Intro to Time Series

- A time series is a single occurrence of a random event.
- The sequence of observations, $\{y_t\}_{t=-\infty}^{t=+\infty}$, is a time-series process.
- There is no counterpart to repeated sampling for a time series.
- We observe realizations of this process in a time window, $t = 1, \ldots, T$.
- The frequency of observations are important, but the length of the window is arguably more important.
- Asymptotics involves considering an increasingly longer window.

Data from Greene (7th ed.), Table F5.2. Macroeconomics data set, quarterly from 1950 to 2000.

```
data <-
read.csv("http://home.cc.umanitoba.ca/~godwinrt/4042/material/gdptime.csv")
attach(data)
```

Variables in the data set include real GDP, M1 money supply, and CPI.

```
plot(time, realgdp, cex= .5)
```

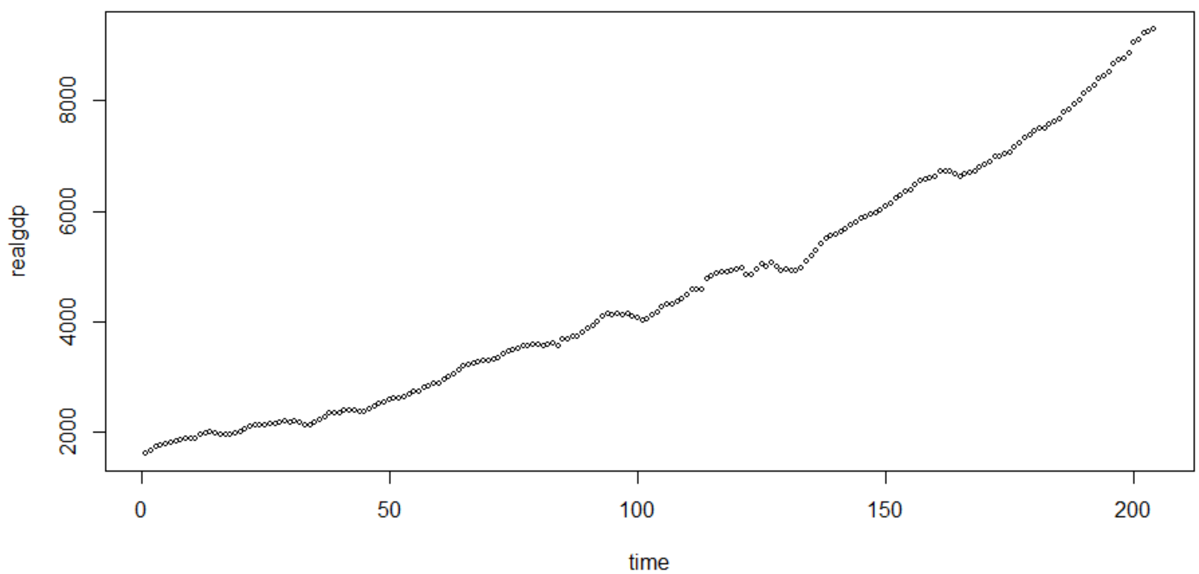

Linear over time?

```
plot(time, log(realgdp), cex = .5)
```

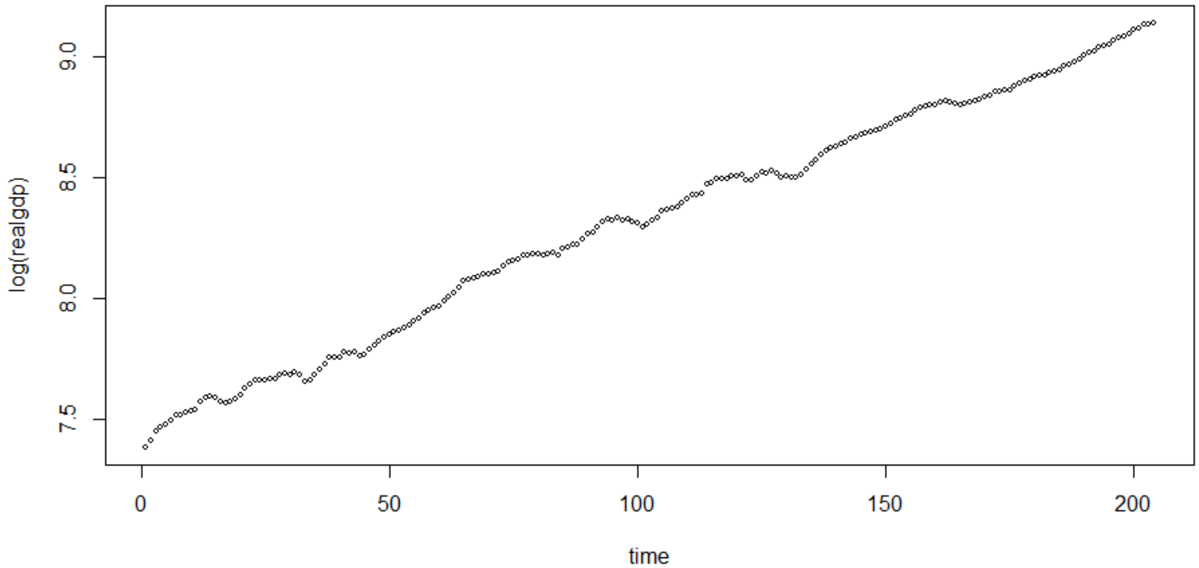

Take logs of the variables:

```
lgdp <- log(realgdp)
lm1 <- log(M1)
lcpi <- log(cpi_u)
```

Regress *log(GDP)* on a *constant* and *time trend*

```
ols1 <- lm(lgdp ~ time)
summary(ols1)

Coefficients:
             Estimate Std. Error t value Pr(>|t|)
(Intercept) 7.468e+00  5.486e-03  1361.3   <2e-16 ***
time        8.236e-03  4.641e-05   177.5   <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.03903 on 202 degrees of freedom
Multiple R-squared:  0.9936,  Adjusted R-squared:  0.9936
F-statistic: 3.149e+04 on 1 and 202 DF,  p-value: < 2.2e-16
```

```
abline(ols1)
```

