Introduction to Time Series Econometrics and Structural Breaks

Ziyodullo Parpiev, PhD

Outline

- **1.** Stochastic processes
- **2.** Stationary processes
- **3.** Purely random processes
- **4.** Nonstationary processes
- **5.** Integrated variables
- **6.** Random walk models
- **7.** Cointegration
- **8.** Deterministic and stochastic trends
- **9.** Unit root tests

What is a time series?

A time series is any series of data that varies over time. For example

- Payroll employment in the U.S.
- Unemployment rate
- 12-month inflation rate
- Daily price of stocks and shares
- Quarterly GDP series
- Annual precipitation (rain and snowfall)

Because of widespread availability of time series databases most empirical studies use time series data.

Some monthly U.S. macro and financial time series

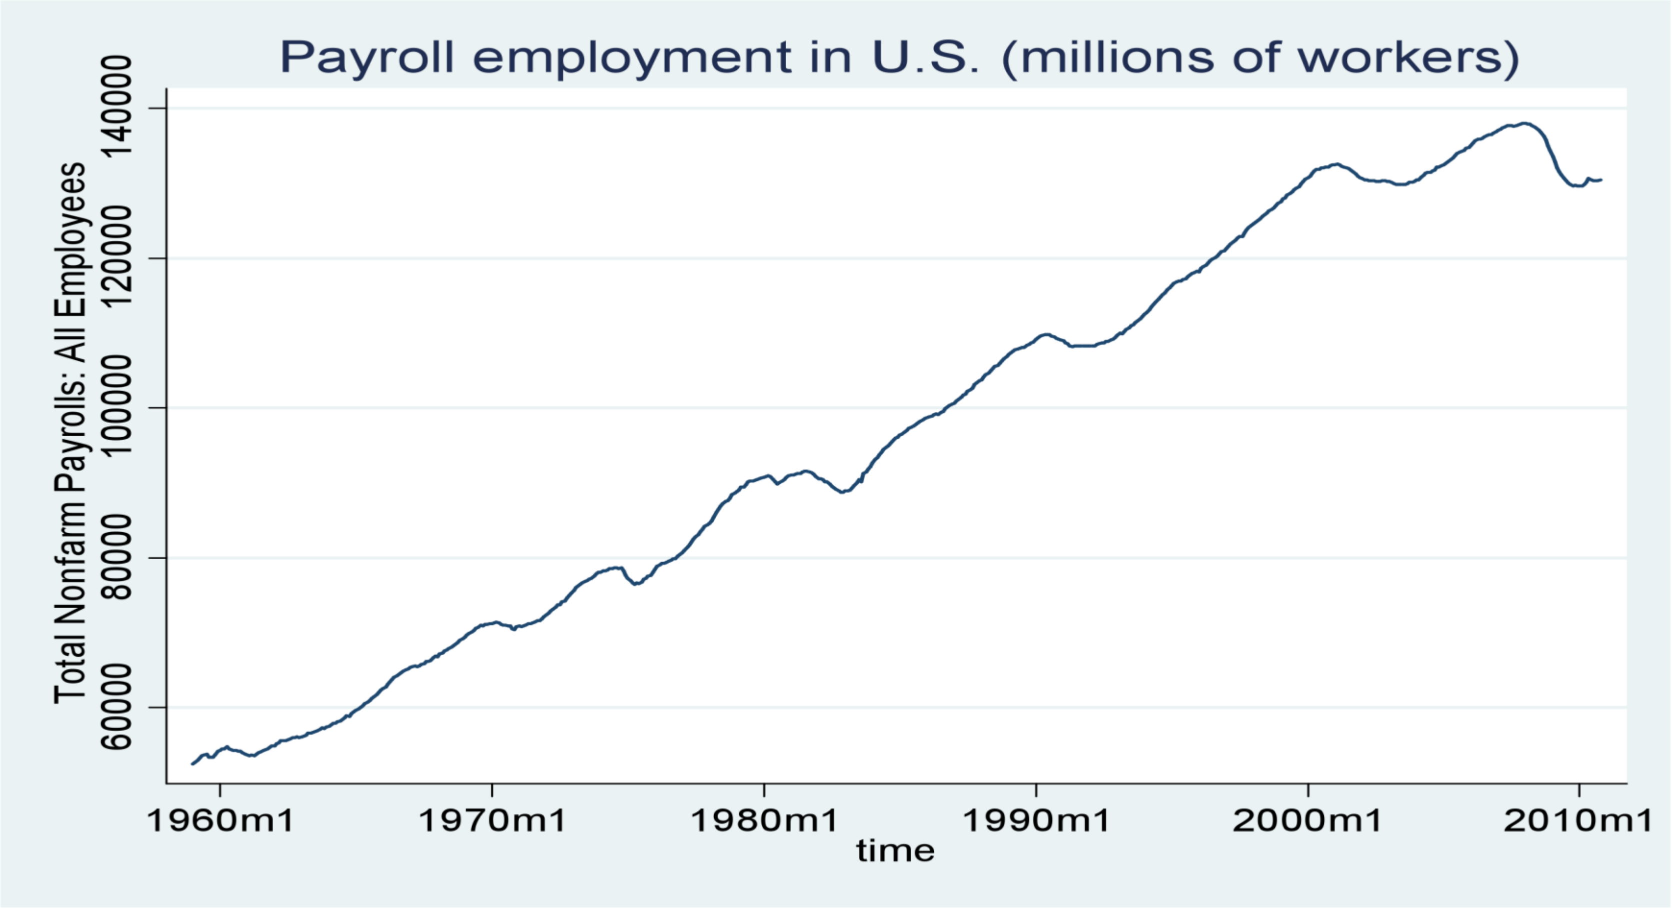

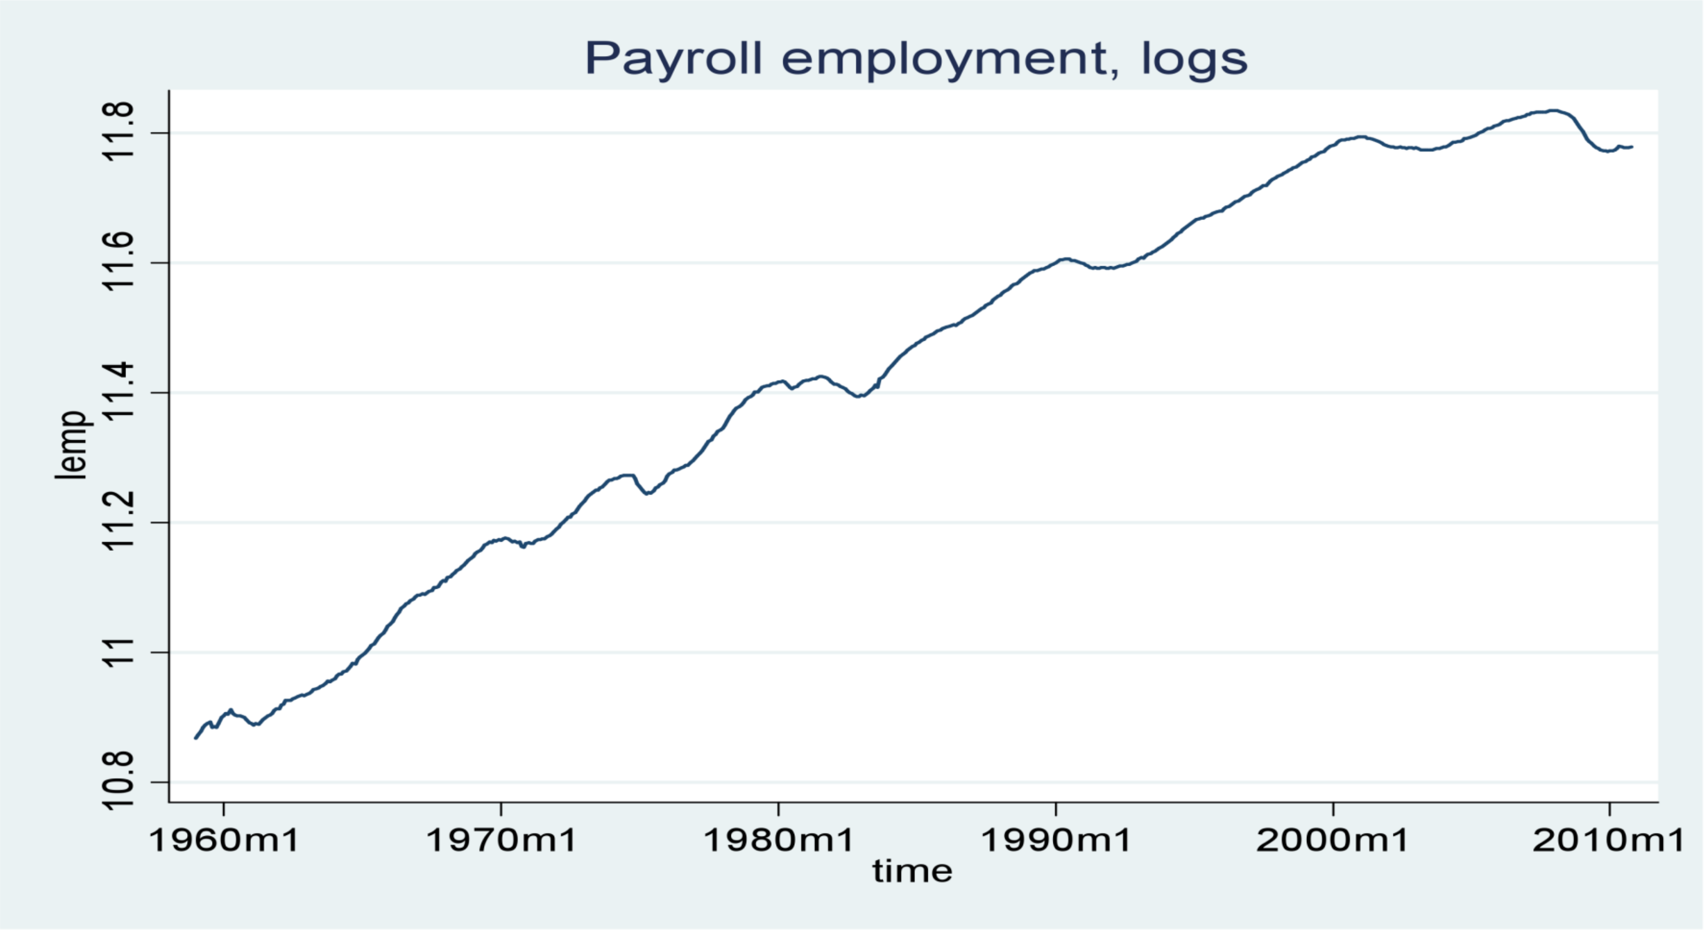
Payroll employment, logs
lemp
10.8
11
11.2
11.4
11.6
11.8
1960m1
1970m1
1980m1
1990m1
2000m1
2010m1
time

