STA 542: Introduction to Time Series Analysis

Lecture 9: Assessing a stationary approximation

Lasse Vuursteen

Duke University

Wed Sep 23, 2026

After removing the mean, can shocks still accumulate?

Begin after the deterministic-mean adjustment from the filtering lecture. For a specified mean $m(t)$, write

$$Y_t = m(t) + X_t, \qquad X_t = Y_t - m(t), \qquad \mathbb{E}[X_t] = 0.$$

Is the centered remainder stationary, or does it accumulate shocks?

For the first models, take $m(t)$ as known. We return to unknown means and trends when specifying the test regressions.

Do oil prices accumulate shocks?

Weekly West Texas Intermediate spot prices, in dollars per barrel.

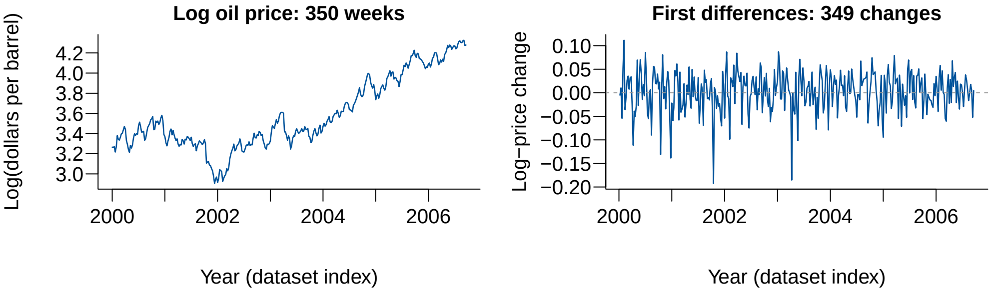

These log prices have not been centered; their deterministic mean is unknown.

Real data: `astsa::oil`, EIA; first 350 weekly prices, 2000–2006.

A long excursion can occur in a stationary process

$X_t = 0.98X_{t-1} + Z_t$ on the left; $X_t = X_{t-1} + Z_t$ on the right.

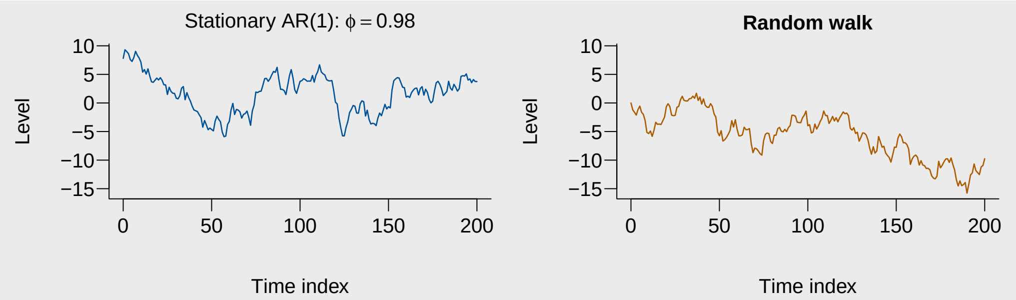

The left process starts in its stationary law; the right starts at zero. How much can we distinguish from 200 transitions?

Simulated data; independent Gaussian innovations of variance one.

What kind of stationary fluctuations are we considering?

Definition ($I(0)$, the convention used here)

A process is $I(0)$ if it is weakly stationary, its ACVF is absolutely summable, and its long-run variance satisfies

$$0 < v := \sum_{h \in \mathbb{Z}} \gamma(h) < \infty.$$

Recall: the variance of a sum of T centered observations is approximately Tv.

We start with mean zero. The definition also allows a constant nonzero mean.

An integrated level accumulates these fluctuations

Definition ($I(1)$, mean-zero case)

A centered level process is $I(1)$ if

$$X_t = X_0 + \sum_{j=1}^{t} V_j, \qquad t \geq 1,$$

where (V_t) is centered and $I(0)$ with long-run variance v, $\mathbb{E}[X_0] = 0$, and $\mathbb{E}[X_0^2] < \infty$.

