In any event, forecasting the past as the 'relative future' means that forecasts can be evaluated as they are made, without having to wait to see what actually happens tomorrow, or on the coming Saturday, and a large sample of forecasts can be generated (with associated outcomes available), which might allow a statistical analysis of the forecast performance of the model. My forecast of rain might turn out to be wrong, but that might just be bad luck. Suppose my forecasting model is that I forecast rain in the afternoon if at 11 a.m. in the morning the cows in a certain field are lying down. Given daily observations on afternoon rainfall and the morning stance of cows over the last year, one could devise a statistical test of whether my forecasting model was a good predictor of meteorological conditions.

Forecasts can be made about anything, and using a variety of means: systems of dynamic equations, back-of-the-envelope calculations, tea-leaf dregs, goats' entrails. Our subject matter will be economic and financial variables, such as output growth rates, unemployment and inflation rates, and stock returns. The means will include econometric models and survey-based expectations. The key issue will be the evaluation of the forecasts. That is, how we judge 'good' in relation to forecasts, and how we decide whether a certain set of forecasts have this property.

The method of evaluation may depend on whether the forecasts are model-based, as well as depending on the type of forecast being made. A forecast defined as 'any statement about the future' includes statements such as: the probability that it will rain is 80% (a probability forecast); that it will rain (an event forecast); that there will be $\frac{1}{4}$ cm of rain (a point forecast); that there is a 75% probability that there will be between 0 cm and $\frac{3}{4}$ cm of rain (an interval forecast).

The material in this book will be structured by the type of forecast. Chapter 2 begins with the evaluation of point forecasts. These are typically quantitative forecasts of the level or rate of change of a continuous variable, such as the level of output or the rate of growth of output, the price level or the inflation rate, but we will also include in this chapter 'direction of change, tests, although these might be more correctly thought of as event forecasts. It is probably fair to say that the traditional concern of economic forecasting has been the production and evaluation of point forecasts, and that it is only relatively recently that there has been a general recognition that some measure of the degree of uncertainty surrounding a 'central tendency' will enhance the value or usefulness of the forecast. For example, the government might react rather differently to a point forecast that inflation will be $2\frac{1}{2}$% next year, but that the forecaster believes there is a 40% chance that it will exceed 5%, compared to the same point forecast and the assertion that the outcome will almost certainly be within $\pm\frac{1}{2}$ percentage point of $2\frac{1}{2}$%. These issues are taken up in subsequent chapters.

Chapter 3 switches attention from the evaluation of the (conditional) mean of the random variable to the evaluation of forecasts of the conditional variance of the process. For a large number of financial time series, as well as some macroeconomic time series (such as inflation), the conditional variance (or volatility) varies over time in a way that is in part predictable from the past of the process. Models of conditional variance are briefly reviewed as a precursor to a discussion of forecast evaluation. A complicating factor is that volatility is not observed.

In Chapter 4 interval forecasts or prediction intervals come under the spotlight. An interval forecast is a formal method of conveying forecast uncertainty. An interval forecast can be used to express the uncertainty surrounding a point forecast of the conditional mean, or of a volatility forecast. Viewed as an estimate of a quantile of the conditional distribution of the random variable, a one-sided interval forecast is an estimate of the 'Value-at-Risk' in the financial risk management literature.

Chapter 5 considers the evaluation of forecast densities, or forecast probability distributions. We review methods of evaluation that make no

recourse to the method of construction of the forecasts. These methods are clearly appropriate when the forecasts come from surveys, or when the models or methods underlying their construction are unknown to the econometrician. A number of recent papers have proposed the evaluation of models' forecast densities as model specification tests, and these are also reviewed.

Finally, Chapter 6 recognizes that forecasts are generally used to guide actions (or decisions) in uncertain environments, and should ideally be evaluated in terms of the benefits (or costs) that result or are expected to result from using them in this way. This approach is still in its infancy in terms of applications in macroeconomics, but an exploration of the 'decision-based' approach and its connections with more standard approaches is illuminating.

Index