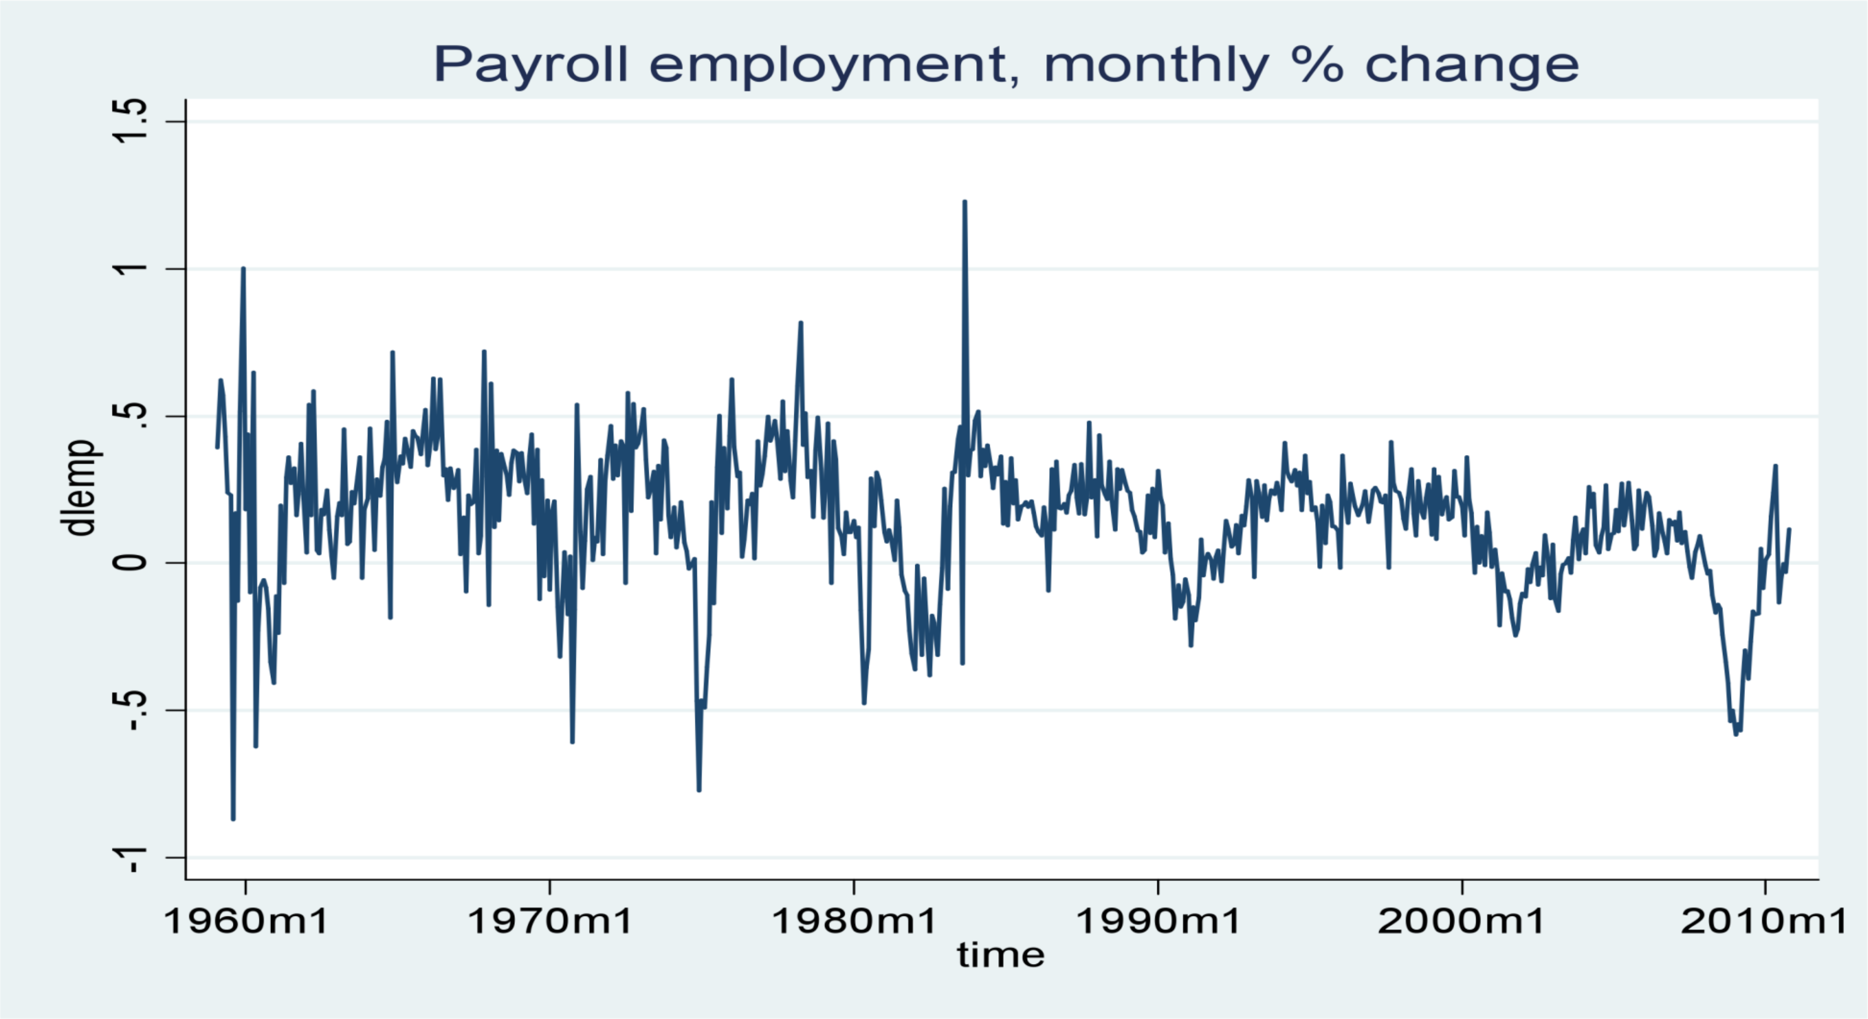
Payroll employment, monthly % change
dlemp
1.5
1
.5
0
-.5
-1
1960m1
1970m1
1980m1
1990m1
2000m1
2010m1
time

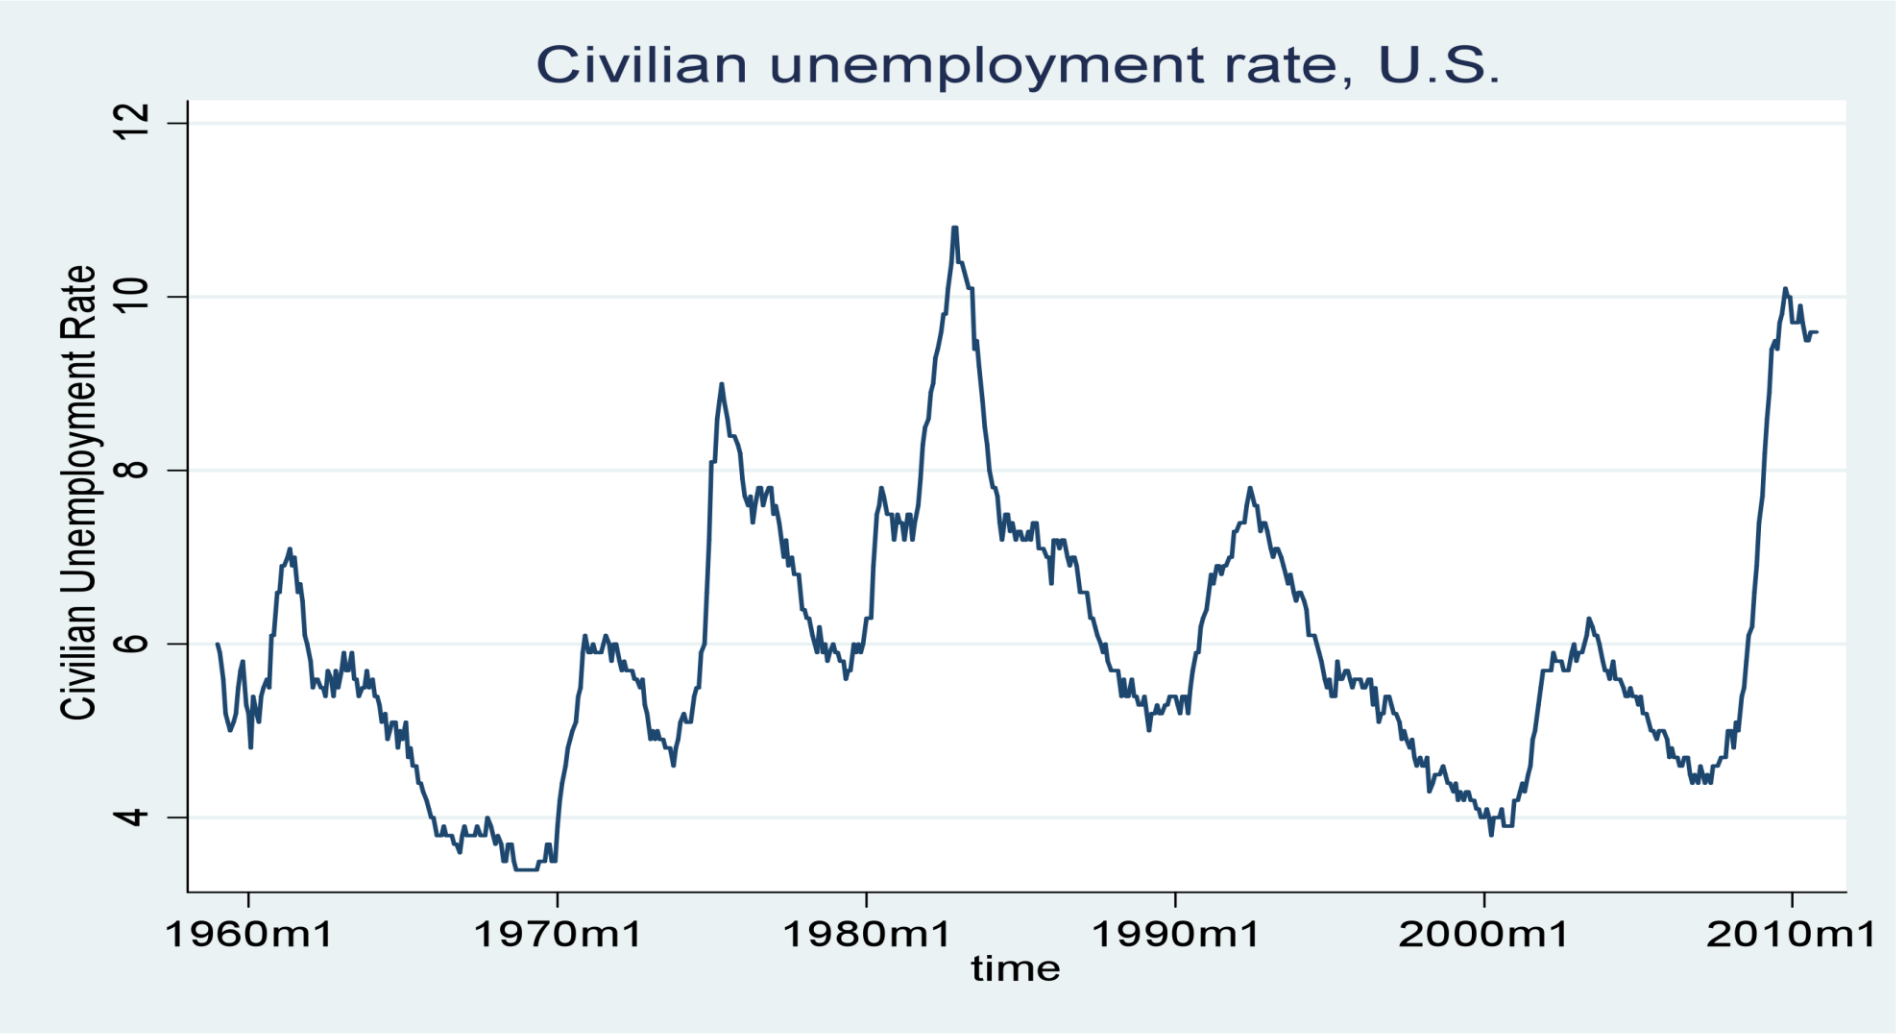

Civilian unemployment rate, U.S.
Civilian Unemployment Rate
12
10
8
6
4
1960m1
1970m1
1980m1
1990m1
2000m1
2010m1
time

