Time Series Econometrics for the 21st Century

Bruce E. Hansen*
University of Wisconsin†

December 2016

Abstract

What topics should be taught to undergraduate students in econometric time series?

*Prepared for the 2017 AEA session on "Updating the Undergraduate Econometrics Curriculum". Research supported by the National Science Foundation.

†Department of Economics, 1180 Observatory Drive, University of Wisconsin, Madison, WI 53706.

1 Introduction

The field of econometrics largely started with time series analysis, since many early data sets were time-series macro-economic data. As the field developed, more cross-sectional and longitudinal data sets were collected, which today dominate the majority of academic empirical research. However, in non-academic (private sector, central bank, and governmental) applications, time-series data is still important. Consequently our undergraduate courses must still teach time series econometric tools if we are to best serve our students. In this paper I review the key models which I belive are most critical for our students.

The core econometrics textbooks of Wooldridge (2013) and Stock and Watson (2015) both provide excellent treatments of econometric time-series. My preference is the treatment in Stock and Watson, as it runs closer to the way I teach the material.

Two excellent texts on time-series econometrics and forecasting are Diebold (2015) and Gonzáles-Rivera (2012), which are appropriate for dedicated courses on economic forecasting. An excellent graduate-level forecasting textbook is Elliott and Timmermann (2016). The classic textbook for graduate-level time-series remains Hamilton (1994).

The data and Stata code used for the empirical results reported in this paper are posted on the author's website http://www.ssc.wisc.edu/~bhansen/.

2 Overview

Most undergraduate economic majors do not persue Ph.D.'s in economics. Many go on to professional schools, some to government jobs, others to the private sector. In many of these positions it is quite common for our graduates to be exposed to economic data and analysis, including formal econometric (e.g. regression) analysis. Many of these applications are time-series in nature. What tools can we give these students to help them succeed?

I believe the core models which undergraduate students should learn are the Autoregressive (AR) model

$$y_t = \alpha + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + e_t, \tag{1}$$

the regression model

$$y_t = \alpha + \delta_0 x_t + e_t \tag{2}$$

the Distributed Lag (DL) model

$$y_t = \alpha + \delta_0 x_t + \delta_1 x_{t-1} + \cdots + \delta_q x_{t-q} + e_t \tag{3}$$

and the Autoregressive-Distributed Lag (ADL) model

$$y_t = \alpha + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + \delta_0 x_t + \delta_1 x_{t-1} + \cdots + \delta_q x_{t-q} + e_t. \tag{4}$$

In these equations, p is the number of lags of the dependent variable y_t, q is the number of lags of the explanatory variable x_t, and e_t is a mean-zero shock. In the ADL model the contemporaneous regressor x_t is often omitted in contexts such as prediction.

With these core models, most of the concepts of single-equation time-series econometrics can be explained and taught. All are special cases of the ADL model, but it is useful to study them separately.

These models allow us to teach the following core insights:

1. The coefficients can be estimated by OLS just like a conventional regression.

2. The appropriate standard error depends on whether or not the dynamics have been explicitly modeled. In AR and ADL models the "robust" standard errors (e.g. the "r" option in Stata) are appropriate so long as the number of lags p is sufficiently large so that the errors are serially uncorrelated. In the regression and DL models the equation error e_t will be serially correlated so we should use HAC (heteroskedasticity-and-autocorrelation) standard errors, e.g. the "Newey" command in Stata.

3. The AR model is useful to understand the serial correlation properties of the series y_t.

4. The coefficients in the DL and ADL can be interpreted as multipliers. The coefficient δ_0 is the impact multiplier. The ratio $(\delta_1 + \cdots + \delta_1)/(1 - \phi_1 - \cdots - \phi_q)$ is the long-run multiplier.

5. These multipliers have structural interpretations when the explanatory variable x_t is strictly exogeneous. This is a rather special situation.

6. When the structural interpretation is taken, the ADL model can be used for counter-factual and policy analysis.

7. The number of lags (p and q) in AR and ADL models may be selected by comparing models using the Akaike Information Criterion (AIC). Test statistics (t and F statistics) should not be used for model selection.

