"Elements of Forecasting" by Francis X. Diebold (2000)

We make this operational as

$$(\hat{\varphi}^2 y_T + \hat{\varphi}\hat{\theta}\varepsilon_T) \pm 1.96\,\hat{\sigma}\sqrt{1 + (\hat{\varphi} + \hat{\theta})^2}$$

5. APPLICATION: FORECASTING EMPLOYMENT

Now we put our forecasting technology to work to produce point and interval forecasts for Canadian employment. Recall that the best moving-average model was an MA(4), whereas the best autoregressive model, as well as the best ARMA model and the best model overall, was an AR(2).

First, consider forecasting with the MA(4) model. In Figure 8.1, we show the employment history together with operational 4-quarter-ahead point and interval extrapolation forecasts. The 4-quarter-ahead extrapolation forecast reverts very quickly to the mean of the employment index, which is 100.2. In 1993.4, the last quarter of historical data, employment is well below its mean, but the forecast calls for a quick rise. The forecast quick rise seems unnatural, because employment dynamics are historically very persistent. If employment is well below its mean in 1993.4, we would expect it to stay well below its mean for some time.

The MA(4) model is unable to capture such persistence. The quick reversion of the MA(4) forecast to the mean is a manifestation of the short

FIGURE 8.1 Employment History and Forecast: MA(4) Model

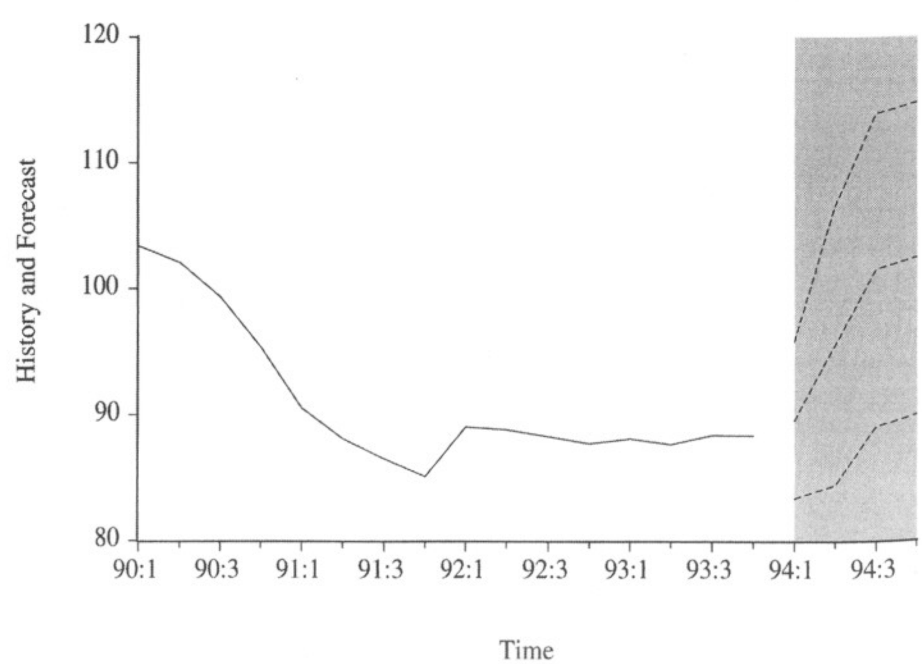

FIGURE 8.2 Employment History and Long-Horizon Forecast: MA(4) Model

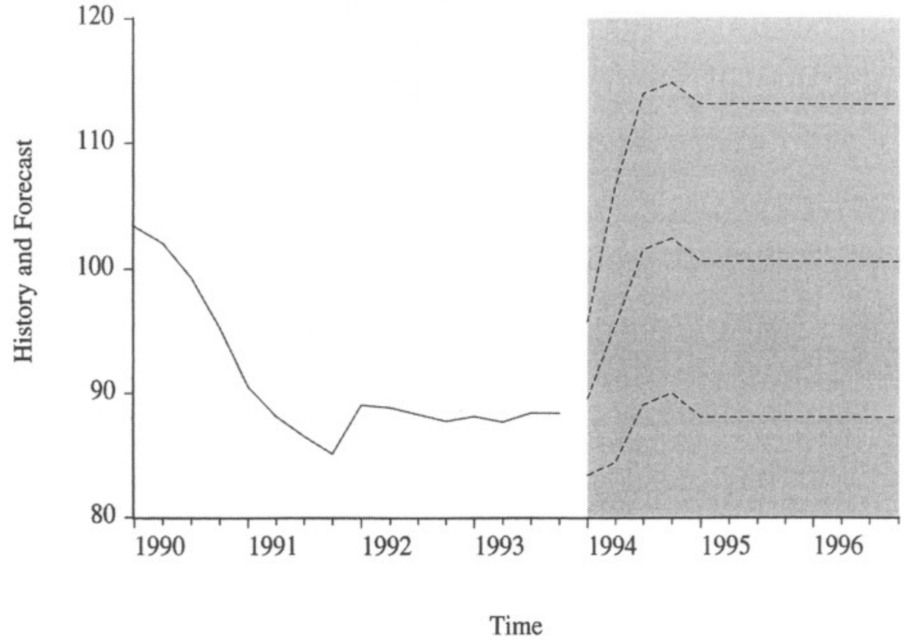

memory of moving-average processes. Recall, in particular, that an MA(4) process has a four-period memory—all autocorrelations are zero beyond displacement 4. Thus, all forecasts more than four steps ahead are simply equal to the unconditional mean (100.2), and all 95% interval forecasts more than four steps ahead are ±1.96 unconditional standard deviations. All of this is made clear in Figure 8.2, in which we show the employment history together with 12-step-ahead point and interval extrapolation forecasts.

In Figure 8.3 we show the 4-quarter-ahead forecast and realization. Our suspicions are confirmed. The actual employment series stays well below its mean over the forecast period, whereas the forecast rises quickly back to the mean. The mean squared forecast error is a large 55.9.

Now consider forecasting with the AR(2) model. In Figure 8.4 we show the 4-quarter-ahead extrapolation forecast, which reverts to the unconditional mean much less quickly, as seems natural given the high persistence of employment. The 4-quarter-ahead point forecast is, in fact, still well below the mean. Similarly, the 95% error bands grow gradually and have not approached their long-horizon values by four quarters out.

Figures 8.5 and 8.6 make clear the very different nature of the autoregressive forecasts. Figure 8.5 presents the 12-step-ahead extrapolation forecast, and Figure 8.6 presents a much longer horizon extrapolation forecast. Eventually the unconditional mean is approached, and eventually the error bands *do* go flat, but only for very long horizon forecasts, due to the high persistence in employment, which the AR(2) model captures.