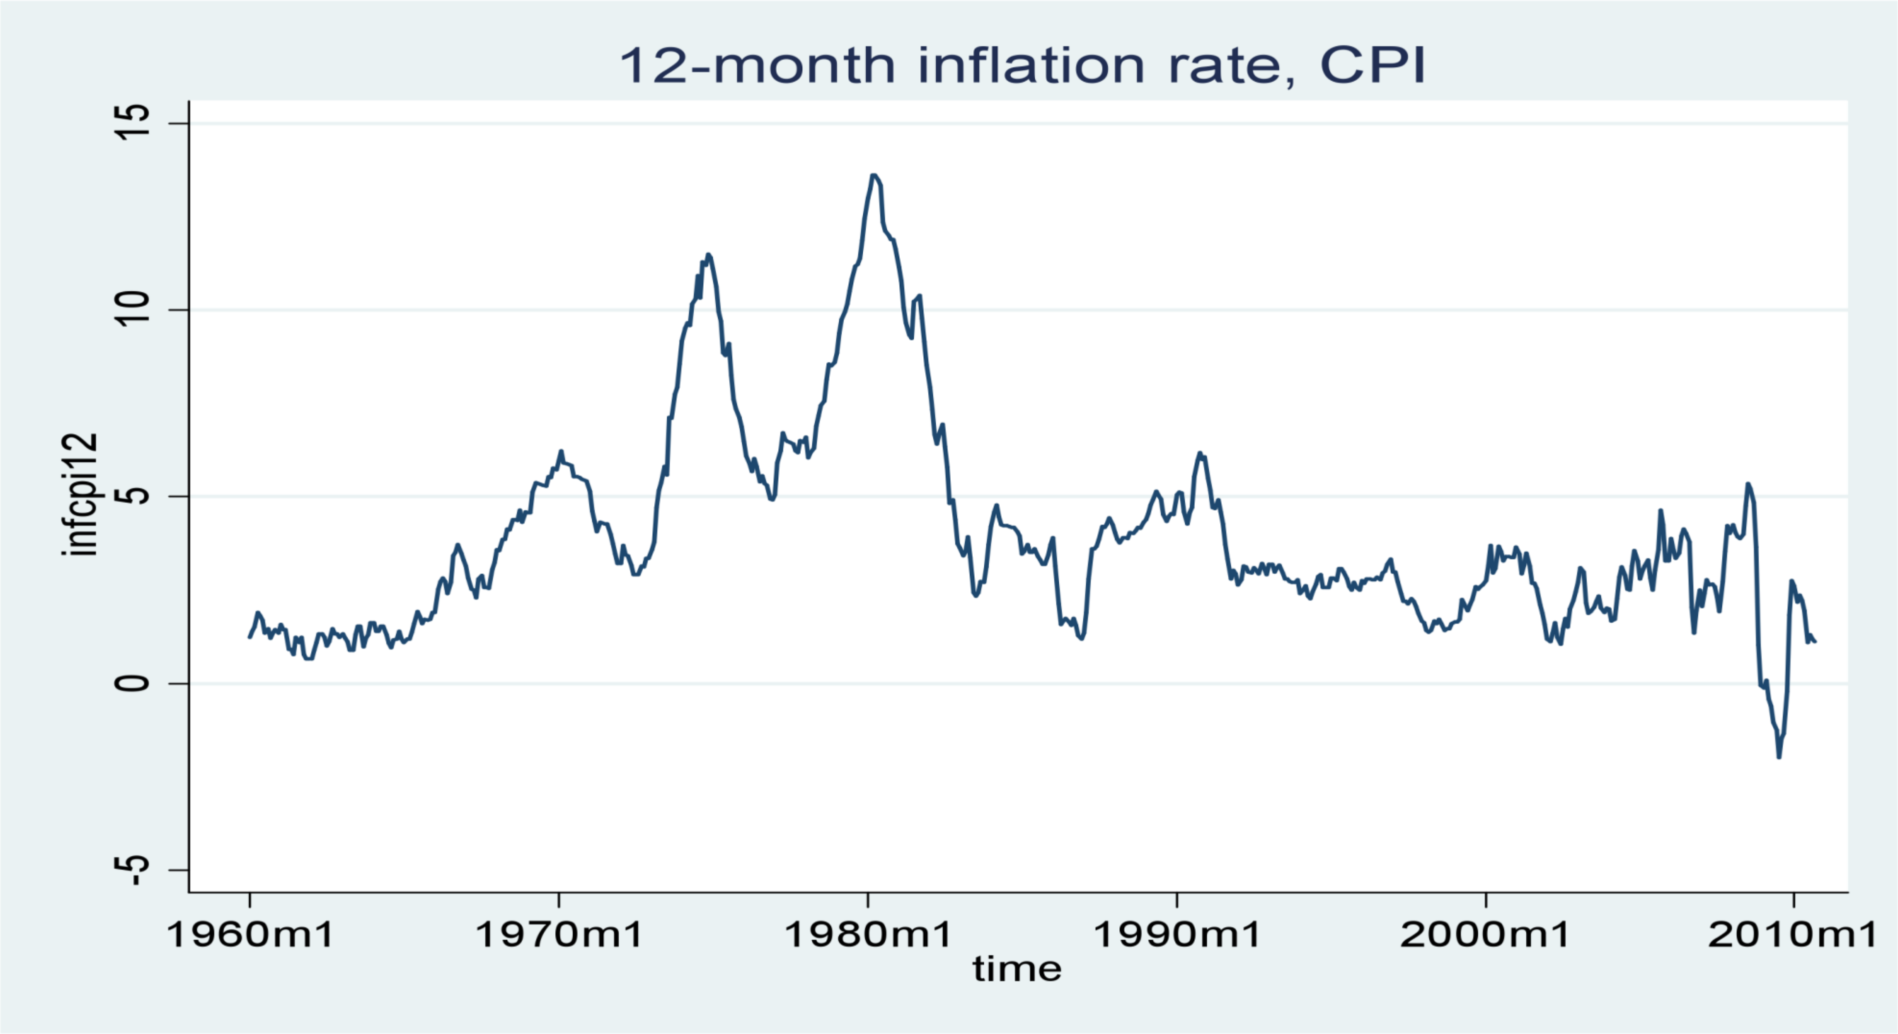
12-month inflation rate, CPI
infcpi12
15
10
5
0
-5
1960m1
1970m1
1980m1
1990m1
2000m1
2010m1
time

A daily financial time series:

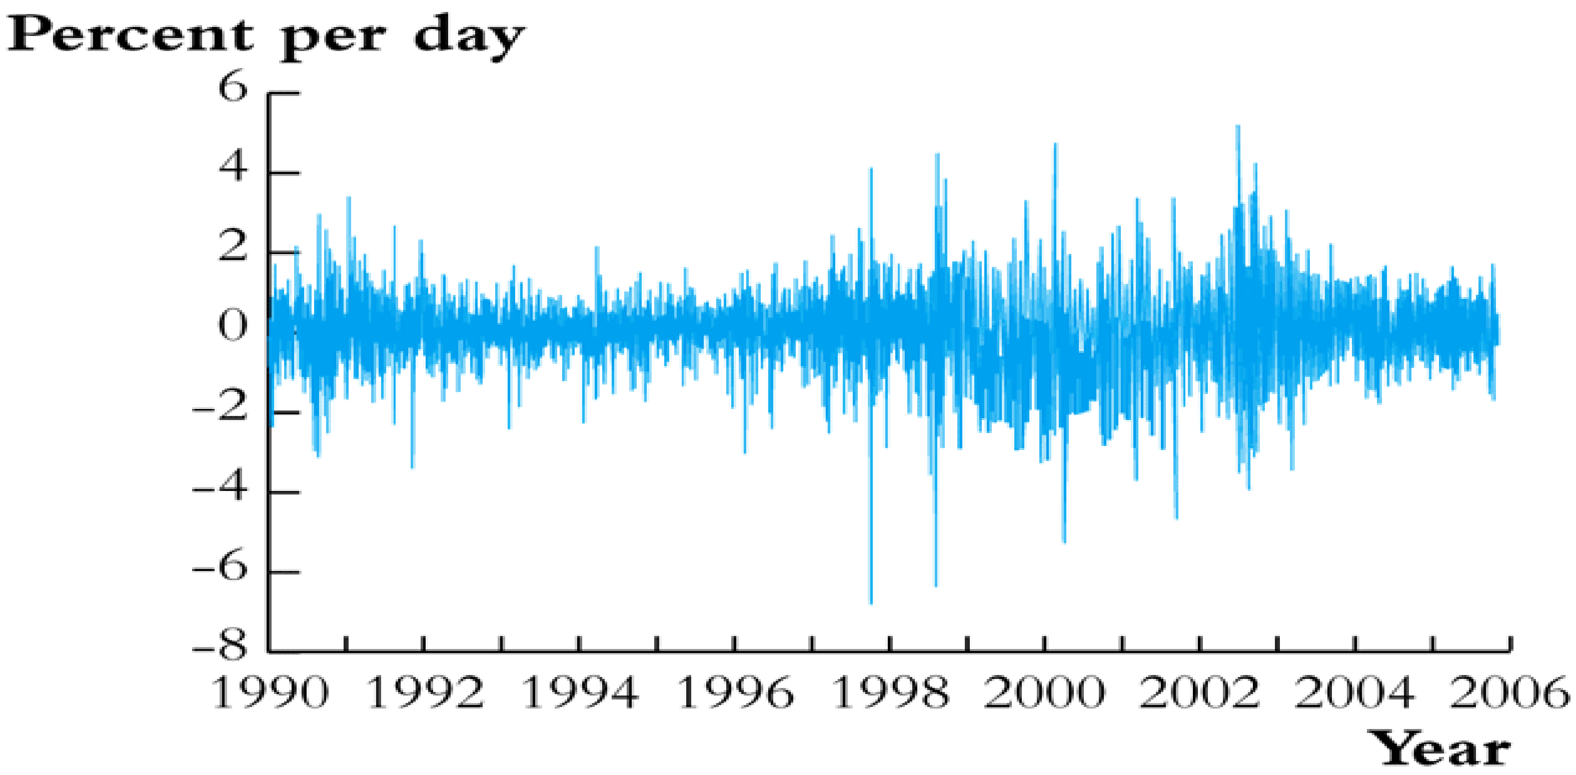

(d) Percentage Changes in Daily Values of the NYSE Composite Stock Index

STOCHASTIC PROCESSES

- *A random or stochastic process is a collection of random variables ordered in time.*
- If we let Y denote a random variable, and if it is continuous, we denote it as $Y(t)$, but if it is discrete, we denoted it as Yt. An example of the former is an electrocardiogram, and an example of the latter is GDP, PDI, etc. Since most economic data are collected at discrete points in time, for our purpose we will use the notation Yt rather than $Y(t)$.
- *Keep in mind that each of these Y's is a random variable.* In what sense can we regard GDP as a stochastic process? Consider for instance the GDP of $2872.8 billion for 1970–I. In theory, the GDP figure for the first quarter of 1970 could have been any number, depending on the economic and political climate then prevailing. The figure of 2872.8 is a particular **realization** of all such possibilities. The distinction between the stochastic process and its realization is akin to the distinction between population and sample in cross-sectional data.

Stationary Stochastic Processes

Forms of Stationarity: weak, and strong

(i) mean: $E(Y_t) = \mu$

(ii) variance: $var(Y_t) = E(Y_t – \mu)^2 = \sigma^2$

(iii) Covariance: $\gamma_k = E[(Y_t – \mu)(Y_{t-k} – \mu)^2$

where *γk*, the covariance (or autocovariance) at lag *k*, is the covariance between the values of *Yt* and *Yt+k*, that is, between two *Y* values *k* periods apart. If *k* = 0, we obtain *γ*0, which is simply the variance of *Y* (= *σ*2); if *k* = 1, *γ*1 is the covariance between two adjacent values of *Y*, the type of covariance we encountered in Chapter 12 (recall the Markov first-order autoregressive scheme).

Types of Stationarity

- A time series is weakly stationary if its mean and variance are constant over time and the value of the covariance between two periods depends only on the distance (or lags) between the two periods.
- A time series if strongly stationary if for any values j_1, j_2,…j_n, the joint distribution of (Y_t, Y_{t+j1}, Y_{t+j2},…Y_{t+jn}) depends only on the intervals separating the dates (j_1, j_2,…,j_n) and not on the date itself (t).
- A weakly stationary series that is Gaussian (normal) is also strictly stationary.
- This is why we often test for the normality of a time series.

Stationarity vs. Nonstationarity

- *A time series is stationary, if its mean, variance, and autocovariance (at various lags) remain the same no matter at what point we measure them; that is, they are time invariant.*
- Such a time series will tend to return to its mean (called **mean reversion**) and fluctuations around this mean (measured by its variance) will have a broadly constant amplitude.
- If a time series is not stationary in the sense just defined, it is called a **nonstationary time series** (keep in mind we are talking only about weak stationarity). In other words, a nonstationary time series will have a *time-varying mean or a time-varying variance or both.*
- Why are stationary time series so important? Because if a time series is nonstationary, we can study its behavior only for the time period under consideration. Each set of time series data will therefore be for a particular episode. As a consequence, it is not possible to generalize it to other time periods. Therefore, for the purpose of forecasting, such time series may be of little practical value.