8. In the regression and DL models we should be greatly concerned about the possibility of a **spurious regression** – the setting where a regression with two unrelated time-series has misleadingly large conventional t-statistics and R^2 values. Understanding the potential for spurious regression may be one of the greatest practical gifts we can teach our students.

9. The parameters of time series models are likely to have changed over time. This requires care and attention.

10. The ADL model without the contemporaneous regressor x_t is useful for prediction. The concept of Granger causality is taught within this model, namely that x_t does not Granger-cause y_t if $\delta_1 = \cdots = \delta_q = 0$.

11. The ADL model without the contemporaneous regressor x_t may be used to produce one-step-ahead point forecasts for y_{n+1}.

12. Point forecasts should be combined with interval forecasts, as the latter convey the degree of uncertainty about future outcomes. Teaching students to appreciate forecast uncertainty and interval forecasts is the most important lesson of prediction.

13. To produce multi-step (h step) forecasts using a ADL model, we can use the multi-step version

$$y_{t+h} = \alpha + \phi_1 y_t + \cdots + \phi_p y_{t-p+1} + \delta_1 x_t + \cdots + \delta_q x_{t-q+1} + e_t.$$

    which can also be estimated by OLS.

14. Multi-step point forecasts should also be accompanied by interval forecasts, and together can be combined into fan charts.

Other time-series issues which can be usefully discussed in an undergraduate course include

1. Trends (deterministic and stochastic)
2. Unit Roots
3. Seasonality
4. Cointegration
5. ARCH
6. Vector autoregressions
7. Out-of-sample and split-sample analysis

Some traditional topics which can be safely omitted include:

1. Durbin Watson statistic
2. Statistical properties under classical assumptions
3. GLS estimation
4. Cochrane-Orcutt estimation

3 Standard Errors and t-statistics

A critical issue in time-series (and econometrics in general) is which standard error to use. There are three popular standard error formulae for applied econometric time-series: classical, heteroskedasticity-robust, and heteroskedasticity and autocorrelation-robust (HAC).

So-called classical (or homoskedastic) standard errors may be useful pedagogically due to their simple form, but are not used in current applied econometric practice. Consequently, courses should quickly move beyond classical standard errors.

Robust standard errors (the "r" option in Stata) are the most commonly used in current applied practice. They are appropriate for time-series models which are dynamically well-modeled, including autoregressive and ADL models.

HAC standard errors (Newey-West, the "Newey" command in Stata) are appropriate for simple (non-dynamic) regressions and DL models. They should be used whenever the serial correlation in the error has not been modeled.

The appropriate way to report regression estimates is coefficient estimates plus standard errors, as the latter can be used to construct confidence intervals or test statistics. Inappropriate reporting methods include coefficient estimates plus t-ratios, as this emphasizes testing (which is not typically the goal of estimation), and coefficient estimates plus asterisks, as this emphasizes statistical significance rather than magnitudes. It is always better for students to focus on parameter estimates, their magnitudes, and interpretations, rather than simple statements about significance.

To illustrate some of these ideas, consider a simple regression of retail gasoline prices on crude oil prices. Let gas_t denote the weekly percentage change in U.S. retail gasoline prices and let oil_t denote the weekly percentage change in the Brendt European spot price for crude oil, for 1991-2016. The estimates are

$$\begin{array}{rccc} gas_t = & 0.029 & + & 0.269 \quad oil_t + \widehat{e}_t \\ & (0.046) & & (0.011) \\ & (0.046) & & (0.015) \\ & (0.073) & & (0.021) \end{array}$$

Here we report the coefficient estimates plus three standard errors. The first set of standard errors are classical (homoskedastic), the second are heteroskedasticity-robust, and the third are Newey-West estimates with 12 lags. We can see that the choice of standard error formula matters greatly, with the Newey-West roughly twice the magnitude of the classical. Since this is a static regression, the Newey-West are the appropriate choice.