FIGURE 8.3 Employment History, Forecast, and Realization: MA(4) Model

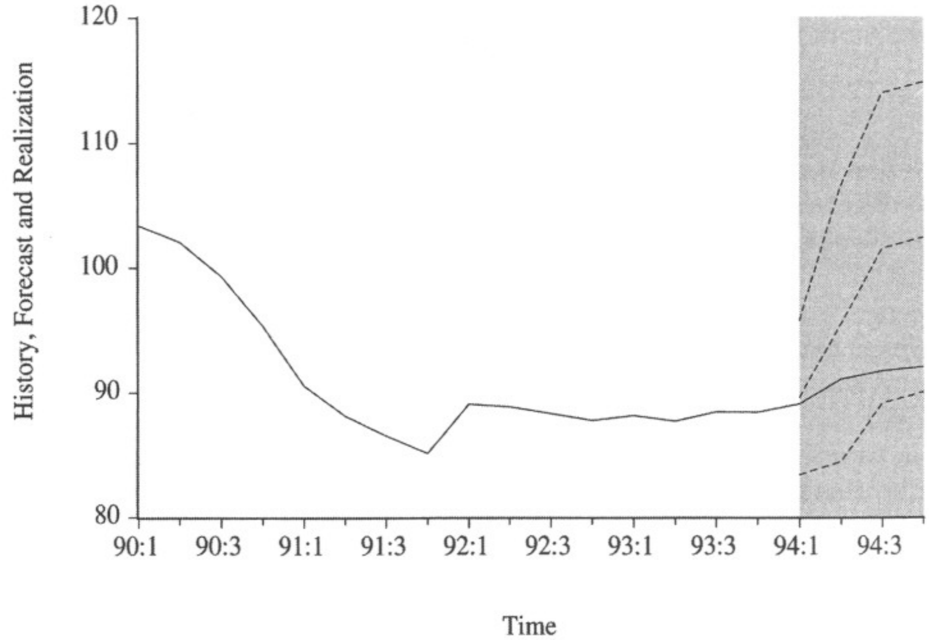

FIGURE 8.4 Employment History and Forecast: AR(2) Model

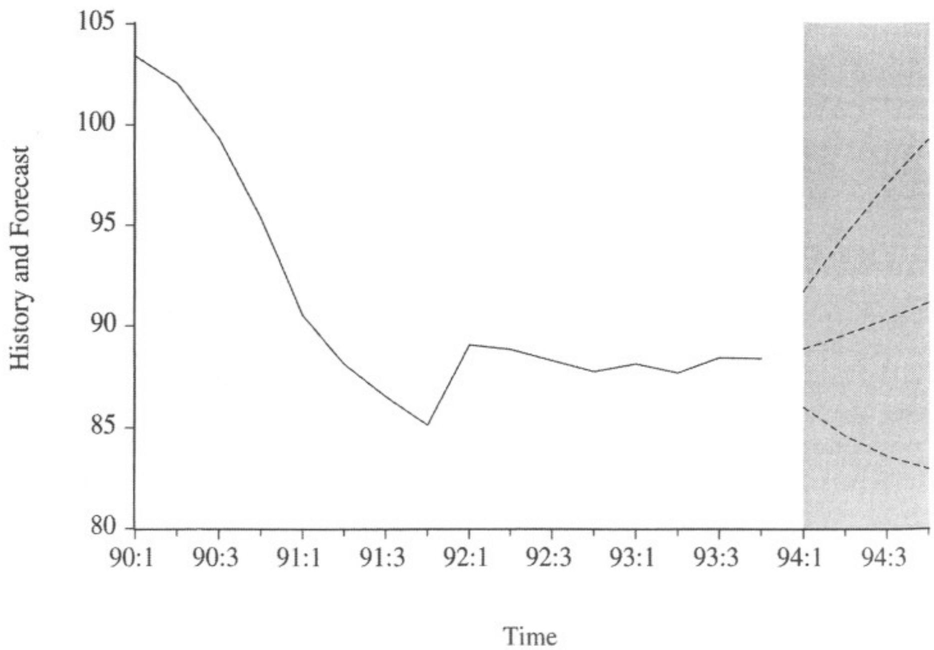

FIGURE 8.5 Employment History and Long Horizon Forecast: AR(2) Model

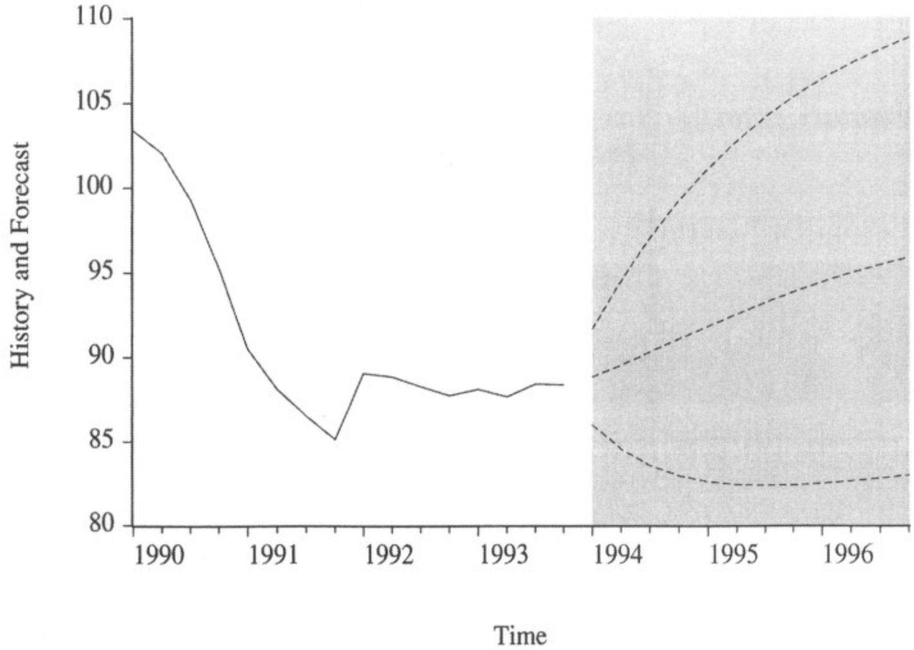

FIGURE 8.6 Employment History and Very Long Horizon Forecast: AR(2) Model

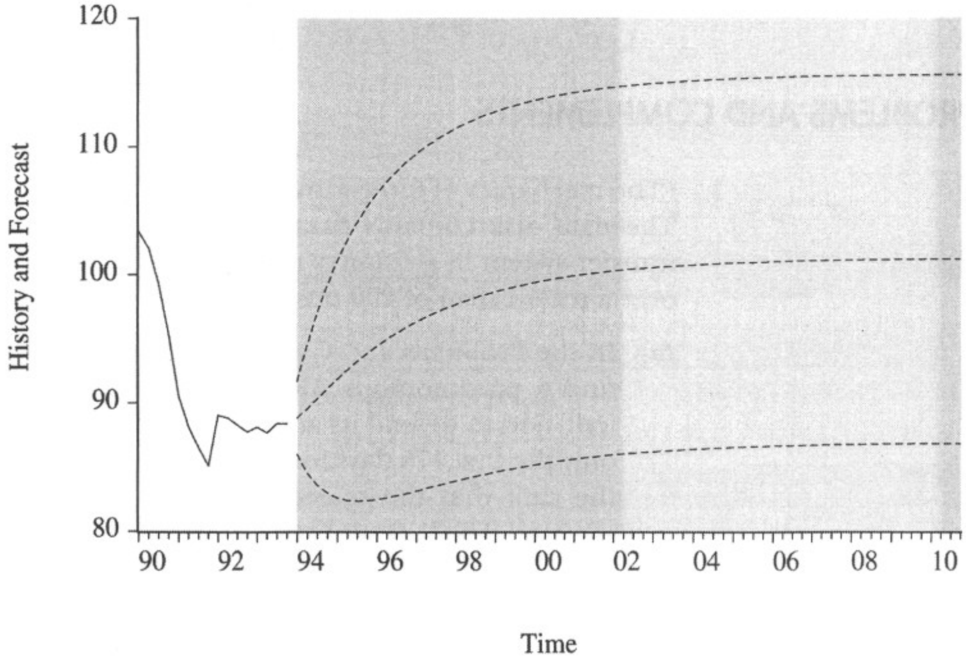

FIGURE 8.7 Employment History, Forecast, and Realization: AR(2) Model

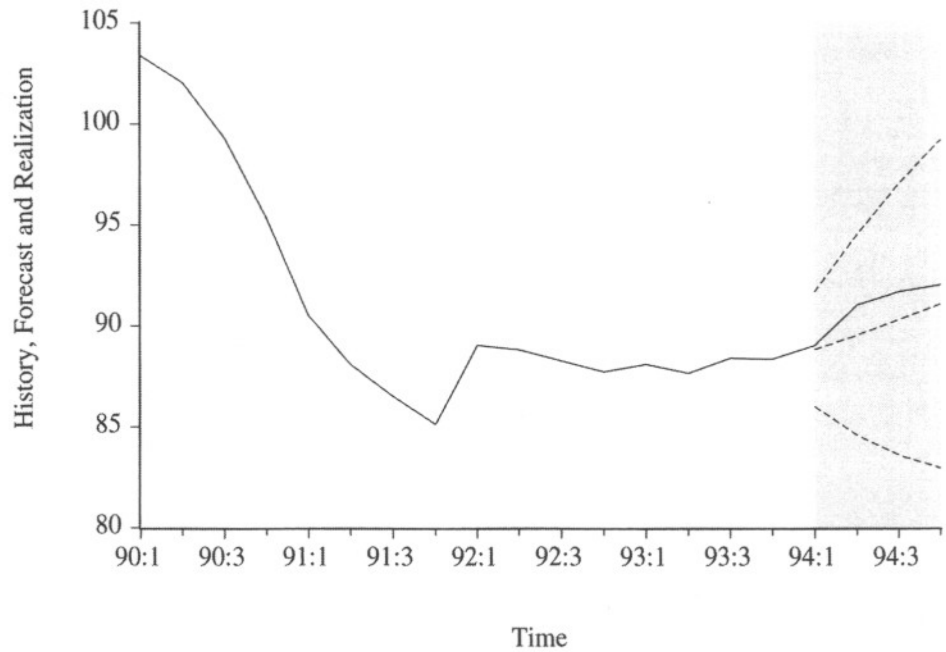

In Figure 8.7 we show the employment history, 4-quarter-ahead AR(2) extrapolation forecast, and the realization. The AR(2) forecast appears quite accurate; the mean squared forecast error is 1.3, drastically smaller than that of the MA(4) forecast.

PROBLEMS AND COMPLEMENTS

1. (The mechanics of forecasting with ARMA models: BankWire continued) The data disk contains data for daily transfers over BankWire, a wire transfer system in a country responsible for much of the world's finance, over a recent span of 200 business days.

   (a) In the Problems and Complements of chapter 7, you were asked to find a parsimonious ARMA(p,q) model that fits the transfer data well, and to defend its adequacy. Repeat the exercise, this time using only the first 175 days for model selection and fitting. Is it necessarily the case that the selected ARMA model will remain the same as when all 200 days are used? Does yours?

   (b) Use your estimated model to produce point and interval forecasts for days 176 through 200. Plot them and discuss the forecast pattern.

   (c) Compare your forecasts to the actual realizations. Do the forecasts perform well? Why or why not?