Its differences satisfy $\nabla X_t = V_t$ and its mean stays zero.

Its uncertainty grows: $\mathrm{Var}(X_t)/t \to v > 0$.

Workbook: *Variance of an accumulated level*

Are stationary differences enough to call a level $I(1)$?

Let (Z_t) be i.i.d., centered, with variance $\sigma^2 > 0$.

$$X_t = Z_t \quad \Longrightarrow \quad \nabla X_t = Z_t - Z_{t-1}.$$

The differences are stationary. Does this describe accumulating uncertainty?

Are stationary differences enough to call a level $I(1)$?

Let (Z_t) be i.i.d., centered, with variance $\sigma^2 > 0$.

$$X_t = Z_t \quad \Longrightarrow \quad \nabla X_t = Z_t - Z_{t-1}.$$

The differences are stationary. Does this describe accumulating uncertainty?

No: $\mathrm{Var}(X_t) = \sigma^2$ at every time. The differenced noise has long-run variance $2\sigma^2 + 2(-\sigma^2) = 0$.

Summing the changes gives $\sum_{j=1}^{t} \nabla X_j = Z_t - Z_0$: the interior terms cancel.

In an AR(1), the question is whether $\phi = 1$

Start with $X_t = \phi X_{t-1} + Z_t$, with i.i.d. centered innovations of finite, positive variance. Under the null, take $X_0 = 0$. Subtract X_{t-1}:

$$\nabla X_t = (\phi - 1)X_{t-1} + Z_t.$$

- $H_0 : \phi = 1$: a unit root; shocks accumulate.
- $H_1 : -1 < \phi < 1$: select the stationary causal solution.

A negative fitted coefficient on X_{t-1} points toward the alternative. How negative is convincing?

Fit changes to the preceding level

Observe $x_0, \ldots, x_T$: T transitions and $T+1$ observed values. Fit $x_t - x_{t-1}$ to x_{t-1}, without an intercept.

$$\hat{\phi} - 1 = \frac{\sum_{t=1}^{T} x_{t-1}(x_t - x_{t-1})}{\sum_{t=1}^{T} x_{t-1}^2}.$$

Write $e_t = (x_t - x_{t-1}) - (\hat{\phi} - 1)x_{t-1}$ for the residual.

Under the null, the population coefficient on the preceding level is zero.

Standardize the fitted coefficient

Definition (Dickey–Fuller statistic, no intercept)

Put $\hat{s}^2 = (T-1)^{-1} \sum_{t=1}^{T} e_t^2$. The statistic is

$$D_T = \frac{(\hat{\phi} - 1)\sqrt{\sum_{t=1}^{T} x_{t-1}^2}}{\hat{s}},$$

when both $\hat{s}$ and $\sum_{t=1}^{T} x_{t-1}^2$ are positive.

It has the form "coefficient divided by its regression standard error."

Can we use the usual regression t critical value?

The regression reuses the accumulating shocks

Consider the null with $X_0 = 0$. Then $X_t = \sum_{j=1}^{t} Z_j$. The shock Z_t enters the current change and every later level.

$$x_t^2 - x_{t-1}^2 = 2x_{t-1}\nabla x_t + (\nabla x_t)^2,$$

$$2\sum_{t=1}^{T} x_{t-1}\nabla x_t = x_T^2 - x_0^2 - \sum_{t=1}^{T}(\nabla x_t)^2.$$

Conditioning on the entire regressor vector does not give the usual fixed-design normal-error regression experiment.

Simulate the unit-root null to find a cutoff

Set $X_0 = 0$, $Z_t \overset{\text{i.i.d.}}{\sim} N(0,1)$; simulate 10,000 paths with 200 transitions and compute D_T.