Furthermore, it is appropriate to report standard errors rather than t-statistics as the former convey the most important information about the regression. In the above regression, we learn that a 1% change in crude oil prices leads to a contemporeneous (within one week) 0.27% change in retail gasoline prices. This is an immediate but partial pass-through. The Newey-West standard

error shows that the degree of pass-through is relatively precisely estimated, with a 95% confidence interval is 23% to 31%. In contrast, the t-statistic for the coefficient is 13, which simply tells us that the true coefficient is non-zero. This is somewhat informative, but much less so than the discussion about magnitudes.

4 Autoregressive Models

A good illustration of an autoregressive model is GDP growth. Let GDP_t denote quarterly real U.S. GDP percentage growth at annual rates, a series which is available for 1947Q2 throgh 2016Q3. An AR(3) is

$$\begin{array}{rccccccccc} GDP_t = & 1.93 & + & 0.34 & GDP_{t-1} + & 0.13 & GDP_{t-2} - & 0.09 & GDP_{t-3} + \widehat{e}_t \\ & (0.32) & & (0.06) & & (0.06) & & (0.06) & \end{array}$$

Here we report heteroskedasticity-robust standard errors since this is a dynamic model.

The estimates show that gdp growth is positively autocorrelated, but mildly. This means that higher-than-average growth is followed in subsequent quarters by higher-than-average growth, but with quick mean reversion.

Another interesting example is stock price returns. Let $return_t$ denote the weekly percentage change in the S&P 500 index for 1950-2016. (For simplicity we ignore dividends. Also, while daily observations are available they have extra complications so it is easier to focus on weekly observations.) The estimates are

$$\begin{array}{rccccccc} return_t = & 0.16 & - & 0.032 & return_{t-1} + & 0.037 & return_{t-2} + \widehat{e}_t \\ & (0.04) & & (0.029) & & (0.025) & & \end{array}$$

A simple form of the efficient market hypothesis suggests that stock returns are unpredictable and thus the AR coefficients should be zero. Thus the F-statistic for the AR coefficients is a simple test of efficient markets. In this example the p-value for the F statistic is 0.12 so we fail to reject the efficient market hypothesis.

To repeat our message about the importance of using robust standard errors, if instead we had used the "old-fashioned" (homoskedastic) formula the standard errors on the AR coefficients would instead be 0.17, the second lag would have a t-statistic of 2.2 and the F statistic for the two AR coefficients would have a p-value of 0.01, which would incorrectly suggest rejection of the efficient market hypothesis. Indeed, using the correct standard error formula makes a huge difference and alters inference.

5 Distributed Lag Models

Distributed lag models are useful when we want to estimate the impact of one variable upon another. As an example, consider the effect of crude oil prices upon retail gasoline prices, using the data from Section 3. A distributed lag model with a contemporeneous effect and six lags takes the form

$$\begin{aligned} gas_t = {} & \underset{(0.057)}{-0.009} + \underset{(0.016)}{0.243}\, oil_t + \underset{(0.012)}{0.112}\, oil_{t-1} + \underset{(0.011)}{0.063}\, oil_{t-2} \\ & + \underset{(0.013)}{0.064}\, oil_{t-3} + \underset{(0.010)}{0.030}\, oil_{t-4} + \underset{(0.011)}{0.032}\, oil_{t-5} + \underset{(0.012)}{0.018}\, oil_{t-6} + \widehat{e}_t \end{aligned}$$

Here, the standard errors are computed using the Newey-West formula with 12 lags.

Under the assumption of strict exogeneity, the coefficients of a DL model are the effects of the regressor on the dependent variable. In this case, we see that a 1% change in crude oil prices leads to an contemporeous (within one week) change in retail gasoline prices of 0.24%, followed by further increases over the following six weeks.

