STA 542: Introduction to Time Series Analysis

Lecture 14: Assessing fitted vector models

Lasse Vuursteen

Duke
UNIVERSITY

Wed Oct 14, 2026

What did we fit?

For a vector series $X_t \in \mathbb{R}^d$, recall the VARMA(p, q) model:

$$X_t = c + \sum_{i=1}^{p} A_i X_{t-i} + Z_t + \sum_{j=1}^{q} C_j Z_{t-j}.$$

The matrices A_i multiply observation lags; C_j multiply innovation lags.

Z_t is vector white noise with mean zero and covariance Σ_Z.

How well does the fitted model describe the data?

Fitting estimates the coefficients and reconstructs the innovations:

$$x_t = \widehat{c} + \sum_{i=1}^{p} \widehat{A}_i x_{t-i} + \widehat{z}_t + \sum_{j=1}^{q} \widehat{C}_j \widehat{z}_{t-j}.$$

Here x_t is observed; $\widehat{z}_t$ estimates the unobserved innovation Z_t.

Today we assess this fit, starting with the residuals.

Is there serial correlation left in the errors?

For weakly stationary errors, write $\Gamma_Z(h) = \text{Cov}(Z_{t+h}, Z_t)$. Choose a maximum lag H before the check.

$$H_0 : \Gamma_Z(h) = 0 \quad \text{for every } h = 1, \ldots, H,$$

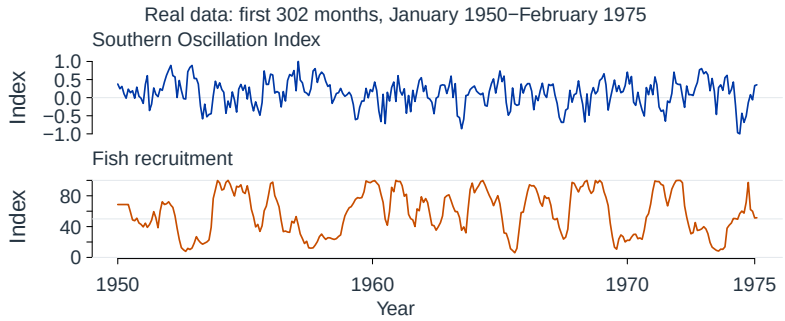
$$H_1 : \Gamma_Z(h) \neq 0 \quad \text{for at least one } h = 1, \ldots, H.$$

The zero matrix rules out both auto- and cross-correlations at these lags.

We assess these restrictions using fitted residuals.

Does climate history help predict fish recruitment?

SOI measures a tropical Pacific air-pressure contrast. Recruitment is an index of new fish entering a population.



Real data: `astsa::soi` and `astsa::rec`, Jan. 1950–Feb. 1975.

What does the fitted model leave unexplained?

Let $X_t = (S_t, R_t)^\top$ collect SOI and recruitment.

$$X_t = b_{m(t)} + \sum_{i=1}^{p} A_i X_{t-i} + Z_t,$$

$$\widehat{z}_t = x_t - \widehat{b}_{m(t)} - \sum_{i=1}^{p} \widehat{A}_i x_{t-i}.$$

- Twelve monthly intercept vectors, estimated jointly with the lag matrices.
- Candidate orders $p \in \{1, 2, 6, 12\}$ on the same $T_* = 290$ response months.

Stationary errors can still be predictable

Weak stationarity makes the mean and lag covariances constant over time. Vector white noise additionally requires

$$\mathbb{E}[Z_t] = 0, \qquad \Gamma_Z(h) := \mathrm{Cov}(Z_{t+h}, Z_t) = \begin{cases} \Sigma_Z, & h = 0, \\ 0, & h \neq 0. \end{cases}$$

Can SOI and recruitment innovations in the same month be correlated?