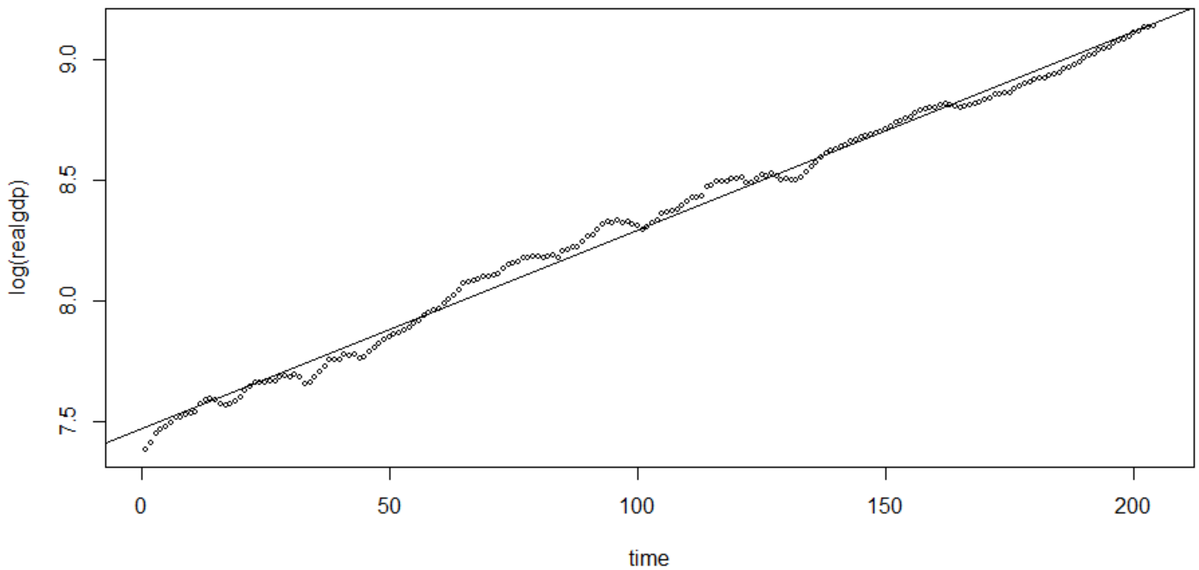

Regress *log(GDP)* on a *constant* and its *lagged value*:

```
lgdpt1 <- lgdp[2:204]
lgdpt <- lgdp[1:203]
summary(lm(lgdpt1 ~ lgdpt))

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) 0.026098   0.011993   2.176   0.0307 *
lgdpt       0.997899   0.001441 692.485   <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.009943 on 201 degrees of freedom
Multiple R-squared:  0.9996,  Adjusted R-squared:  0.9996
F-statistic: 4.795e+05 on 1 and 201 DF,  p-value: < 2.2e-16
```

What does this tell you?

A time-series model typically describes a variable, y_t, in terms of:

- contemporaneous factors, $\boldsymbol{x}_t$
- lagged factors, $\boldsymbol{x}_{t-1}, \boldsymbol{x}_{t-2}, \ldots$
- its own past, $y_{t-1}, y_{t-2}, \ldots$
- disturbances (innovations), ε_t.