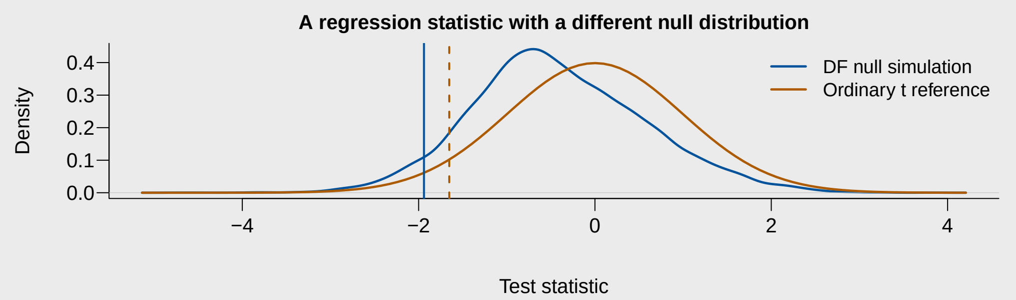

Reject for D_T below the simulated 5% quantile.

Simulated data; Gaussian random walks; calibration seed 54209.

Does the cutoff reject about 5% of true nulls?

Evaluate the fixed cutoffs on 10,000 new Gaussian random walks.

Reference	Cutoff	Reject (%)	Simulation SE
Simulated DF	-1.939	5.2	0.22
Ordinary t_{199}	-1.653	9.3	0.29

Standard errors are in percentage points, conditional on the cutoffs.

Scaling a whole path by a nonzero constant leaves D_T unchanged. The same calibration covers other positive Gaussian innovation variances.

It calibrates this null experiment. It does not establish power against a stationary process with ϕ close to one.

Simulated data; 200 transitions per path; evaluation seed 54309.


YOU CAN'T
HIDE THE
TRUTH
DOUGLAS
ROUILLARD
TRUST
NO 1

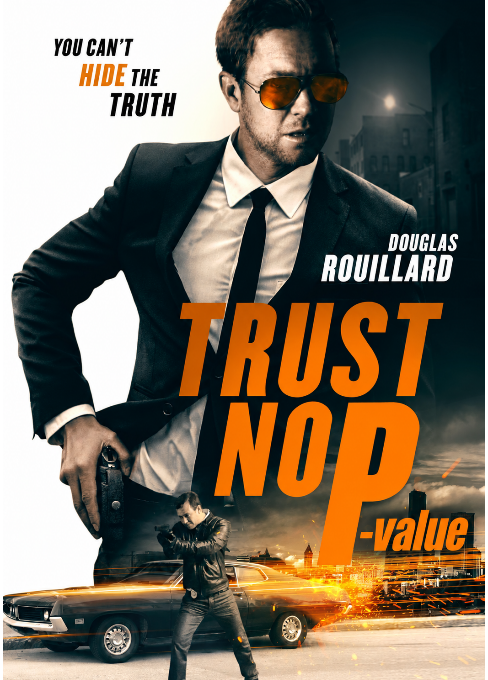
YOU CAN'T
HIDE THE
TRUTH
DOUGLAS
ROUILLARD
TRUST
NO P
-value

What if the increments are correlated?

Suppose $X_t = X_{t-1} + V_t$, where (V_t) is the stationary causal solution of $V_t = cV_{t-1} + Z_t$, $|c| < 1$. The innovations are i.i.d., centered, with finite, positive variance.

$$\nabla X_t = c\nabla X_{t-1} + Z_t.$$

The levels still have a unit root.

Adding lagged differences to the test regression accounts for short-run dependence in the increments.

Allow lagged changes and deterministic terms

Adding $a + \delta t$ to a centered $I(1)$ process gives an $I(1)$ level with drift. Unknown deterministic terms are fitted within the test:

Definition (Augmented Dickey–Fuller regression)

For a chosen lag count $p \geq 0$, fit on the available common rows

$$\nabla x_t = \alpha + \beta t + (\phi - 1)x_{t-1} + \sum_{j=1}^{p} c_j \nabla x_{t-j} + e_t.$$

Write the fitted level coefficient as $\hat{\phi} - 1$. The ADF statistic is this coefficient divided by its conventional regression standard error.