Examples of Non-Stationary Time Series

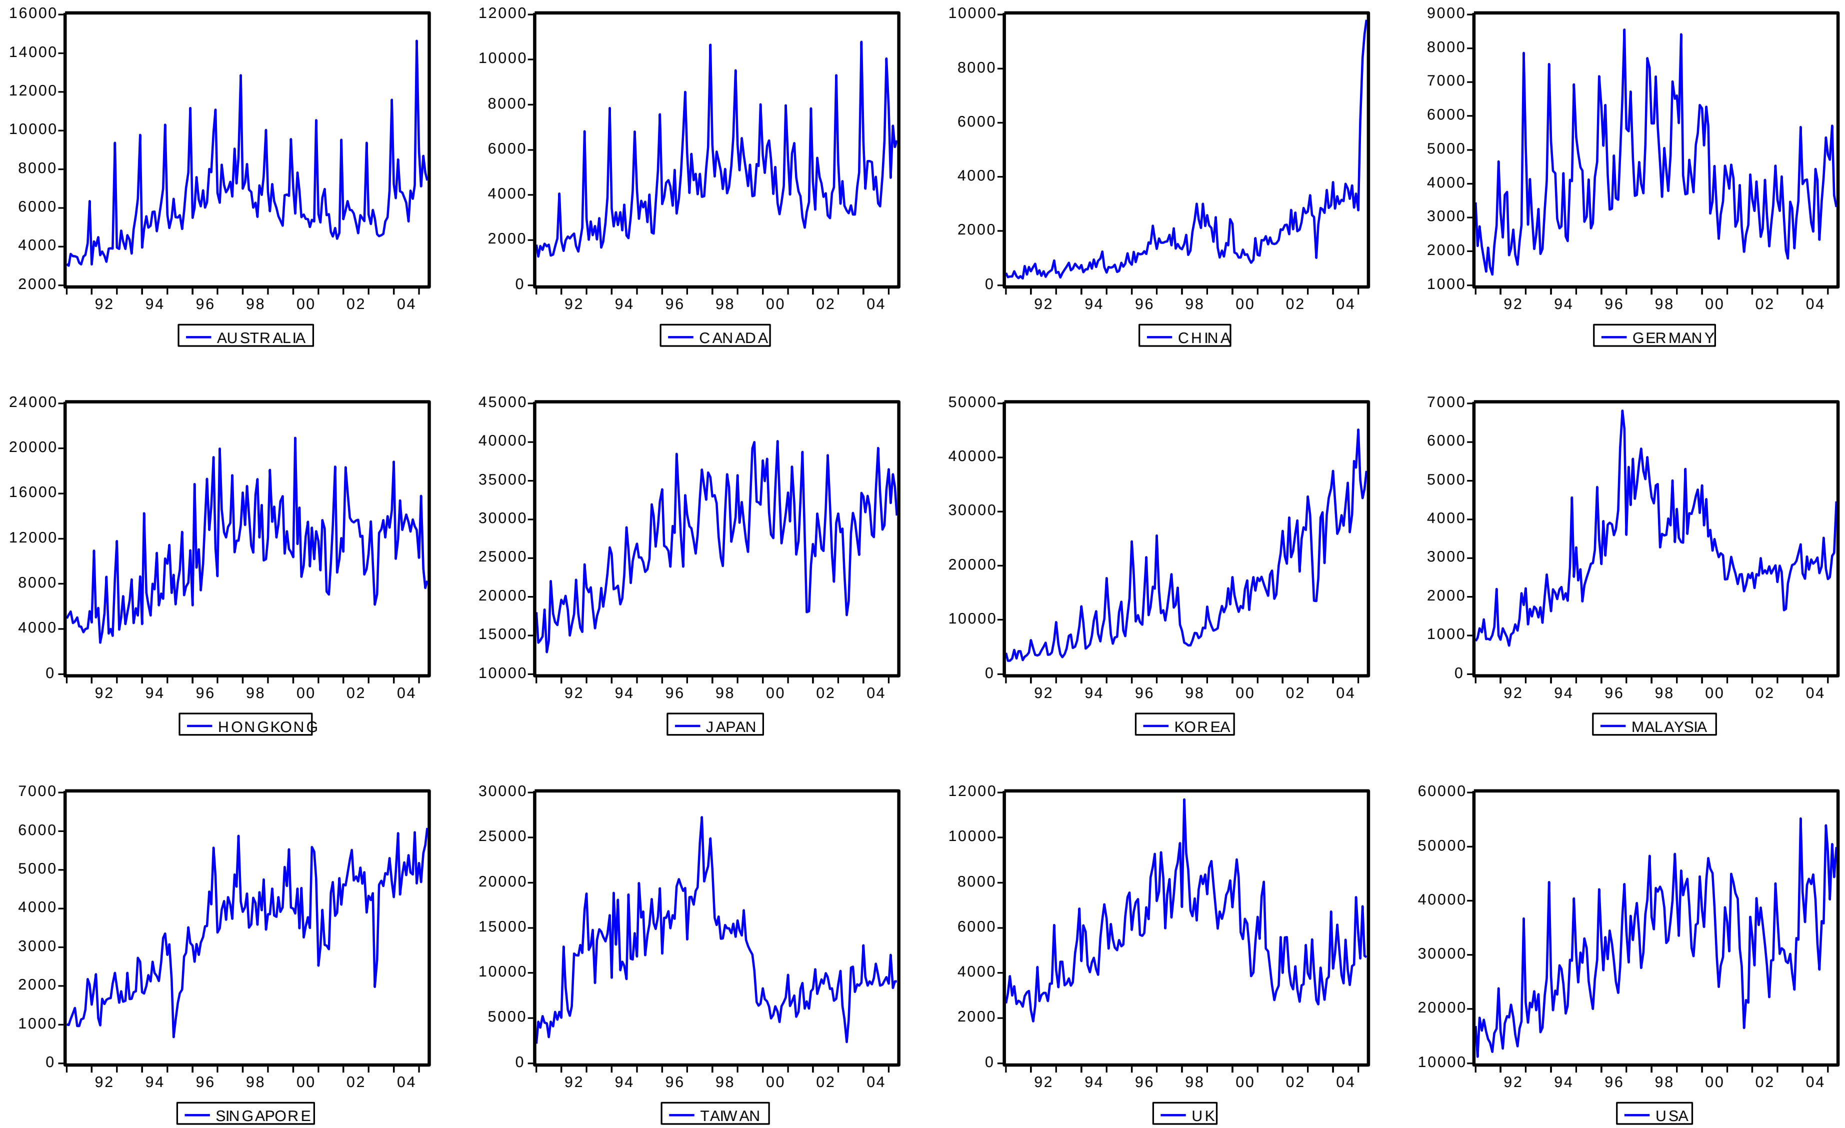

Examples of Stationary Time Series

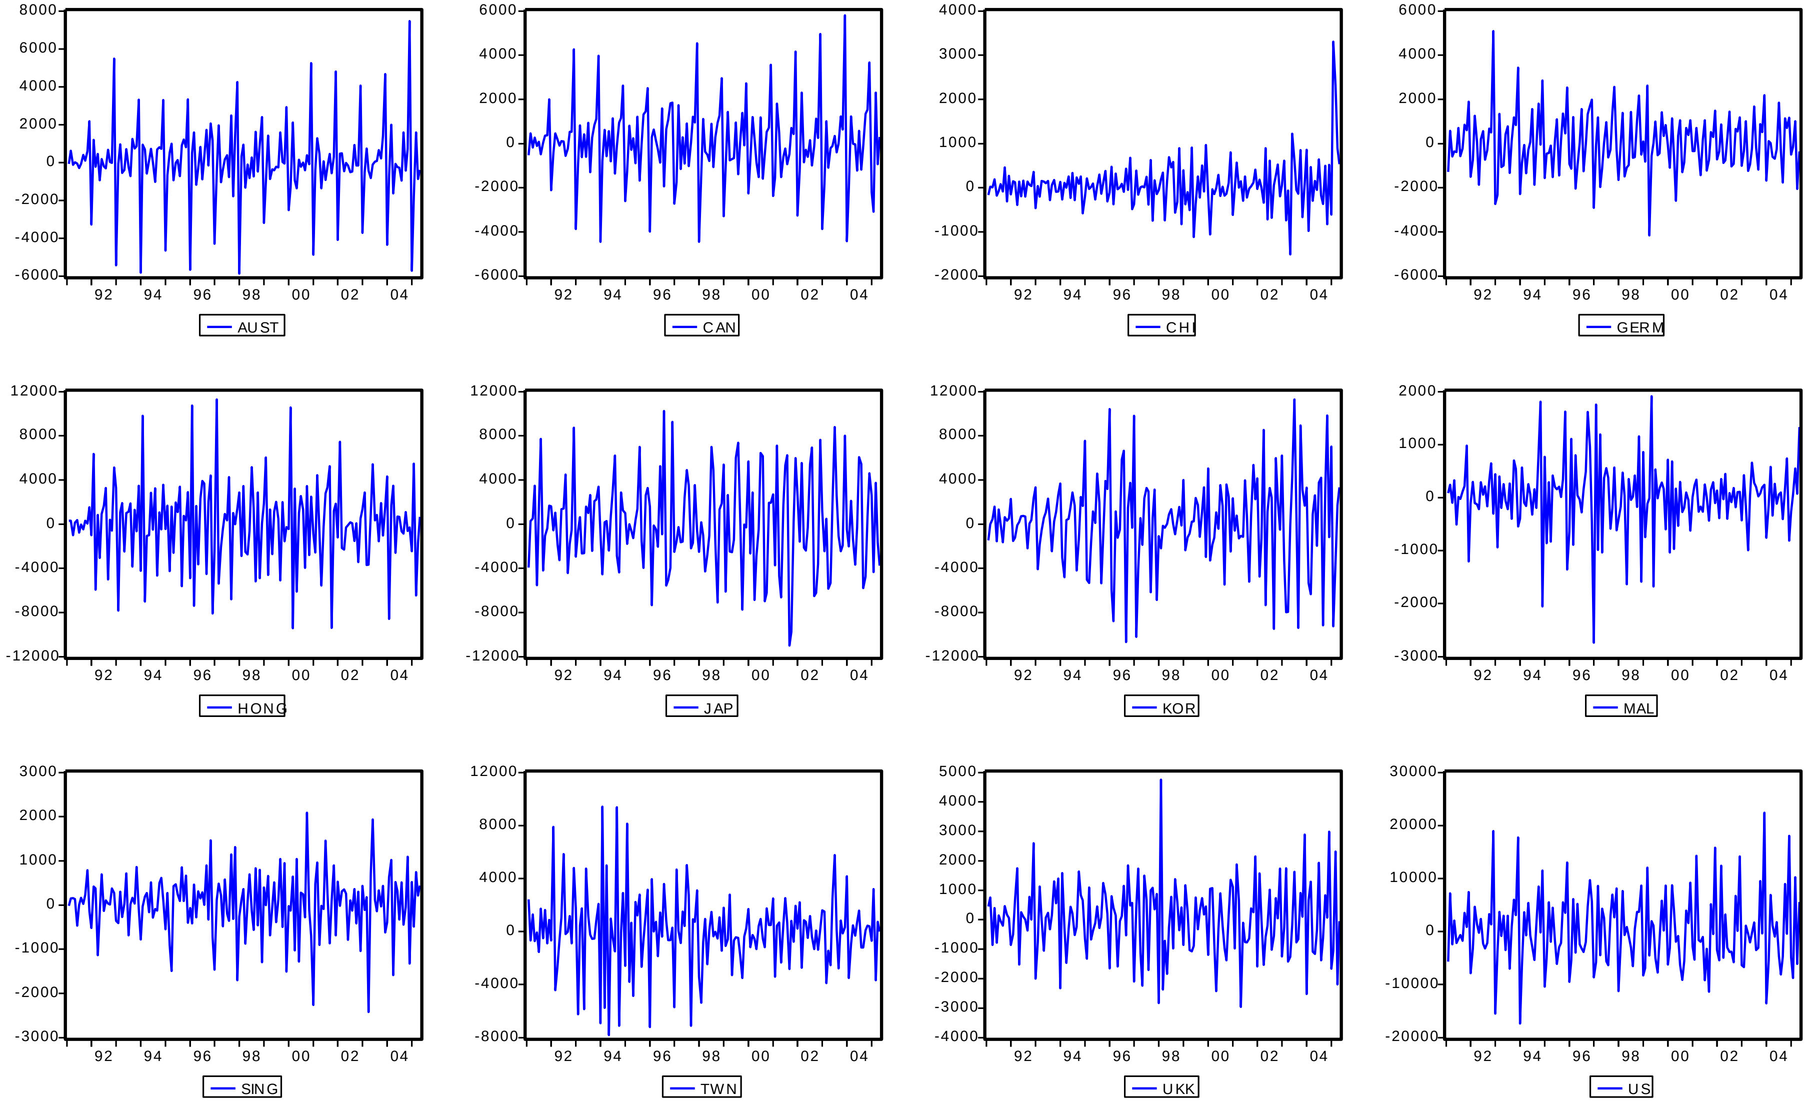

“Unit Root” and order of integration

If a Non-Stationary Time Series Y_t has to be “differenced” d times to make it stationary, then Y_t is said to contain d “Unit Roots”. It is customary to denote Y_t ~ I(d) which reads “Y_t is integrated of order d”