Stationary errors can still be predictable

Weak stationarity makes the mean and lag covariances constant over time. Vector white noise additionally requires

$$\mathbb{E}[Z_t] = 0, \qquad \Gamma_Z(h) := \mathrm{Cov}(Z_{t+h}, Z_t) = \begin{cases} \Sigma_Z, & h = 0, \\ 0, & h \neq 0. \end{cases}$$

Can SOI and recruitment innovations in the same month be correlated?

Yes. The off-diagonal entries of Σ_Z may be nonzero.

Each cross-correlation has a direction

Write the fitted residual vectors as $e_1, \ldots, e_{T_*}$.

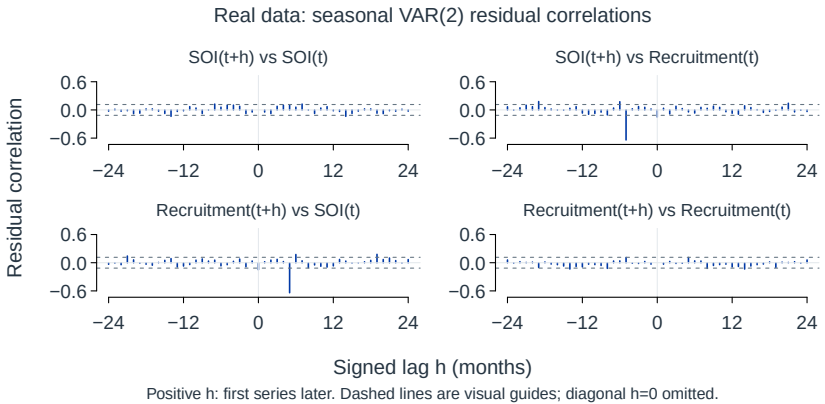
$$\widehat{\Gamma}_e(h) = \frac{1}{T_*} \sum_{t=1}^{T_*-h} e_{t+h} e_t^\top, \qquad 0 \leq h < T_*,$$

$$\widehat{R}_{ij}(h) = \frac{[\widehat{\Gamma}_e(h)]_{ij}}{\sqrt{[\widehat{\Gamma}_e(0)]_{ii}[\widehat{\Gamma}_e(0)]_{jj}}}.$$

At positive lag h: row i is later, column j is earlier.

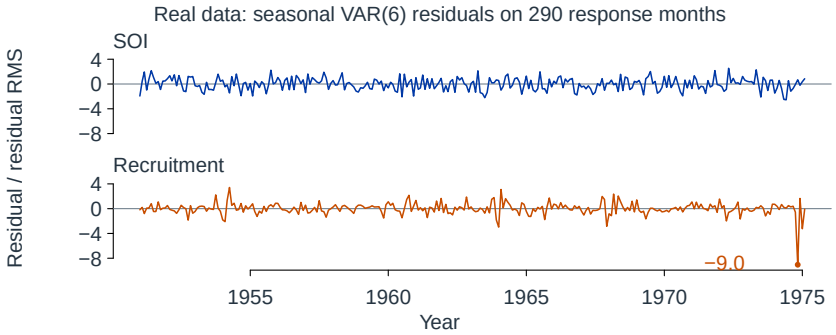
$\widehat{R}_{21}(5)$ pairs recruitment errors with SOI errors five months earlier.

What has the two-lag model missed?

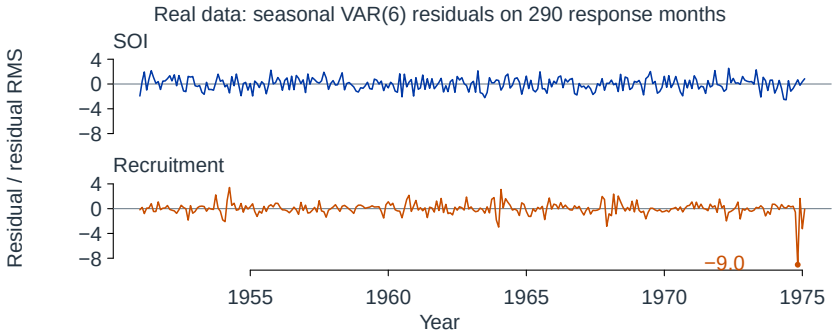


Seasonal VAR(2). Dashed lines: rough $\pm 1.96/\sqrt{290}$ guides, not a joint test.