Small values count against the unit-root null, using matching ADF critical values under the workbook's innovation and increment assumptions.

The regression specification determines the reference

Terms fitted	Test regression	Alternative allows
None	$\alpha = \beta = 0$	Zero-mean stationary level
Constant	Fit α; $\beta = 0$	A constant mean
Linear trend	Fit α and β	A linear mean trend

Choose p to account for increment dependence. Estimate an unknown mean or trend within the chosen test; fitted residuals cannot be treated as observations whose mean was known to be zero.

The conventional constant-only reference uses a zero-drift unit-root null. Fitting an intercept does not make that reference valid for every drift.

Workbook: *How Do Serial Correlation and a Trend Change the Test?*

What if we start by assuming stationarity?

Suppose $X_t = \mu + U_t$, with (U_t) an $I(0)$ remainder. Should deviations from the fitted mean accumulate this much?

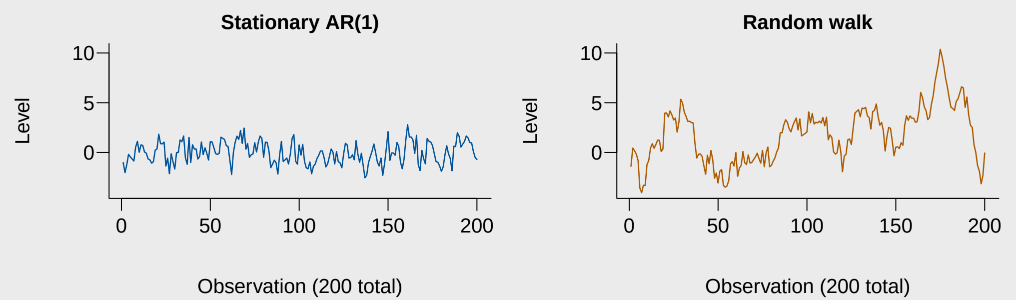

KPSS asks whether cumulative residuals are unusually large under this null.

Simulated data; 200 observations; AR coefficient 0.5; seed 54296.

The final residual sum is zero in both cases

Fit the constant by $\bar{x}_T$ and let $e_t = x_t - \bar{x}_T$.

$$S_t := \sum_{j=1}^{t} e_j = \sum_{j=1}^{t} x_j - \frac{t}{T} \sum_{j=1}^{T} x_j.$$

Why is $S_T = 0$, regardless of whether the process is stationary?

The final residual sum is zero in both cases

Fit the constant by $\bar{x}_T$ and let $e_t = x_t - \bar{x}_T$.

$$S_t := \sum_{j=1}^{t} e_j = \sum_{j=1}^{t} x_j - \frac{t}{T} \sum_{j=1}^{T} x_j.$$

Why is $S_T = 0$, regardless of whether the process is stationary?

OLS with an intercept forces the residuals to sum to zero.

The evidence is in the intermediate excursions of S_t.

How large is a cumulative fluctuation under the null?

For an $I(0)$ remainder (U_t) with long-run variance v, $\mathrm{Var}(\sum_{j=1}^{T} U_j) \approx Tv$. Estimate v from the fitted residuals using Bartlett:

$$\widehat{v}_H = \hat{\gamma}_e(0) + 2\sum_{h=1}^{H}\left(1 - \frac{h}{H+1}\right)\hat{\gamma}_e(h).$$

Here $H < T$ and $\hat{\gamma}_e(h) = T^{-1}\sum_{t=1}^{T-h} e_t e_{t+h}$.

The marginal variance alone omits the contribution from serial dependence.

Workbook: *Estimating long-run variance; Can Cumulative Residuals Challenge Stationarity?*

Compare the entire cumulative paths

Plot $S_t/\sqrt{T\widehat{v}_H}$, using $H = 8$ for each record.