If Y_t ~ I(0), then Y_t is Stationary

If Y_t ~ I(1), then $Z_t = Y_t - Y_{t-1}$ is Stationary

If Y_t ~ I(2), then $Z_t = Y_t - Y_{t-1} - (Y_t - Y_{t-2})$is Stationary

Unit Roots

- Consider an AR(1) process:

$y_t = a_1 y_{t-1} + \varepsilon_t$ (Eq. 1)

$\varepsilon_t \sim N(0, \sigma^2)$

- Case #1: Random walk ($a_1 = 1$)

$y_t = y_{t-1} + \varepsilon_t$

$\Delta y_t = \varepsilon_t$

Unit Roots

- In this model, the variance of the error term, ε_t, increases as t increases, in which case OLS will produce a downwardly biased estimate of a_1 (Hurwicz bias).
- Rewrite equation 1 by subtracting y_{t-1} from both sides:

$y_t – y_{t-1} = a_1y_{t-1} – y_{t-1} + \varepsilon_t$ (Eq. 2)

$\Delta y_t = \delta\, y_{t-1} + \varepsilon_t$

$\delta = (a_1 – 1)$

Unit Roots

- H_0: $\delta = 0$ (there is a unit root)
- H_A: $\delta \neq 0$ (there is not a unit root)
- If $\delta = 0$, then we can rewrite Equation 2 as

$\Delta y_t = \varepsilon_t$

Thus first differences of a random walk time series are stationary, because by assumption, ε_t is purely random.

In general, a time series must be differenced *d* times to become stationary; it is integrated of order *d* or I(d). A stationary series is I(0). A random walk series is I(1).

Tests for Unit Roots

- Dickey-Fuller test
  - Estimates a regression using equation 2
  - The usual t-statistic is not valid, thus D-F developed appropriate critical values.
  - You can include a constant, trend, or both in the test.
  - If you accept the null hypothesis, you conclude that the time series has a unit root.
  - In that case, you should first difference the series before proceeding with analysis.

Tests for Unit Roots

- Augmented Dickey-Fuller test (dfuller in STATA)
  - We can use this version if we suspect there is autocorrelation in the residuals.
  - This model is the same as the DF test, but includes lags of the residuals too.
- Phillips-Perron test (pperron in STATA)
  - Makes milder assumptions concerning the error term, allowing for the ε_t to be weakly dependent and heterogenously distributed.
- Other tests include KPSS test, Variance Ratio test, and Modified Rescaled Range test.
- There are also unit root tests for panel data (Levin et al 2002, Pesaran et al).

Tests for Unit Roots

- These tests have been criticized for having low power (1-probability(Type II error)).
- They tend to (falsely) accept H_o too often, finding unit roots frequently, especially with seasonally adjusted data or series with structural breaks. Results are also sensitive to # of lags used in the test.
- Solution involves increasing the frequency of observations, or obtaining longer time series.

Trend Stationary vs. Difference Stationary

- Traditionally in regression-based time series models, a time trend variable, t, was included as one of the regressors to avoid spurious correlation.
- This practice is only valid if the trend variable is deterministic, not stochastic.
- A trend stationary series has a data generating process (DGP) of:

$$y_t = a_0 + a_1 t + \varepsilon_t$$

Trend Stationary vs. Difference Stationary

- A difference stationary time series has a DGP of:

$$y_t - y_{t-1} = a_0 + \varepsilon_t$$

$$\Delta y_t = a_0 + \varepsilon_t$$

- Run the ADF test with a trend. If the test still shows a unit root (accept H_o), then conclude it is difference stationary. If you reject H_o, you could simply include the time trend in the model.

What is a Spurious Regression?

A Spurious or Nonsensical relationship may result when one Non-stationary time series is regressed against one or more Non-stationary time series

The best way to guard against Spurious Regressions is to check for "Cointegration" of the variables used in time series modeling

Symptoms of Likely Presence of Spurious Regression

- If the R^2 of the regression is greater than the Durbin-Watson Statistic

- If the residual series of the regression has a Unit Root

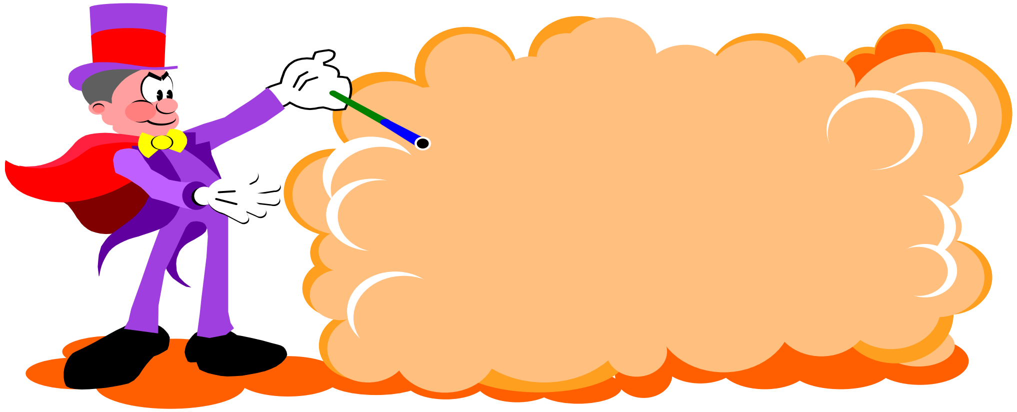

Examples of spurious relationships

- For school children, shoe size is strongly correlated with reading skills.
- Amount of ice cream sold and death by drowning in monthly data
- Number of doctors and number of people dying of disease in cities
- Number of libraries and number of people on drugs in annual data
- **Bottom line: Correlation measures association. But association is not the same as causation.**
- More complicated case:
- Fat in the diet seems to be correlated with cancer. Can we say the diagram is some evidence for the theory?
- But the evidence is quite weak, because other things aren't equal. For example, the countries with lots of fat in the diet also have lots of sugar. A plot of colon cancer rates against sugar consumption would look just like figure 8, and nobody thinks that sugar causes colon cancer. As it turns out, fat and sugar are relatively expensive. In rich countries, people can afford to eat fat and sugar rather than starchier grain products. Some aspects of the diet in these countries, or other factors in the life-style, probably do cause certain kinds of cancer and protect against other kinds. So far, epidemiologists can identify only a few of these factors with any real confidence. Fat is not among them

Figure 8. Cancer rates plotted against fat in the diet, for a sample of countries.

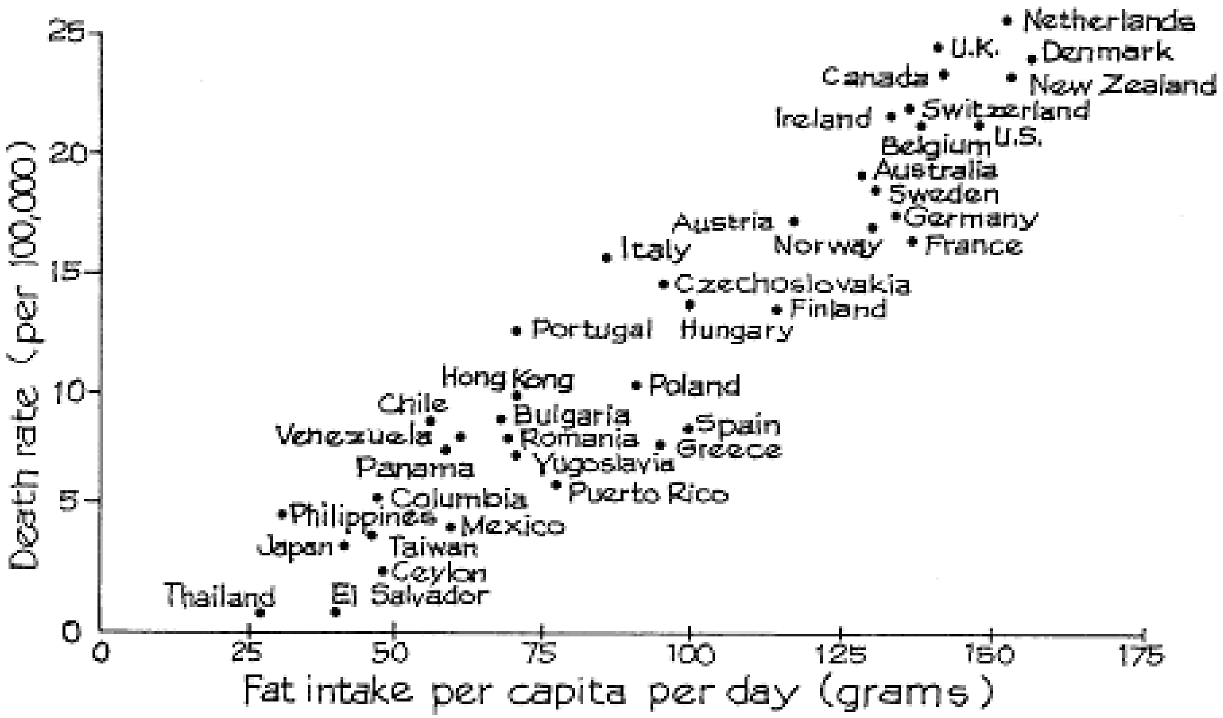

Source: K. Carroll, "Experimental evidence of dietary factors and hormone-dependent cancers," *Cancer Research* vol. 35 (1975) p. 3379. Copyright by *Cancer Research*. Reproduced by permission.