For example:

$$y_t = \beta_1 + \beta_2 x_t + \beta_3 y_{t-1} + \varepsilon_t$$

Let's try regressing the (logs) of U.S. money stock (M1) on real GDP and CPI:

```
ols2 <- lm(lm1 ~ lgdp + lcpi)
summary(ols2)
Coefficients:
            Estimate Std. Error t value Pr(>|t|)    
(Intercept) -1.63306    0.22857  -7.145 1.62e-11 ***
lgdp         0.28705    0.04738   6.058 6.68e-09 ***
lcpi         0.97181    0.03377  28.775  < 2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.08288 on 201 degrees of freedom
Multiple R-squared:  0.9895,  Adjusted R-squared:  0.9894 
F-statistic:  9489 on 2 and 201 DF,  p-value: < 2.2e-16
```

Plot the residuals over time:

```
resids <- ols2$residuals
plot(time, resids)
```

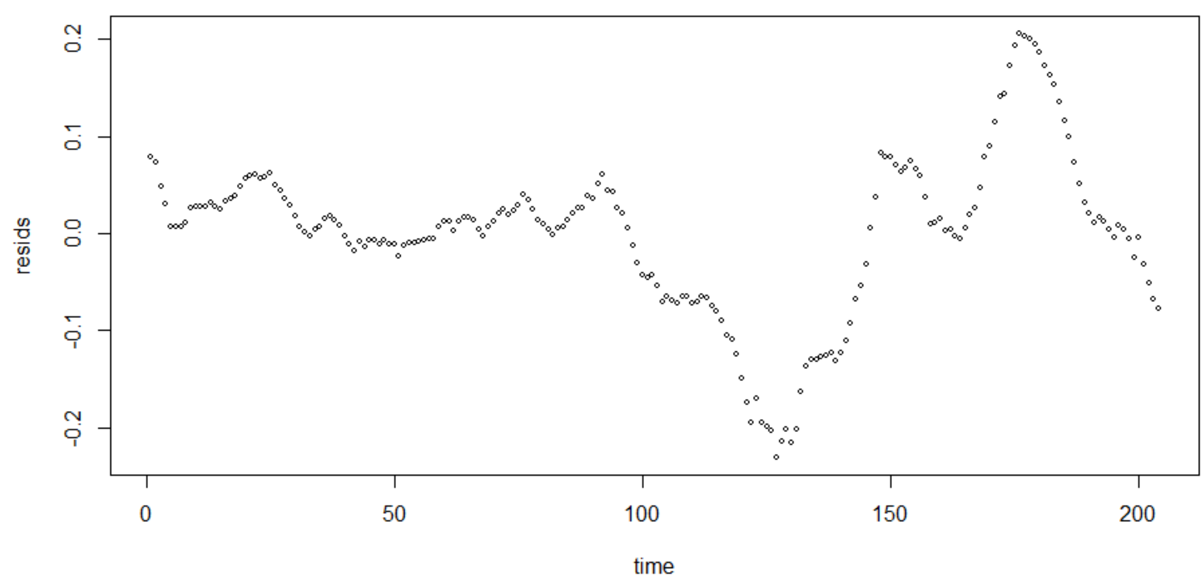

What does this tell you?