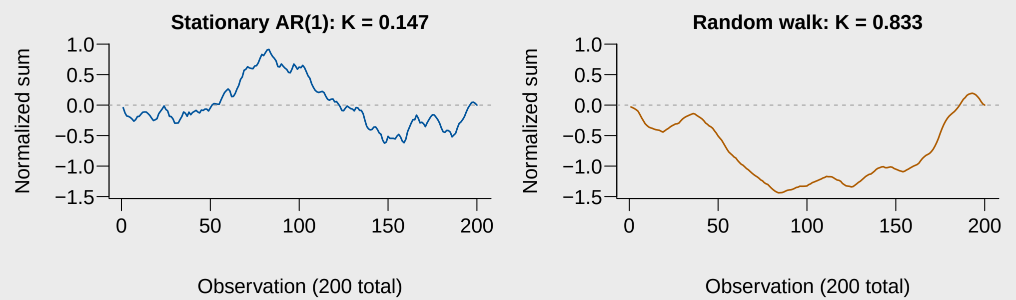

Both paths end at zero. Which has the larger average squared excursion?

Simulated data; the same 200 observations; each path uses its own Bartlett estimate.

Measure the average squared cumulative excursion

Definition (KPSS statistic)

Fit a constant, or a constant and a linear trend, and form residuals e_t. Let $S_t = \sum_{j=1}^{t} e_j$. When $\widehat{v}_H > 0$, define

$$K_T = \frac{1}{T} \sum_{t=1}^{T} \left(\frac{S_t}{\sqrt{T \widehat{v}_H}} \right)^2 = \frac{\sum_{t=1}^{T} S_t^2}{T^2 \widehat{v}_H}.$$

Large values count against an $I(0)$ remainder around the fitted deterministic part, under the KPSS reference assumptions.

For the two plotted records: $K_T = 0.147$ (AR) and $K_T = 0.833$ (walk).

KPSS rejects above its critical value

Null: an $I(0)$ remainder around...	5% critical value
A constant	0.463
A linear trend	0.146

These are supplied asymptotic references under the workbook's causal linear-process assumptions and suitable Bartlett bandwidth growth. Weak stationarity alone is insufficient.

In a finite record, compare a few plausible bandwidths: H changes the estimated null scale and can change the decision.

Reference: Kwiatkowski, Phillips, Schmidt, and Shin (1992).

The tests start from opposite nulls

ADF: a unit root. KPSS: an $I(0)$ remainder.

Reject ADF?	Reject KPSS?	Evidence within the chosen models
Yes	No	Supports a stationary remainder
No	Yes	Supports an integrated level
No	No	May have too little power to distinguish
Yes	Yes	Reconsider specification and calibration

Nonrejection does not establish the null. A stationary alternative close to a unit root can be difficult to detect.

Neither test checks every requirement of weak stationarity

Imagine a process whose marginal distribution is $N(0, 1)$ at every time, but whose lag-one covariance changes halfway through the record.

Its mean and variance are constant. Is it weakly stationary?

Neither test checks every requirement of weak stationarity

Imagine a process whose marginal distribution is $N(0,1)$ at every time, but whose lag-one covariance changes halfway through the record.

Its mean and variance are constant. Is it weakly stationary?

No. A unit-root test does not directly test that covariance change, and KPSS is not a general test of covariance stability.

We still need to compare fluctuations and lag relationships across time.

Workbook: *Unchanged marginals do not settle stationarity*

Return to the oil record with explicit choices

Analyze the first 350 log prices and their 349 first differences.

- Primary comparison: fit a constant in both tests.
- ADF: use four lagged differences ($p = 4$).
- KPSS: use Bartlett bandwidth $H = 8$.

Compare these choices with $p = 0, 8$, $H = 4, 16$, and a fitted linear trend. These examine dependence over several weeks and sensitivity to a changing mean.

Report the comparison even if some choices disagree.