Cointegration

- Is the existence of a long run equilibrium relationship among time series variables
- Is a property of two or more variables moving together through time, and despite following their own individual trends will not drift too far apart since they are linked together in some sense

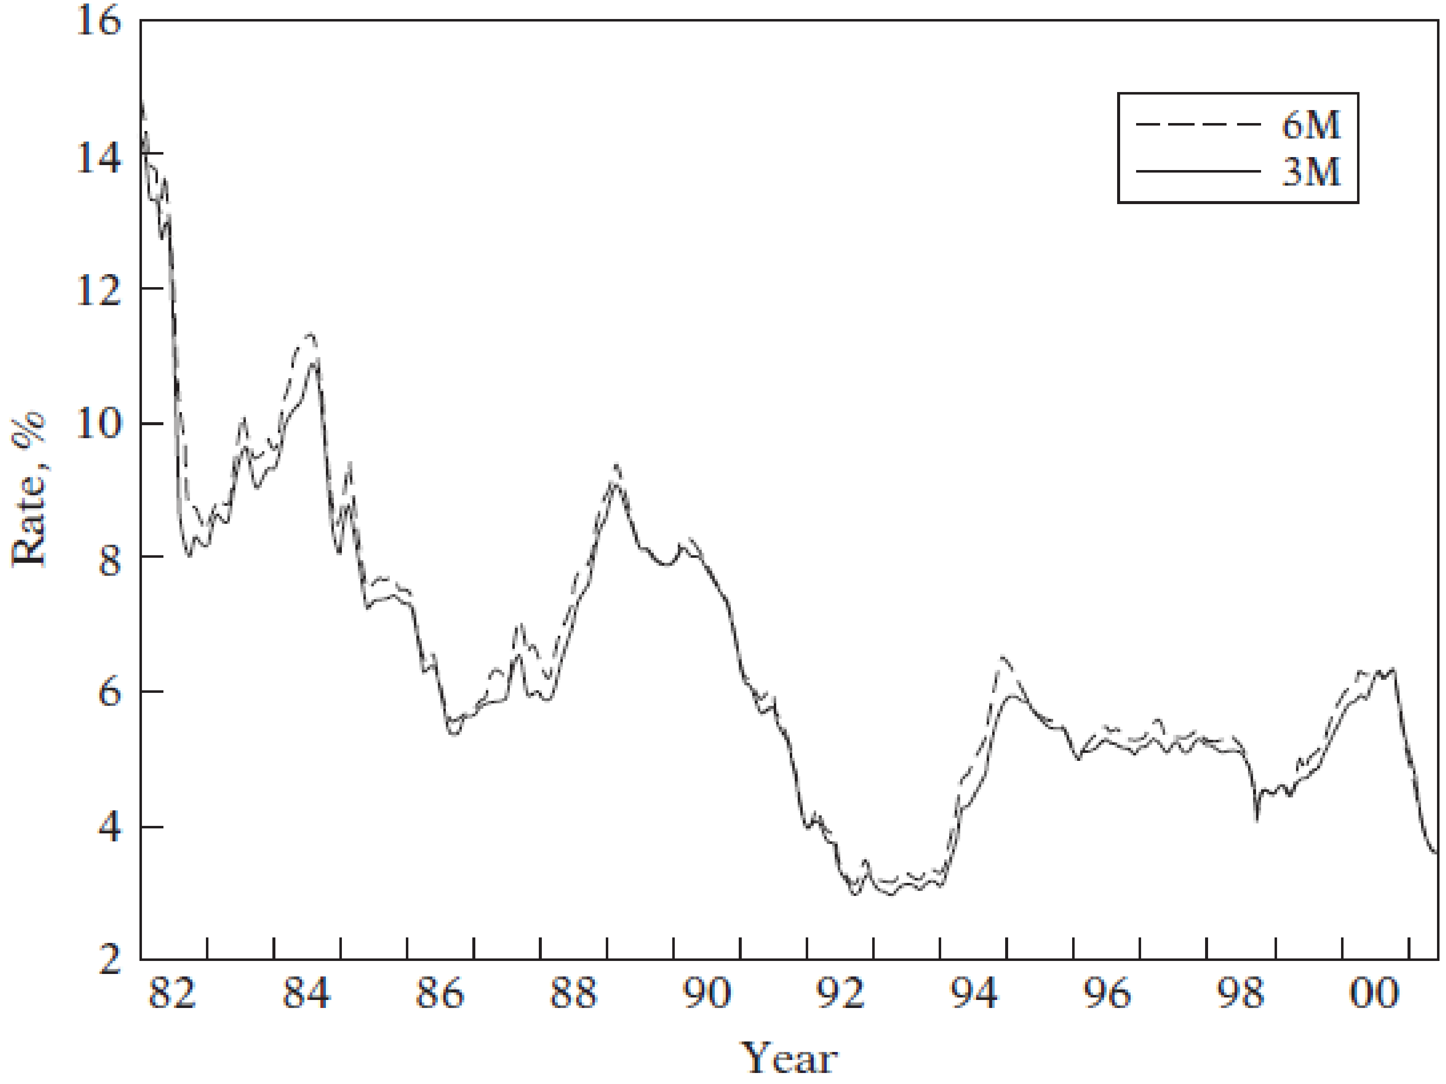

FIGURE 21.13
Three- and six-month Treasury bill rates (constant maturity).

Cointegration Analysis: Formal Tests

- Cointegrating Regression Durbin-Watson (CRDW) Test
- Augmented Engle-Granger (AEG) Test
- Johansen Multivariate Cointegration Tests or the Johansen Method

Error Correction Mechanism (ECM)

- Reconciles the Static LR Equilibrium relationship of Cointegrated Time Series with its Dynamic SR disequilibrium
- Based on the Granger Representation Theorem which states that “If variables are cointegrated, the relationship among them can be expressed as ECM”.

6. Nonstationarity I: Trends

So far, we have assumed that the data are stationary, that is, the distribution of $(Y_{s+1},\ldots, Y_{s+T})$ doesn't depend on s.

If stationarity doesn't hold, the series are said to be ***nonstationary***.

Two important types of nonstationarity are:

- Trends
- Structural breaks (model instability)

Outline of discussion of trends in time series data:

A. What is a trend?
B. Deterministic and stochastic (random) trends
C. How do you detect stochastic trends (statistical tests)?

A. What is a trend?

A trend is a persistent, long-term movement or tendency in the data. Trends need not be just a straight line!

Which of these series has a trend?