Are Gaussian innovations plausible?



Are Gaussian innovations plausible?



One recruitment error is about nine fitted root mean squares below zero.

Seasonal VAR(6). Each coordinate is divided by its residual root mean square.

Large residual lag covariances increase the statistic

For nonsingular $\widehat{\Gamma}_e(0)$, the adjusted portmanteau statistic is

$$Q_H = T_*^2 \sum_{h=1}^{H} \frac{\operatorname{tr}\left\{\widehat{\Gamma}_e(h)^\top \widehat{\Gamma}_e(0)^{-1} \widehat{\Gamma}_e(h) \widehat{\Gamma}_e(0)^{-1}\right\}}{T_* - h}.$$

- The trace measures covariance after standardizing the innovation coordinates.
- Sum over lags $1, \ldots, H$; reject for large Q_H.

The reference distribution needs stronger assumptions

Stable VAR(p), fitted by OLS, with i.i.d. mean-zero innovations, constant positive-definite covariance, and finite fourth moments:

$$Q_H \;\approx\; \chi^2_{d^2(H-p)}.$$

Fish model	Q_{24}	Reference df	p-value
VAR(2)	246.3	88	$< 10^{-4}$
VAR(6)	63.1	72	0.763

At 5%, no remaining serial correlation is detected for VAR(6) through 24 months.

Reject when the upper-tail p-value is below .05. $p < H \ll T_*$.
Specify the order and deterministic regressors before the check.

Can residuals be dependent without autocorrelation?

For one innovation coordinate $Z_{t,j}$, the BDS test assesses

$$H_0: \quad Z_{t,j} \overset{\text{iid}}{\sim} F_j, \qquad F_j \text{ unspecified.}$$

No Gaussian law is assumed. We use the fitted residuals for this coordinate.

Diagnostic	What it checks
Residual ACF / Ljung–Box	Correlation at selected lags
Squared-residual ACF / ARCH	Predictable squared magnitudes
BDS	Broader departures from i.i.d.

BDS can detect nonlinear dependence left after fitting the mean.

R: `tseries::bds.test()` on one residual series.

Can a coefficient be set to zero?

Return to regression with fixed predictors $x_i \in \mathbb{R}^m$:

$$Y_i = x_i^\top \beta + Z_i, \qquad Z_i \overset{\text{iid}}{\sim} N(0, \sigma^2), \qquad i = 1, \ldots, n.$$

Test $H_0 : C\beta = c$ against $H_1 : C\beta \neq c$. The specified matrix C has r independent rows.

Examples: set a slope to zero, or require two slopes to be equal.

After imposing the restriction, what happens to the other estimates?

$n > m$; the design has full column rank; $\sigma^2 > 0$ is unknown.

Can a coefficient be set to zero?

Return to regression with fixed predictors $x_i \in \mathbb{R}^m$:

$$Y_i = x_i^\top \beta + Z_i, \qquad Z_i \overset{\text{iid}}{\sim} N(0, \sigma^2), \qquad i = 1, \ldots, n.$$

Test $H_0 : C\beta = c$ against $H_1 : C\beta \neq c$. The specified matrix C has r independent rows.

Examples: set a slope to zero, or require two slopes to be equal.

After imposing the restriction, what happens to the other estimates?

Refit the remaining coefficients and the error variance.

$n > m$; the design has full column rank; $\sigma^2 > 0$ is unknown.

How much likelihood do the restrictions cost?