The following equivalent regression can be used to estimate the cumulative multipliers

$$\begin{aligned} gas_t = {} & \underset{(0.057)}{-0.009} + \underset{(0.016)}{0.243}\, \Delta oil_t + \underset{(0.024)}{0.355}\, \Delta oil_{t-1} + \underset{(0.028)}{0.418}\, \Delta oil_{t-2} \\ & + \underset{(0.037)}{0.482}\, \Delta oil_{t-3} + \underset{(0.040)}{0.512}\, \Delta oil_{t-4} + \underset{(0.045)}{0.544}\, \Delta oil_{t-5} + \underset{(0.047)}{0.562}\, oil_{t-6} + \widehat{e}_t \end{aligned}$$

These results show that a 1% change in crude oil prices lead to a 0.56% cumulative change in retail gasoline prices after six weeks, with most of the change incorporated within the first four weeks. These estimates show how quickly changes in crude oil prices translate into retail prices.

6 Autoregressive Distributed Lag Models

By combining the autoregressive and distributed lag models we obtain a dynamically sound model useful for both policy analysis and forecasting. As an example consider a short-run quarterly Phillips curve which models the change in U.S. CPI inflation as a function of the U.S. unemployment

rate

$$\begin{aligned}\Delta Inf_t = & \underset{(0.42)}{0.44} - \underset{(0.11)}{0.34}\,\Delta Inf_{t-1} - \underset{(0.09)}{0.39}\,\Delta Inf_{t-2} - \underset{(0.11)}{0.02}\,\Delta Inf_{t-3} - \underset{(0.07)}{0.17}\,\Delta Inf_{t-4} \\ & - \underset{(0.56)}{1.53}\,Unemp_{t-1} + \underset{(1.06)}{1.58}\,Unemp_{t-2} + \underset{(1.03)}{0.11}\,Unemp_{t-3} - \underset{(0.47)}{0.23}\,Unemp_{t-4} + \widehat{e}_t\end{aligned}$$

We estimate this equation without the contemporeneous umployment rate so that it can be used as a forecasting equation.

The coefficients on lagged changes in inflation are large and negative, indicating that quarterly changes in inflation tend to be followed by a reversal over the next 6 months. The impact effect of the unemployment rate is negative, indicating that when the unemployment rate is higher than normal, then the inflation rate tends to fall. Conversely, when the unemployment rate is lower than normal, then the inflation rate tends to rise. This is the classic Phillips curve relationship. The sum of all four coefficients, however, is close to zero, indicating that the long-run impact of the unemployment rate on inflation rate changes is zero. This means that when the unemployment rate has been higher than normal and roughly constant for multiple periods (or lower than normal and roughly constant) then there is no effect on the inflation rate.

To assess if the overall effect is statistically significant, we can perform a joint statistical test on all four lags of the unemployment rate. This F statistic has a p-value of 0.03, so is marginally significant. This is known as a Granger Causality test, and the evidence here suggests that the U.S. unemployment rate "Granger causes" U.S. inflation rates. This means that the unemployment rate helps to forecast inflation, not that it is causal in a strict sense.

7 Model Selection

Economic theory does not provide guidance regarding the number of lags in an AR, DL or ADL model. Rather it is a matter of statistical fit. The more lags are included, the smaller the bias yet the larger the estimation variance. Lag selection is inherently a bias-variance trade-off.

Consequently, statistical tests (Durbin-Waton, t and F tests) are not appropriate for lag selection. There is no economic hypothesis to test. There is no null hypothesis. Testing is the incorrect lens through which to view the model.

Instead, information criterion are appropriate. For undergraduate courses the AIC is a good criterion as it is simple and readily available. AR and ADL models can be easily compared via the AIC. This is a simple and effective method to select the number of lags in a dynamic model.

As one example, take the GDP autoregressive model presented in Section 4. Comparing AR(0) through AR(6) models, the AIC criterion selects the AR(3) model as presented.