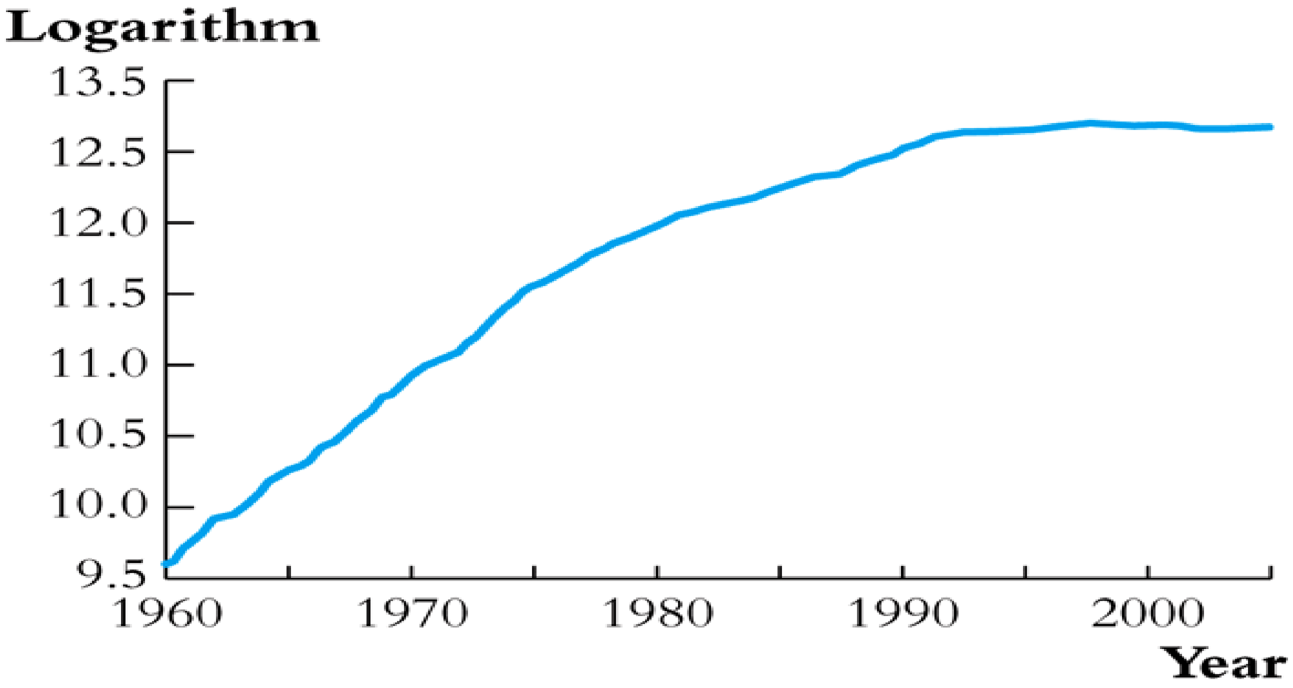

(c) Logarithm of GDP in Japan

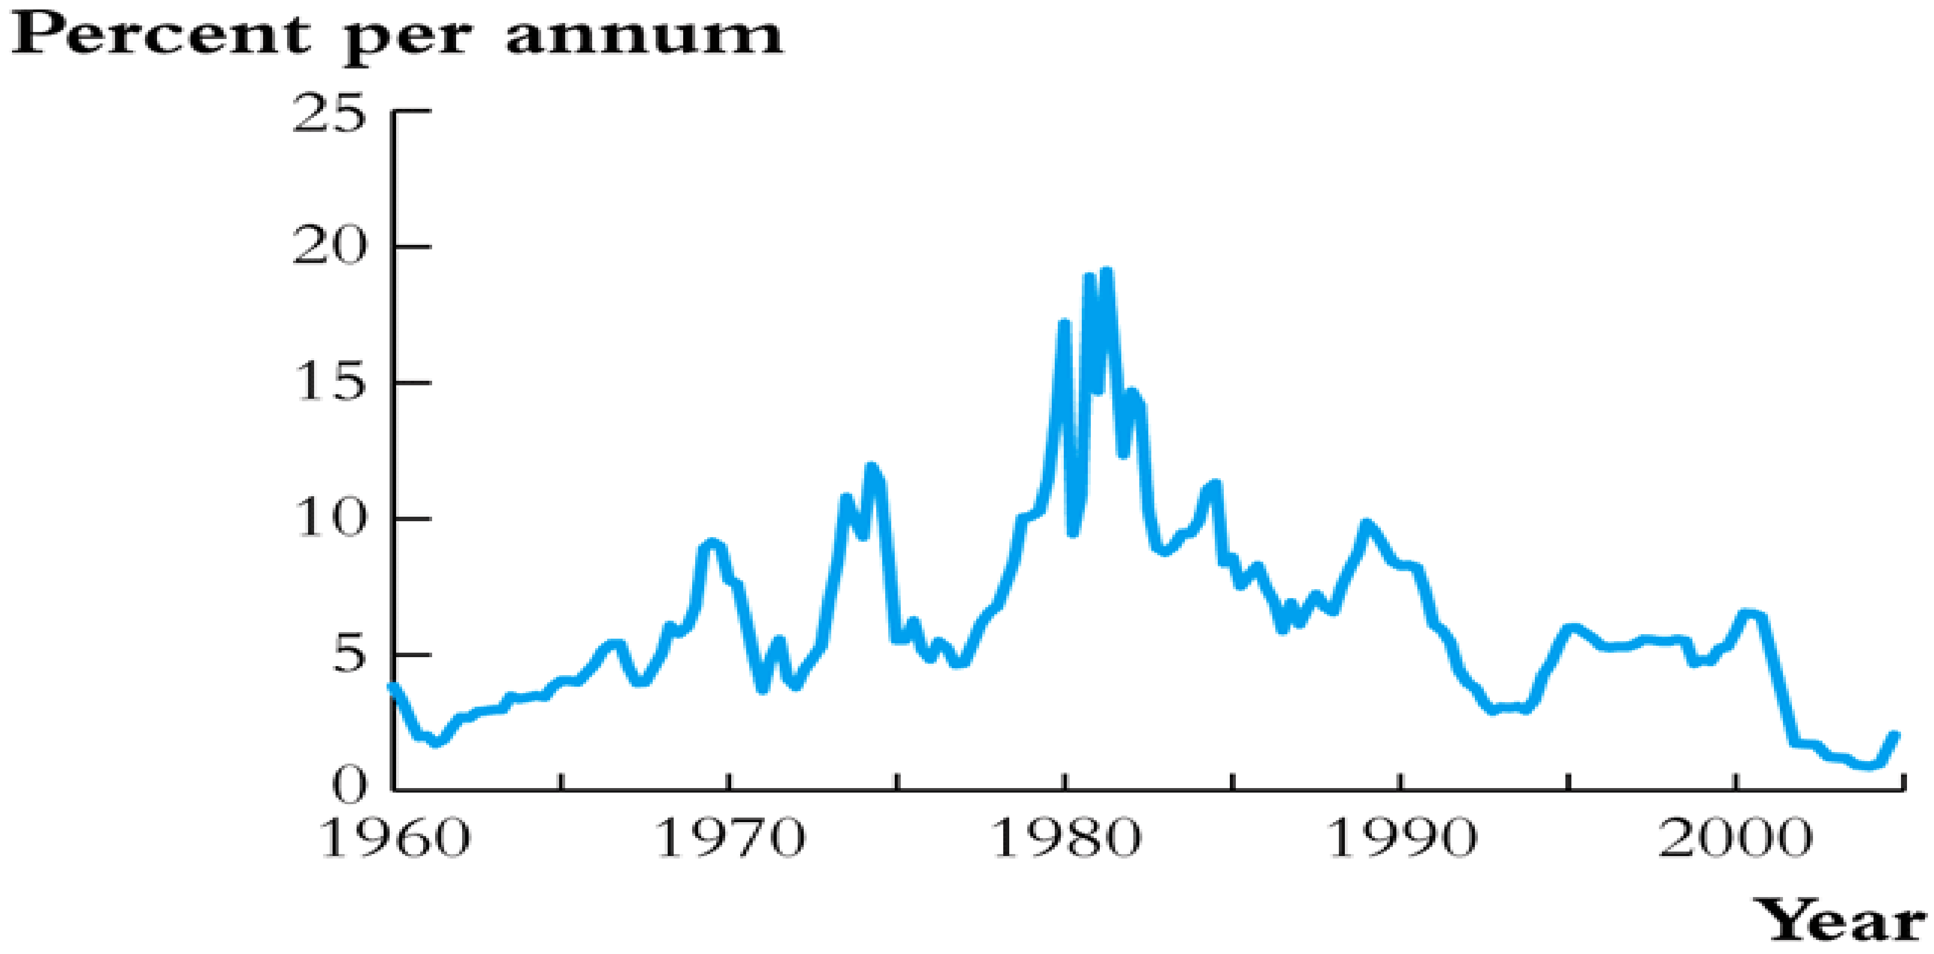

(a) Federal Funds Interest Rate

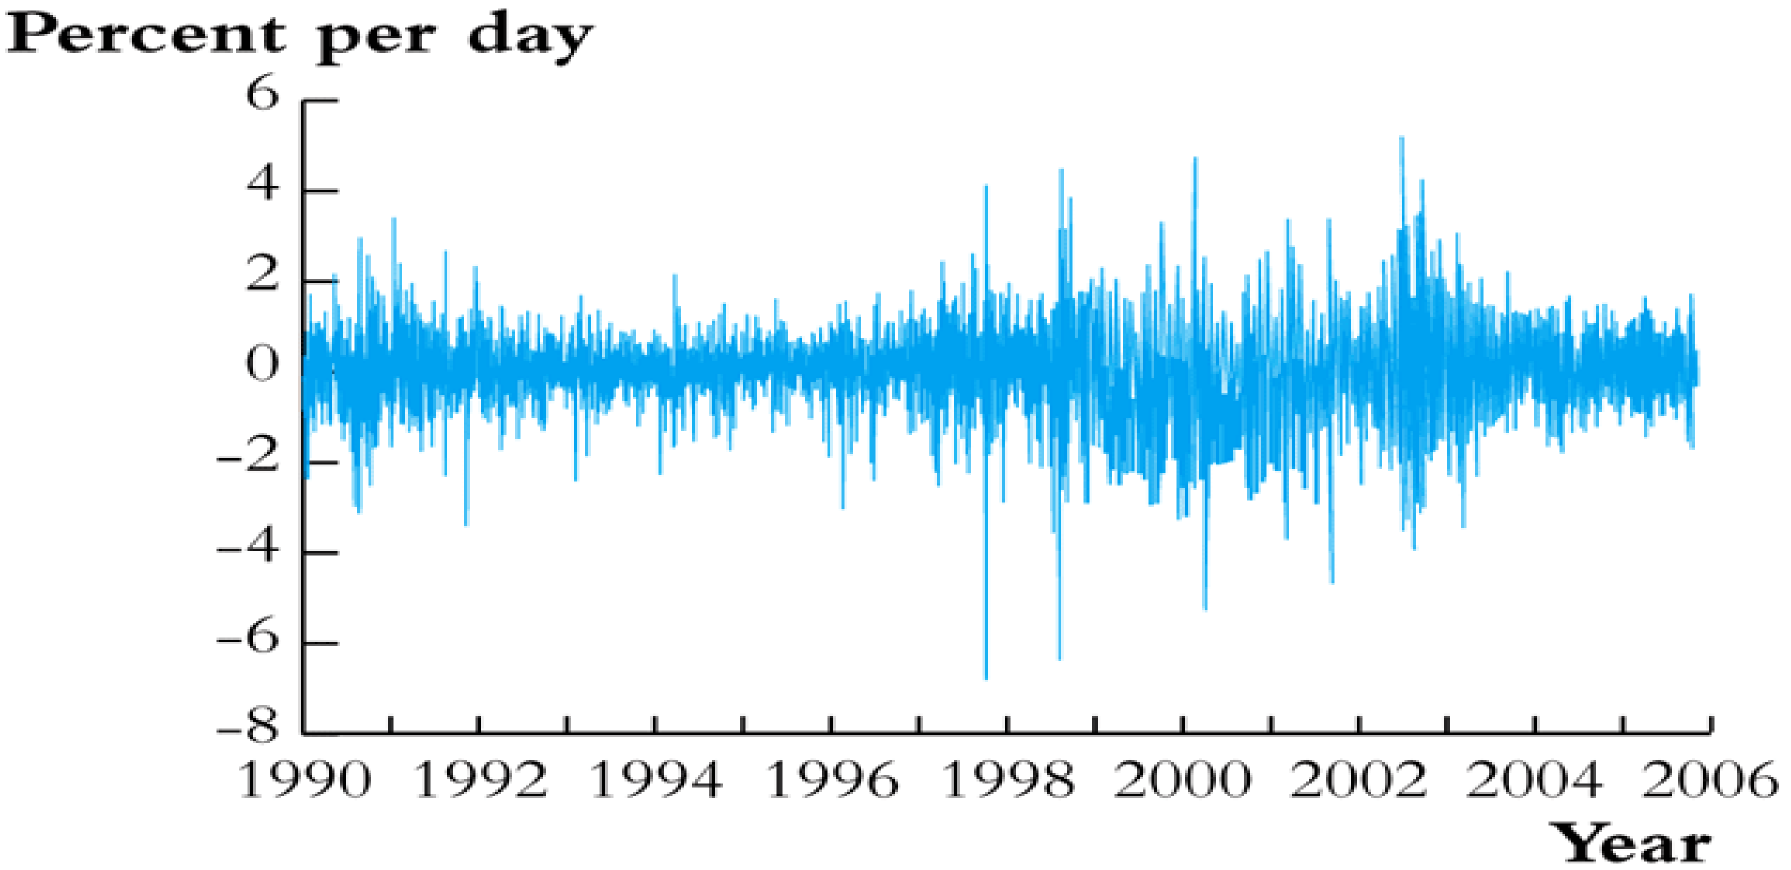

(d) Percentage Changes in Daily Values of the NYSE Composite Stock Index

What is a trend, ctd.

The three series:

- Log Japan GDP clearly has a long-run trend – not a straight line, but a slowly decreasing trend – fast growth during the 1960s and 1970s, slower during the 1980s, stagnating during the 1990s/2000s.
- Inflation has long-term swings, periods in which it is persistently high for many years (’70s/early ’80s) and periods in which it is persistently low. Maybe it has a trend – hard to tell.
- NYSE daily changes has no apparent trend. There are periods of persistently high volatility – but this isn’t a trend.

B. Deterministic and stochastic trends

A trend is a long-term movement or tendency in the data.

- A **deterministic trend** is a nonrandom function of time (e.g. $y_t = t$, or $y_t = t^2$).
- A **stochastic trend** is random and varies over time
- An important example of a stochastic trend is a **random walk**:

$$Y_t = Y_{t-1} + u_t, \text{ where } u_t \text{ is serially uncorrelated}$$

If Y_t follows a random walk, then the value of Y tomorrow is the value of Y today, plus an unpredictable disturbance.

Deterministic and stochastic trends, ctd.

Two key features of a random walk:

(i) $Y_{T+h|T} = Y_T$

- Your best prediction of the value of *Y* in the future is the value of *Y* today
- To a first approximation, log stock prices follow a random walk (more precisely, stock returns are unpredictable)

(ii) Suppose $Y_0 = 0$. Then var(Y_t) = $t\sigma^2$.

- This variance depends on *t* (increases linearly with *t*), so Y_t isn't stationary (recall the definition of stationarity).

Deterministic and stochastic trends, ctd.

A **random walk with drift** is

$$Y_t = \beta_0 + Y_{t-1} + u_t, \text{ where } u_t \text{ is serially uncorrelated}$$

The "drift" is β_0: If $\beta_0 \neq 0$, then Y_t follows a random walk around a linear trend. You can see this by considering the *h*-step ahead forecast:

$$Y_{T+h|T} = \beta_0 h + Y_T$$

The random walk model (with or without drift) is a good description of stochastic trends in many economic time series.

C. How do you detect stochastic trends?

1. Plot the data – are there persistent long-run movements?
2. Use a regression-based test for a random walk: the Dickey-Fuller test for a unit root.

<u>The Dickey-Fuller test in an AR(1)</u>

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + u_t$$

or

$$\Delta Y_t = \beta_0 + \delta Y_{t-1} + u_t$$

H_0: $\delta = 0$ (that is, $\beta_1 = 1$) v. H_1: $\delta < 0$

(note: this is 1-sided: $\delta < 0$ means that Y_t is stationary)

DF test in AR(1), ctd.

$\Delta Y_t = \beta_0 + \delta Y_{t-1} + u_t$

H_0: $\delta = 0$ (that is, $\beta_1 = 1$) v. H_1: $\delta < 0$

DF test: compute the t-statistic testing $\delta = 0$

- Under H_0, this t statistic does **not** have a normal distribution!
- You need to use the table of Dickey-Fuller critical values. There are two cases, which have different critical values:

  (a) $\Delta Y_t = \beta_0 + \delta Y_{t-1} + u_t$ (intercept only)

  (b) $\Delta Y_t = \beta_0 + \mu t + \delta Y_{t-1} + u_t$ (intercept & time trend)

The Dickey-Fuller Test in an AR(p)

In an AR(p), the DF test is based on the rewritten model,

$\Delta Y_t = \beta_0 + \delta Y_{t-1} + \gamma_1 \Delta Y_{t-1} + \gamma_2 \Delta Y_{t-2} + \ldots + \gamma_{p-1} \Delta Y_{t-p+1} + u_t$ (*)

where $\delta = \beta_1 + \beta_2 + \ldots + \beta_p - 1$. If there is a unit root (random walk trend), $\delta = 0$; if the AR is stationary, $\delta < 1$.

<u>The DF test in an AR(p) (intercept only):</u>

1. Estimate (*), obtain the t-statistic testing $\delta = 0$
2. Reject the null hypothesis of a unit root if the t-statistic is less than the DF critical value

When should you include a time trend in the DF test?

The decision to use the intercept-only DF test or the intercept & trend DF test depends on what the alternative is – and what the data look like.

- In the intercept-only specification, the alternative is that Y is stationary around a constant – no long-term growth in the series
- In the intercept & trend specification, the alternative is that Y is stationary around a linear time trend – the series has long-term growth.

Example: Does U.S. inflation have a unit root?