Let $\widehat{\ell}_0, \mathrm{RSS}_0$ come from the restricted fit and $\widehat{\ell}_1, \mathrm{RSS}_1$ from the unrestricted fit. Here $\widehat{\ell}_j$ is maximized log likelihood and RSS is residual sum of squares.

$$\mathrm{LR} = 2(\widehat{\ell}_1 - \widehat{\ell}_0) = n \log\left(\frac{\mathrm{RSS}_0}{\mathrm{RSS}_1}\right) \geq 0.$$

Under H_0, with m, r fixed, $\mathrm{LR} \xrightarrow{d} \chi^2_r$ as $n \to \infty$.

An approximate 5% test rejects above the 95th percentile of χ^2_r.

This Gaussian regression also has an exact F calibration. The χ^2 reference is asymptotic.

With lagged predictors, calibration is asymptotic

A stable, causal, ergodic VAR(p) with an intercept permits consistent OLS with dependent white noise.

For likelihood-ratio inference with fixed p, d, assume i.i.d. $N_d(0, \Sigma_Z)$ innovations with $\Sigma_Z > 0$.

For r independent linear restrictions on its coefficients,

$$2(\widehat{\ell}_1 - \widehat{\ell}_0) \xrightarrow{d} \chi^2_r \qquad \text{under } H_0.$$

- Refit all free coefficients and Σ_Z on the same response dates.
- Excluding the SOI lags from the recruitment equation imposes $r = p$ zero restrictions.

Later predictor rows depend on earlier innovations.
The exact fixed-design F calibration does not carry over.

Does the better fit pay for more parameters?

A likelihood-ratio test starts from a specified restriction. How should we compare several specifications, including models that are not restrictions of one another?

$$\mathrm{AIC} = -2\widehat{\ell} + 2k,$$

$$\mathrm{BIC} = -2\widehat{\ell} + k \log T_*.$$

$\widehat{\ell}$: maximized log likelihood.
k: free estimated parameters. T_*: scored time points.

Lower scores are preferred. BIC has the larger per-parameter cost when $T_* > e^2$.

Compare the same response vectors, dates, units, and compatible likelihood conventions.

How many parameters does another lag add?

For the bivariate fish model,

$$k = \underbrace{4p}_{\text{lag matrices}} + \underbrace{24}_{\text{monthly intercepts}} + \underbrace{3}_{\text{innovation covariance}}.$$

Moving from two to six lags adds how many parameters?

How many parameters does another lag add?

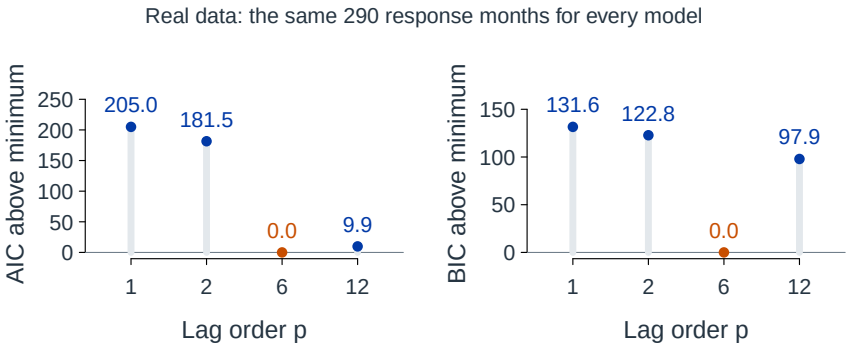
For the bivariate fish model,

$$k = \underbrace{4p}_{\text{lag matrices}} + \underbrace{24}_{\text{monthly intercepts}} + \underbrace{3}_{\text{innovation covariance}}.$$

Moving from two to six lags adds how many parameters?

Sixteen. AIC requires a log-likelihood gain above 16; BIC above $\frac{1}{2}(16)\log(290) \approx 45.4$.

Both criteria favor six lags



The improvement from six to twelve lags does not pay for 24 more coefficients.

Same 290 response vectors. Each panel subtracts its smallest score.