As a second example, take the Phillips curve example from Section 6. Comparing models with

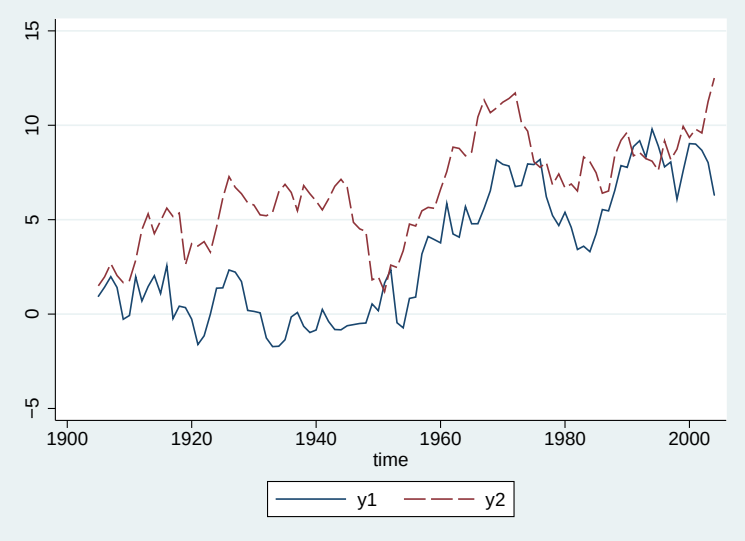

Figure 1: Plot of y_1 and y_2

1 through 6 lags of the change in the inflation rate and 1 through 4 lags of the unemployment rate, the AIC selects the model with 4 lags of the change in the inflation rate and 2 lags of the unemployment rate. We will use this selected model for a forecast application in Section 10.

8 Spurious Regression

One of the dangers of estimating simple regressions or distributed lag models is that conventional standard errors can be off by several orders of magnitude. This danger is routinely ignored in non-academic (e.g. journalistic) writing, and can also be seen in some academic (e.g. macroeconomic) articles. This is a serious practical issue facing economists working in industry.

To illustrate, Figure 1 displays a plot of two annual time-series, labeled y_1 and y_2, for the period 1906-2015. The two appear to track each other fairly well. Their sample correlation is 0.73. Estimation of a linear regression of y_{1t} on y_{2t} (with conventional "robust" standard errors) is reported below. The coefficient for y_{2t} is highly significant (the t-statistic is about 13) and $R^2 = 0.54$. Clearly the relationship seems strong.

$$\begin{array}{rccc} y_{1t} = & -2.95 & + & 0.95 \quad y_{2t} + \widehat{e}_t \\ & (0.52) & & (0.07) \end{array}$$

However, the truth is that relationship is completely spurious. The two time-series were generated using a random number generator, and are statistically independent. The two series are completely unrelated, yet have a high sample correlation, R^2, and t-statistics.

The reason behind the deception is the strong serial correlation in the variables. In this example they were generated as random walks. The seemingly large sample correlation, R^2, and t-statistic

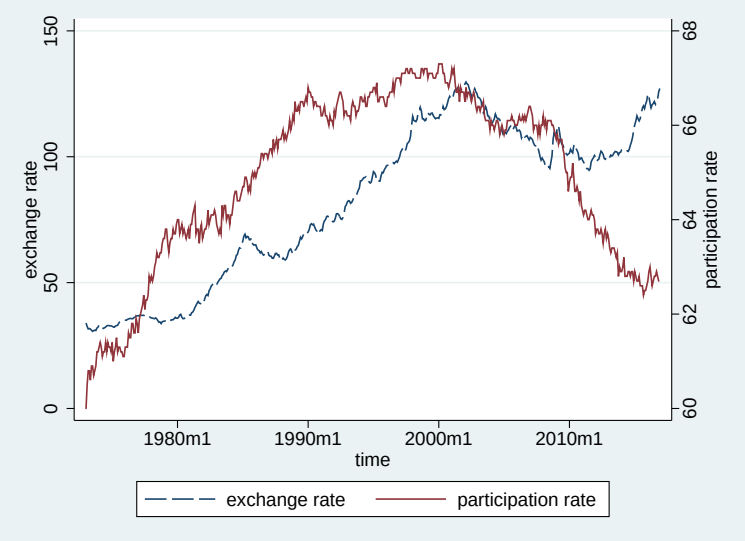

Figure 2: Plot of U.S. exchange rate and labor force participation rate

are all artifacts of unmodeled serial correlation.