Autocorrelated Errors

- Often, current values of the error term are correlated with past values.
- We often say they are "*Serially Correlated* ".
- In this case, the off-diagonal elements of $V(\boldsymbol{\varepsilon})$ will be non-zero.
- The particular values they take will depend on the *form of autocorrelation*.
- That is, they will depend on the *pattern of the correlations* between the elements of the error vector.
- $$V(\boldsymbol{\varepsilon}) = \begin{bmatrix} \sigma^2 & \sigma_{12} & \sigma_{13} \\ \sigma_{12} & \sigma^2 & \sigma_{23} \\ \sigma_{13} & \sigma_{23} & \sigma^2 \end{bmatrix}$$
- If the *errors* themselves are autocorrelated, often this will be reflected in the regression *residuals* also being autocorrelated.
- That is, the residuals will follow some sort of *pattern*, rather than just being random.
- If the errors of our model are autocorrelated, then the OLS estimator of $\boldsymbol{\beta}$ usually will be unbiased and consistent, but it will be inefficient.
- In addition $V(\boldsymbol{b})$ will be computed incorrectly, and the standard errors, *etc*., will be ***inconsistent***. (Same situation as with heteroskedasticity).
- In general, concern lies in formal methods for
  1. Testing for the presence/absence of autocorrelation.
  2. Estimating models when the errors are autocorrelated.
- However, in this introduction, we will not look at these methods. We will consider to ways of modelling autocorrelation: an AR process and an MA process. We will also consider a limiting form of an AR process.

Autoregressive Process

$$\varepsilon_t = \rho\varepsilon_{t-1} + u_t \quad ; \quad u_t \sim i.i.d.N[0\,,\sigma_u^2] \quad ; \quad |\rho| < 1$$

This is an AR(1) model for the error process.

More generally:

$$\varepsilon_t = \rho_1\varepsilon_{t-1} + \rho_2\varepsilon_{t-2} + \cdots + \rho_p\varepsilon_{t-p} + u_t \quad ; \quad u_t \sim i.i.d.N[0\,,\sigma_u^2]$$

This is an AR(p) model for the error process. [*e.g.*, p = 4 with quarterly data.]

Moving Average Process

$$\varepsilon_t = u_t + \phi u_{t-1} \quad ; \quad u_t \sim i.i.d.N[0\,,\sigma_u^2]$$

This is an MA(1) model for the error process.

More generally:

$$\varepsilon_t = u_t + \phi_1\varepsilon_{t-1} + \cdots + \phi_q u_{t-q} \quad ; \quad u_t \sim i.i.d.N[0\,,\sigma_u^2]$$

This is an MA(q) model for the error process.

We can combine both types of process into an ARMA(p , q) model:

$$\varepsilon_t = \rho_1\varepsilon_{t-1} + \rho_2\varepsilon_{t-2} + \cdots \rho_p\varepsilon_{t-p} + u_t + \phi_1 u_{t-1} + \cdots + \phi_q u_{t-q}$$

where: $u_t \sim i.i.d.N[0\,,\sigma_u^2]$.

- Note that in the AR(1) process, we said that $|\rho| < 1$.
- This condition is needed to ensure that the process is "stationary".
- Let's see what this actually means, more generally.

Stationarity

Suppose that the following 3 conditions are satisfied:

1. $E[\varepsilon_t] = 0$; for all t
2. $var.[\varepsilon_t] = \sigma^2$; for all t
3. $cov.[\varepsilon_t , \varepsilon_s] = \gamma_{|t-s|}$; for all t, s; $t \neq s$

Then we say that the time-series sequence, $\{\varepsilon_t\}$ is "Covariance Stationary"; or "Weakly Stationary".