An innovation lag can replace an observation lag

GNP growth and unemployment changes, on the same 103 response quarters:

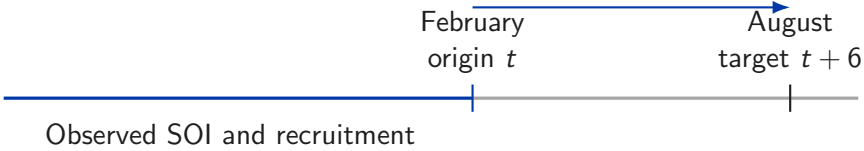
Model	k	AIC	BIC
VAR(1)	9	395.5	419.2
VAR(2)	13	384.8	419.1
VARMA(1,1)	13	363.2	397.4

VAR(2) and VARMA(1,1) have equal parameter cost here. The difference comes from fitted likelihood.

Real data: `astsa::econ5`. Identified VARMA(1,1); common initialized Gaussian scores.

What was known when the forecast was issued?

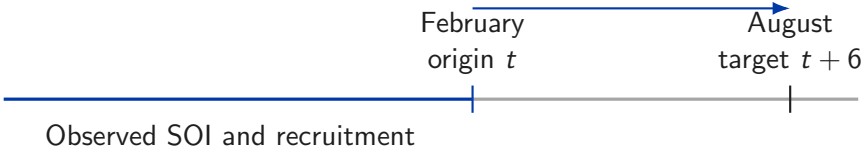
Fit through February 1975. Forecast recruitment in August 1975.



May this forecast use March's observed SOI?

What was known when the forecast was issued?

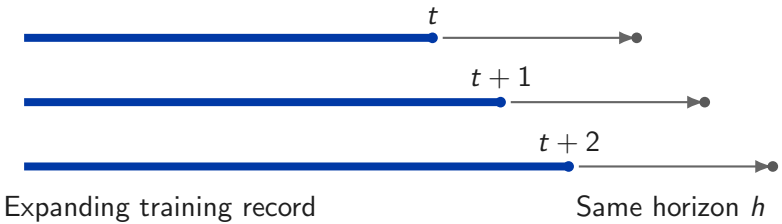
Fit through February 1975. Forecast recruitment in August 1975.



May this forecast use March's observed SOI?

No. Recursively forecast both coordinates from February's information.

Advance the origin, then refit



- Refit every coefficient at each origin.
- If AIC or BIC chooses the order, repeat that choice using only that origin's training data.

Compare errors for the same target and horizon

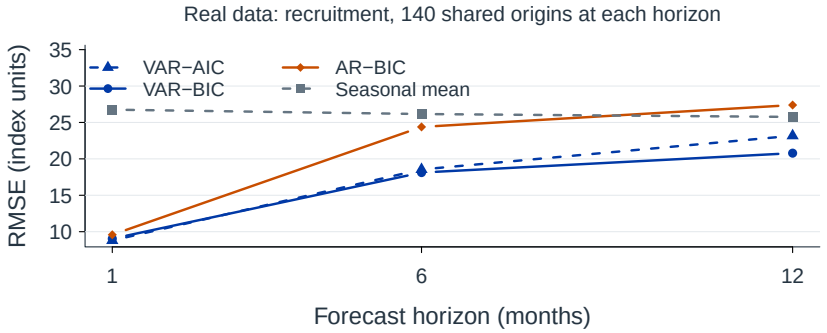
For recruitment at horizon h, average over a common set of origins $\mathcal{O}_h$:

$$\mathrm{RMSE}_R(h) = \left\{ \frac{1}{|\mathcal{O}_h|} \sum_{t \in \mathcal{O}_h} \left(r_{t+h} - \widehat{r}_{t+h|t} \right)^2 \right\}^{1/2}.$$

- The score has recruitment-index units.
- Joint VAR and recruitment-only AR forecasts can be compared on this target.