The secret to detecting spurious regressions of this form is to first examine the time-series plot. If strong serial correlation is present in the series, then static regressions and distributed-lag models need (at a minimum) HAC-type standard errors. If the serial correlation is sufficiently strong (as is the case for a random walk process such as in this example) then HAC standard errors will not be a sufficient correction either. Instead, a better check is to estimate an ADL model. Typically, it is sufficient to include a lagged dependent variable. We report such estimates here

$$\begin{array}{rccccc} y_{1t} = & -0.22 & + & 0.91 & y_{1t-1} + & 0.09 \quad y_{2t} + \widehat{e}_t \\ & (0.35) & & (0.04) & & (0.06) \end{array}$$

The point estimate of the long-run multiplier is similar to the static regression, but the effect is no longer statistically significant. As a general rule, this is a simple method to break a spurious regression.

As an example using actual data, in Figure 2 we plot monthly observations on the U.S. trade-weighted exchange rate, along with the U.S. labor force participation rate. The two seem to have a common relationship. The sample correlation is 0.59. The R^2 from a linear regression is 0.35 and the labor force participation rate has a highly significant t-statistic of 18.

$$\begin{array}{rccc} ExchangeRate_t = & -578 & + & 10.2 \quad LaborParticipation_t + \widehat{e}_t \\ & (37) & & (0.56) \end{array}$$

However, the relationship appears to be spurious once we account for the serial correlation. We show this by re-estimating the relationship include four lags of the dependent variable. (Four lags

were included to ensure that the serial correlation in the exchange rate was fully modeled.) In this ADL regression, the labor force participation rate has a t-statistic of just 0.4, so is statistically insignificant. This is a more reasonable finding, as an economist would naturally be puzzled by a claim that the labor participation rate effects the exchange rate!

$$\begin{aligned} ExchangeRate_t = & \underset{(1.94)}{-0.58} + \underset{(0.05)}{1.43}\ ExchangeRate_{t-1} - \underset{(0.08)}{0.59}\ ExchangeRate_{t-2} \\ & + \underset{(0.00)}{0.18}\ ExchangeRate_{t-3} - \underset{(0.05)}{0.05}\ ExchangeRate_{t-4} \\ & - \underset{(0.032)}{0.013}\ LaborParticipation_t + \widehat{e}_t \end{aligned}$$

9 Structural Change

Time-series relationships often change. In practical applications this is important to monitor. As a simple example, take U.S. GDP growth. There is considerable concern that growth rates may have slowed down in recent decades. To illustrate, we compute the average growth rate, standard deviation, and OLS AR(1) coefficient estimate over four historical periods.

U.S. GDP Growth Rates

	Mean	Standard Deviation	AR(1) Coefficient
1947-1956	4.0	5.3	0.44
1957-1976	3.6	4.2	0.30
1977-1996	3.2	3.5	0.31
1997-2016	2.3	2.5	0.41

The table shows that average growth rates have decreased from an average near 4% from 1947 through 1976, to 2.3% over the past decade. Simultaneously, the standard deviation has decreased by half. This means that average growth has slowed and has become less volatile. The AR(1) coefficient estimate are all quite similar (between 0.3 and 0.4) and show no pattern, suggesting that there is no meaningful change in the serial correlation patterns.

These calculations, while simple, have important implications for economic analysis and forecasting. If the mean growth rate has changed, then using the full sample to compute forecasts will lead to systematic errors.