The primary results favor differencing the log prices

Constant fitted; ADF $p = 4$; KPSS $H = 8$; significance level 5%.

Series	Test	Statistic	Critical value	Reject?
Log level	ADF ($<$)	-0.391	-2.870	No
	KPSS ($>$)	3.124	0.463	Yes
Log change	ADF ($<$)	-8.428	-2.870	Yes
	KPSS ($>$)	0.131	0.463	No

Which null is rejected in each row?

Real data: `astsa::oil`, EIA; first 350 weekly prices, 2000–2006.

Do these conclusions depend on the chosen tuning values?

ADF: $p \in \{0, 4, 8\}$; KPSS: $H \in \{4, 8, 16\}$. For each test, compare both constant and linear-trend specifications.

	ADF rejections	KPSS rejections
Log levels	$0/6$	$6/6$
Log changes	$6/6$	$0/6$

The conclusions agree across these choices. These are sensitivity comparisons on the same record, not independent confirmations.

Real data: `astsa::oil`, EIA; first 350 weekly prices, 2000–2006.

Do the changes look equally variable throughout?

Split the 349 weekly log changes into two consecutive portions.

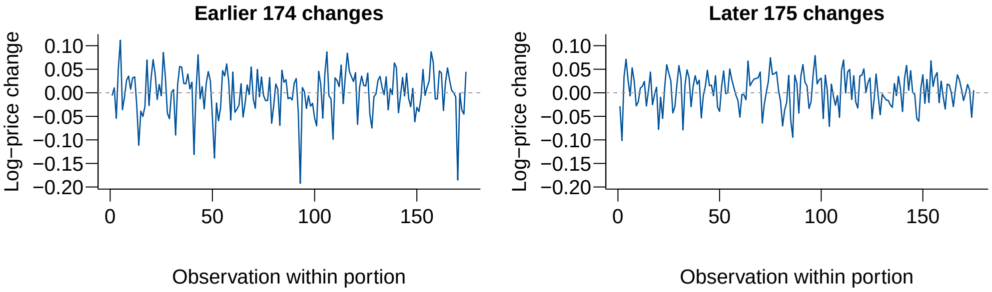

Sample standard deviations: 0.047 earlier and 0.036 later. The plots use the same vertical scale.

Real data: `astsa::oil`, EIA; first 350 weekly prices, 2000–2006.

Are the lag relationships comparable?

Each portion has its own mean and sample ACF.

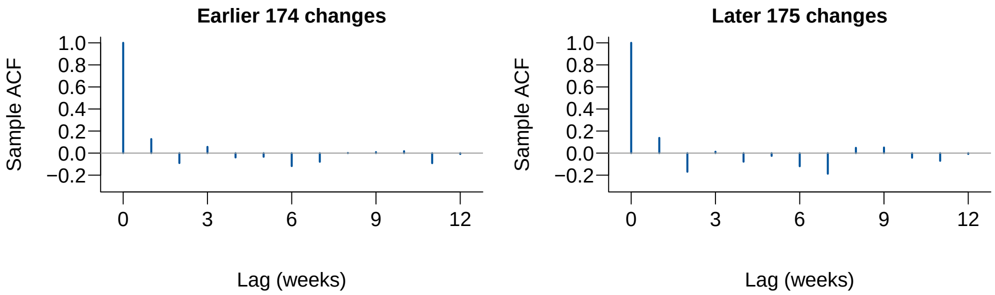

Differences between these estimates can also be sampling variation. These plots address a question the unit-root decision leaves open.

Real data: `astsa::oil`, EIA; first 350 weekly prices, 2000–2006.

An integrated log-price model is worth investigating

The test comparisons support a stationary model for log changes over a stationary model for log levels, within the stated specifications.

The changing size of the fluctuations remains a concern. Differencing and nonrejection do not settle the stationary approximation.

What if oil and gasoline log prices both accumulate shocks, but a combination of those log prices does not? That is the cointegration question.