- More generally, this can apply to *any* time-series – not just the error process.
- Unless a time-series is stationary, we can't identify & estimate the parameters of the process that is generating its values.
- Let's see how this notion relates to the AR(1) model, introduced above.
- We have: $\varepsilon_t = \rho\varepsilon_{t-1} + u_t$

$$E[u_t] = 0$$
$$var.[u_t] = E[u_t^2] = \sigma_u^2$$
$$cov.[u_t, u_s] = 0 \quad ; \quad t \neq s$$

- So,

$$\begin{aligned}\varepsilon_t &= \rho[\rho\varepsilon_{t-2} + u_{t-1}] + u_t \\ &= \rho^2\varepsilon_{t-2} + \rho u_{t-1} + u_t \\ &= \rho^2[\rho\varepsilon_{t-3} + u_{t-2}] + \rho u_{t-1} + u_t \\ &= \rho^3\varepsilon_{t-3} + \rho^2 u_{t-2} + \rho u_{t-1} + u_t \\ &etc.\end{aligned}$$

- Continuing in this way, eventually, we get:

$$\varepsilon_t = u_t + \rho u_{t-1} + \rho^2 u_{t-2} + \cdots \qquad \textbf{(1)}$$

[This is an infinite-order MA process.]

The value of ε_t embodies the entire past history of the u_t values.

- From **(1)**, $E(\varepsilon_t) = 0$, and

$$\begin{aligned}var.(\varepsilon_t) &= var.(u_t) + var.(\rho u_{t-1}) + var.(\rho^2\varepsilon_{t-2}) + \cdots \\ &= \sigma_u^2 + \rho^2\sigma_u^2 + \rho^4\sigma_u^2 + \cdots\end{aligned}$$

- Now, under what conditions will this series converge?

  The series will converge to $\sigma_u^2(1-\rho^2)^{-1}$, as long as $|\rho^2| < 1$, and this in turn requires that $|\rho| < 1$.
- This is a necessary condition needed to ensure that the process, $\{\varepsilon_t\}$ is stationary, because if this condition isn't satisfied, then $var.[\varepsilon_t]$ is ***infinite***.
- So, for the AR(1) process, as long as $|\rho| < 1$, then $var.[\varepsilon_t] = \sigma_u^2(1-\rho^2)^{-1}$.
- In addition, stationarity implies that $var.[\varepsilon_t] = var.[\varepsilon_{t-s}]$, for all 's'.
- It can be shown that the full covariance matrix for $\boldsymbol{\varepsilon}$ is:

$$V(\boldsymbol{\varepsilon}) = \sigma_u^2\Omega = \frac{\sigma_u^2}{(1-\rho^2)}\begin{bmatrix} 1 & \rho & \cdots & \rho^{n-1} \\ \rho & 1 & \ddots & \rho^{n-2} \\ \vdots & \ddots & \ddots & \vdots \\ \rho^{n-1} & \rho^{n-2} & \dots & 1 \end{bmatrix}$$

- If we can find a matrix, P, such that $\Omega^{-1} = P'P$, and if the value of ρ were *known*, then we could apply GLS estimation.

Random Walk

Consider a random variable that follows an AR(1) process, but where $\rho = 1$. This variable is said to be *nonstationary*, to be *integrated of order one* I(1), or to follow a *random walk*.

What is the variance for a variable that follows a random walk?

What happens if we regress one random walk on another?

See:

Granger, C. W., & Newbold, P. (1974). Spurious regressions in econometrics. *Journal of econometrics*, *2*(2), 111-120.

```
n <- 100
y <- x <- 0

for(i in 2:n){
  y[i] <- y[i - 1] + rnorm(1)
  x[i] <- x[i - 1] + rnorm(1)
}

plot(y, type = "l", col = "red", ylim = c(min(x,y),max(x,y)))
points(x, type = "l", col = "blue")
summary(lm(y ~ x))
```