10 Forecasting

A common application for applied time series is forecasting. Given the ADL model (3) with the contemporeneous effect omitted, an estimated direct h-step forecasting equation is

$$y_{t+h} = \widehat{\alpha} + \widehat{\phi}_1 y_t + \cdots + \widehat{\phi}_p y_{t-p+1} + \widehat{\delta}_1 x_t + \cdots + \widehat{\delta}_q x_{t-q+1} + \widehat{e}_t. \tag{5}$$

The key is that all variables on the right-hand-side are dated at time t or before. Estimation is OLS using the sample from $t = 1, ..., n$. The point forecast for y_{n+h} is

$$\widehat{y}_{n+h} = \widehat{\alpha} + \widehat{\phi}_1 y_n + \cdots + \widehat{\phi}_p y_{n-p+1} + \widehat{\delta}_1 x_n + \cdots + \widehat{\delta}_q x_{n-q+1}.$$

A simple $1 - \alpha$ forecast interval is based on the normal approximation

$$\widehat{y}_{n+h} \pm \widehat{s}_{n+h} z_{1-\alpha/2}$$

where z_α is the $1 - \alpha/2$ normal quantile (e.g. $z_\alpha = 1.645$ for an 80% interval) and $\widehat{s}_{n+h}$ is the standard error of the forecast. The latter is generated in Stata by the command `predict` *`name`*`, stdf`. Alternatively, a reasonable approximation is the standard deviation of the regression residual, $\widehat{s}_{n+h} \approx \left(n^{-1} \sum_{t=1}^{n} \widehat{e}_t^2\right)^{1/2}$.

For direct forecasts, a separate regression is estimated for each forecast horizon h. Thus for multiple forecasts we need to estimate multiple OLS regressions.

To illustrate, let's take quarterly U.S. inflation, and use the ADL Phillips curve model selected by the AIC (four autoregressive lags and two lags of the unemployment rate) to make forecasts for 2017-2018. This requires estimation of eight regressions of the form (5) for $h = 1, ..., 8$. The point forecasts and 80% forecast intervals are plotted in Figure 3. These are forecasts for quarterly inflation at an annual rate. You can see that the Phillips curve equation predicts that the inflation rate will continue at its current 2% rate. The 80% forecast intervals show that there is considerable uncertainty in the inflation realizations.

References

[1] Diebold, Francis X (2015): *Forecasting in Economics, Business, Finance and Beyond*, http://www.ssc.upenn.edu/~fdiebold/

[2] Elliott, Graham and Allan Timmermann (2016) *Economic Forecasting*, Princeton University Press.

[3] Gonzáles-Rivera, Gloria (2012): *Forecasting for Economics and Business*, Routledge.

[4] Hamilton, James D. (1994): *Time Series Analysis*, Princeton University Press.

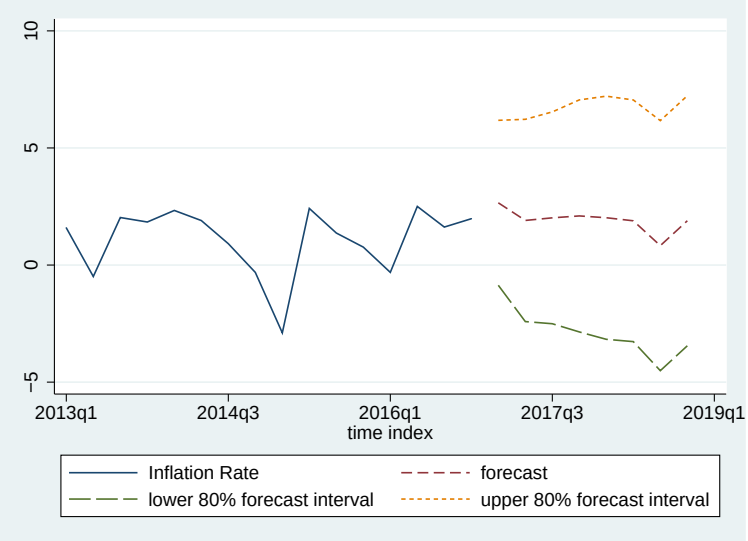

Figure 3: U.S. Inflation Rate Forecasts for 2017-2018

[5] Stock, James H. and Mark W. Watson (2015): *Introduction to Econometrics*, 5^{th} Edition, Pearson.

[6] Wooldridge, Jeffrey M. (2013): *Introductory Econometrics: A Modern Approach*, 5^{th} Edition. South-Western.