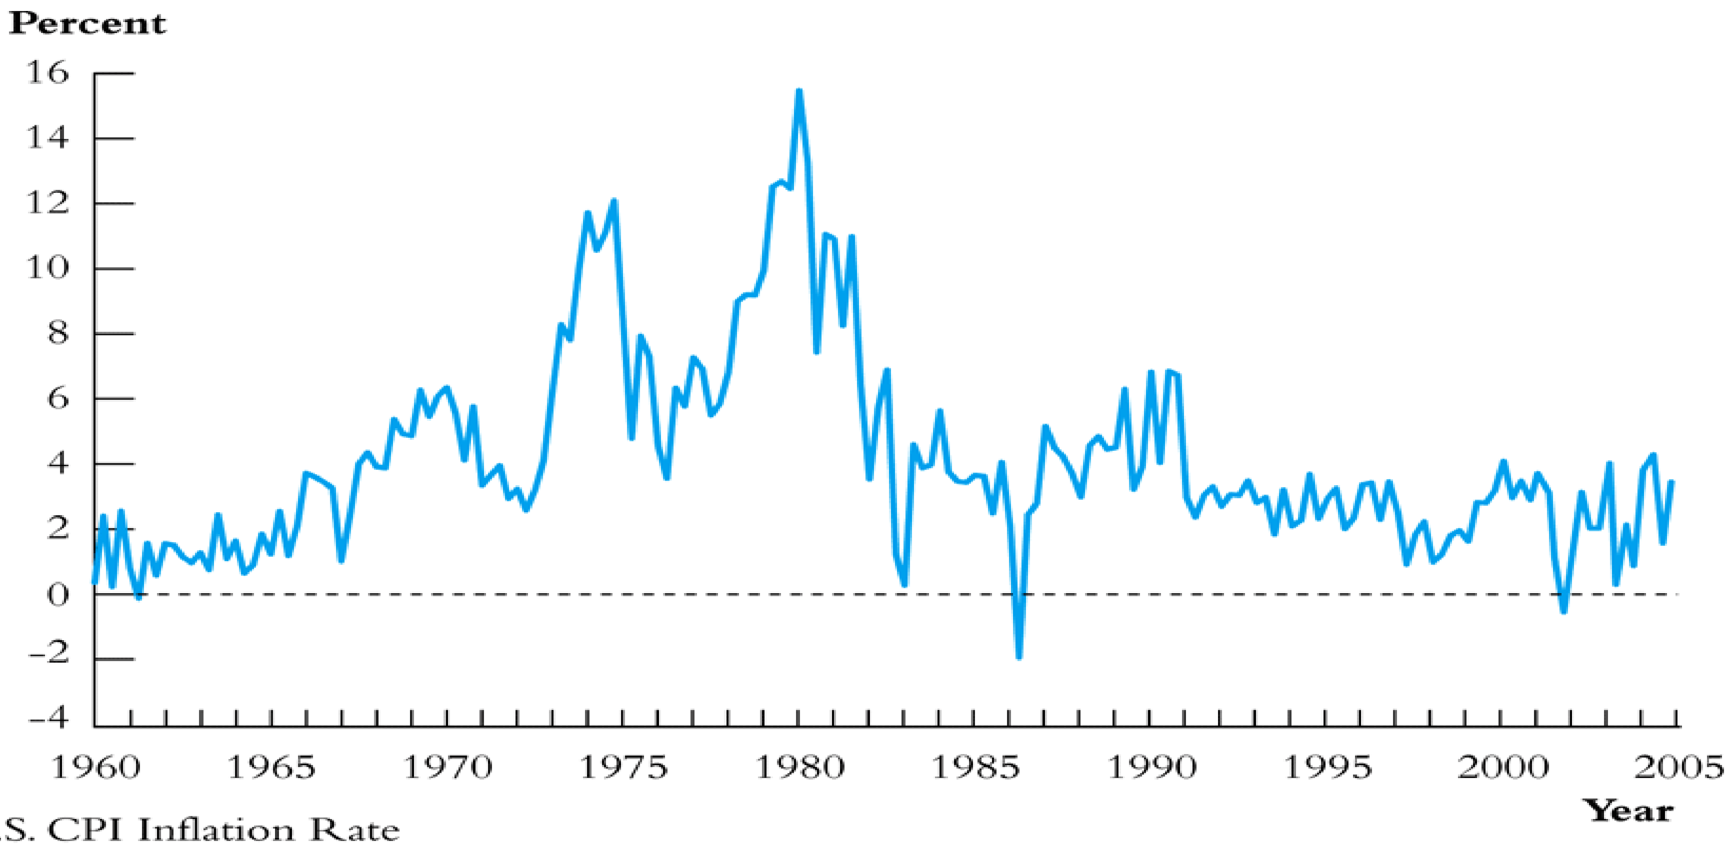

(a) U.S. CPI Inflation Rate

The alternative is that inflation is stationary around a constant

Does U.S. inflation have a unit root? Ctd

DF test for a unit root in U.S. inflation – using *p* = 4 lags

```
. reg dinf L.inf L(1/4).dinf if tin(1962q1,2004q4);

      Source |       SS       df       MS              Number of obs =     172
-------------+------------------------------           F(  5,   166) =   10.31
       Model |  118.197526     5  23.6395052           Prob > F      =  0.0000
    Residual |  380.599255   166   2.2927666           R-squared     =  0.2370
-------------+------------------------------           Adj R-squared =  0.2140
       Total |  498.796781   171  2.91694024           Root MSE      =  1.5142

------------------------------------------------------------------------------
        dinf |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         inf |
         L1. |  -.1134149   .0422339    -2.69   0.008    -.1967998     -.03003
        dinf |
         L1. |  -.1864226   .0805141    -2.32   0.022    -.3453864   -.0274589
         L2. |   -.256388   .0814624    -3.15   0.002     -.417224   -.0955519
         L3. |    .199051   .0793508     2.51   0.013     .0423842    .3557178
         L4. |   .0099822   .0779921     0.13   0.898     -.144002    .1639665
       _cons |   .5068071    .214178     2.37   0.019     .0839431     .929671
------------------------------------------------------------------------------
```

DF *t*-statistic = –2.69 (intercept-only):

TABLE 14.5 Large-Sample Critical Values of the Augmented Dickey–Fuller Statistic

Deterministic Regressors	10%	5%	1%
Intercept only	−2.57	−2.86	−3.43
Intercept and time trend	−3.12	−3.41	−3.96

Reject if the DF *t*-statistic (the *t*-statistic testing $\delta = 0$) is less than the specified critical value. This is a 1-sided test of the null hypothesis of a unit root (random walk trend) vs. the alternative that the autoregression is stationary.

t = –2.69 rejects a unit root at 10% level but not the 5% level

- Some evidence of a unit root – not clear cut.
- Whether the inflation rate has a unit root is hotly debated among empirical monetary economists.

7. Nonstationarity II: Breaks

The second type of nonstationarity we consider is that the coefficients of the model might not be constant over the full sample. Clearly, it is a problem for forecasting if the model describing the historical data differs from the current model – you want the current model for your forecasts!

So we will:

- Go over the way to detect changes in coefficients: tests for a break
- Work through an example: the U.S. Phillips curve

A. Tests for a break (change) in regression coefficients

Case I: The break date is known

Suppose the break is known to have occurred at date τ. Stability of the coefficients can be tested by estimating a fully interacted regression model. In the ADL(1,1) case:

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + \delta_1 X_{t-1}$$

$$+ \gamma_0 D_t(\tau) + \gamma_1[D_t(\tau)\times Y_{t-1}] + \gamma_2[D_t(\tau)\times X_{t-1}] + u_t$$

where $D_t(\tau) = 1$ if $t \geq \tau$, and $= 0$ otherwise.

If $\gamma_0 = \gamma_1 = \gamma_2 = 0$, then the coefficients are constant over the full sample.

If at least one of γ_0, γ_1, or γ_2 are nonzero, the regression function changes at date τ.

$Y_t = \beta_0 + \beta_1 Y_{t-1} + \delta_1 X_{t-1}$

$+ \gamma_0 D_t(\tau) + \gamma_1[D_t(\tau)\times Y_{t-1}] + \gamma_2[D_t(\tau)\times X_{t-1}] + u_t$

where $D_t(\tau) = 1$ if $t \geq \tau$, and $= 0$ otherwise

The ***Chow test statistic*** for a break at date τ is the (heteroskedasticity-robust) F-statistic that tests:

H_0: $\gamma_0 = \gamma_1 = \gamma_2 = 0$

vs. H_1: at least one of γ_0, γ_1, or γ_2 are nonzero

- Note that you can apply this to a subset of the coefficients, e.g. only the coefficient on X_{t-1}.
- Unfortunately, you often don't have a candidate break date, that is, you don't know τ …

Case II: The break date is unknown

Why consider this case?

- You might suspect there is a break, but not know when
- You might want to test the null hypothesis of coefficient stability against the general alternative that there has been a break sometime.
- Even if you think you know the break date, if that "knowledge" is based on prior inspection of the series then you have in effect "estimated" the break date. This invalidates the Chow test critical values.

The Quandt Likelihood Ratio (QLR) Statistic
(also called the "sup-Wald" statistic)

The QLR statistic = the maximum Chow statistic

- Let $F(\tau)$ = the Chow test statistic testing the hypothesis of no break at date τ.
- The *QLR* test statistic is the ***maximum*** of all the Chow *F*-statistics, over a range of τ, $\tau_0 \leq \tau \leq \tau_1$:

$$QLR = \max[F(\tau_0), F(\tau_0+1), \ldots, F(\tau_1-1), F(\tau_1)]$$

- A conventional choice for τ_0 and τ_1 are the inner 70% of the sample (exclude the first and last 15%).
- Should you use the usual $F_{q,\infty}$ critical values?

TABLE 14.6 Critical Values of the QLR Statistic with 15% Trimming

Number of Restrictions (q)	10%	5%	1%
1	7.12	8.68	12.16
2	5.00	5.86	7.78
3	4.09	4.71	6.02
4	3.59	4.09	5.12
5	3.26	3.66	4.53
6	3.02	3.37	4.12
7	2.84	3.15	3.82
8	2.69	2.98	3.57
9	2.58	2.84	3.38
10	2.48	2.71	3.23
11	2.40	2.62	3.09
12	2.33	2.54	2.97
13	2.27	2.46	2.87
14	2.21	2.40	2.78
15	2.16	2.34	2.71
16	2.12	2.29	2.64
17	2.08	2.25	2.58
18	2.05	2.20	2.53
19	2.01	2.17	2.48
20	1.99	2.13	2.43

These critical values apply when $\tau_0 = 0.15T$ and $\tau_1 = 0.85T$ (rounded to the nearest integer), so the F-statistic is computed for all potential break dates in the central 70% of the sample. The number of restrictions q is the number of restrictions tested by each individual F-statistic. Critical values for other trimming percentages are given in Andrews (2003).

Note that these critical values are larger than the $F_{q,\infty}$ critical values – for example, $F_{1,\infty}$ 5% critical value is 3.84.

Example: Has the postwar U.S. Phillips Curve been stable?

Recall the ADL(4,4) model of ΔInf_t and $Unemp_t$ – the empirical backwards-looking Phillips curve, estimated over (1962 – 2004):

$$\Delta Inf_t = \underset{(.44)}{1.30} - \underset{(.08)}{.42}\Delta Inf_{t-1} - \underset{(.09)}{.37}\Delta Inf_{t-2} + \underset{(.08)}{.06}\Delta Inf_{t-3} - \underset{(.08)}{.04}\Delta Inf_{t-4}$$

$$- \underset{(.46)}{2.64}Unem_{t-1} + \underset{(.86)}{3.04}Unem_{t-2} - \underset{(.89)}{0.38}Unem_{t-3} + \underset{(.45)}{.25}Unemp_{t-4}$$

Has this model been stable over the full period 1962-2004?

QLR tests of the stability of the U.S. Phillips curve.

dependent variable: ΔInf_t

regressors: intercept, ΔInf_{t-1},…, ΔInf_{t-4}, $Unemp_{t-1}$,…, $Unemp_{t-4}$

- test for constancy of intercept only (other coefficients are assumed constant): QLR = 2.865 (q = 1).
  - 10% critical value = 7.12 → don't reject at 10% level
- test for constancy of intercept and coefficients on $Unemp_t$,…, $Unemp_{t-3}$ (coefficients on ΔInf_{t-1},…, ΔInf_{t-4} are constant): QLR = 5.158 (q = 5)
  - 1% critical value = 4.53 → reject at 1% level
  - Estimate break date: maximal F occurs in 1981:IV
- Conclude that there is a break in the inflation – unemployment relation, with estimated date of 1981:IV

F-Statistics Testing for a Break at Different Dates