Their raw likelihoods concern different response vectors and cannot be compared by AIC/BIC.

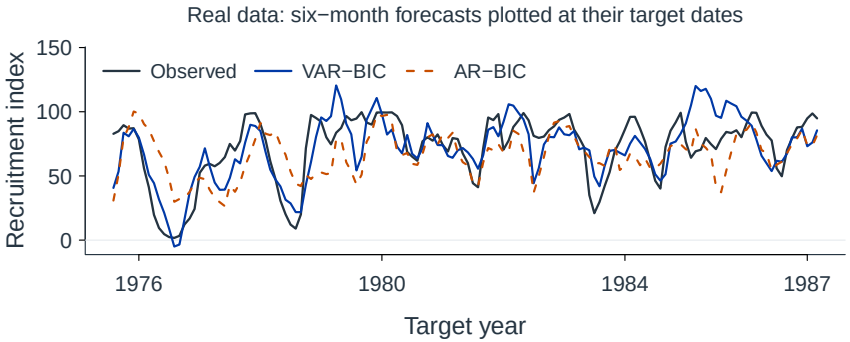
SOI improves the longer-horizon recruitment forecasts



At six months, VAR-BIC reduces RMSE by about 26% relative to AR-BIC.

140 common origins. Orders reselected at each origin.
Seasonal mean: training average for the target calendar month.

What does that gain look like in the forecasts?

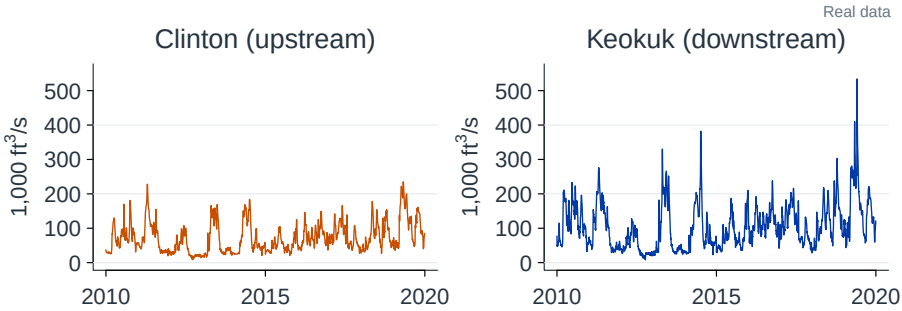


Each prediction was issued six months before its plotted target date.

Two joint forecasts are negative. They remain in the reported errors.

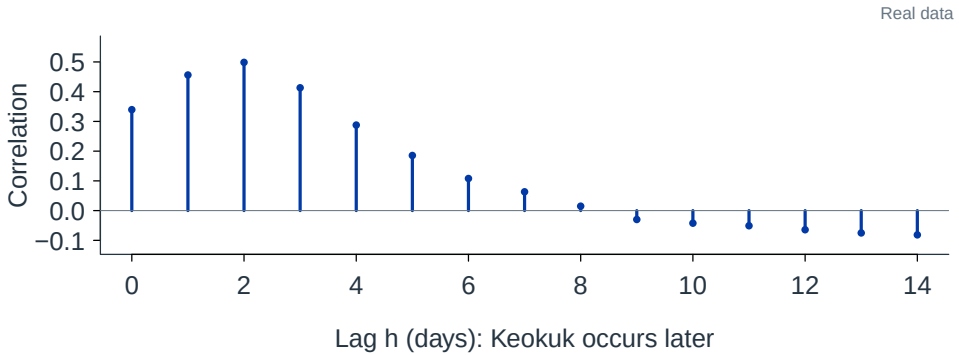
Can upstream flow improve a downstream forecast?

Daily Mississippi discharge at Clinton and Keokuk, Iowa. Clinton is about 148 river miles upstream.



Real data: USGS 05420500 and 05474500. Initial fitting period: 2010–2019.

Changes at Clinton precede changes at Keokuk

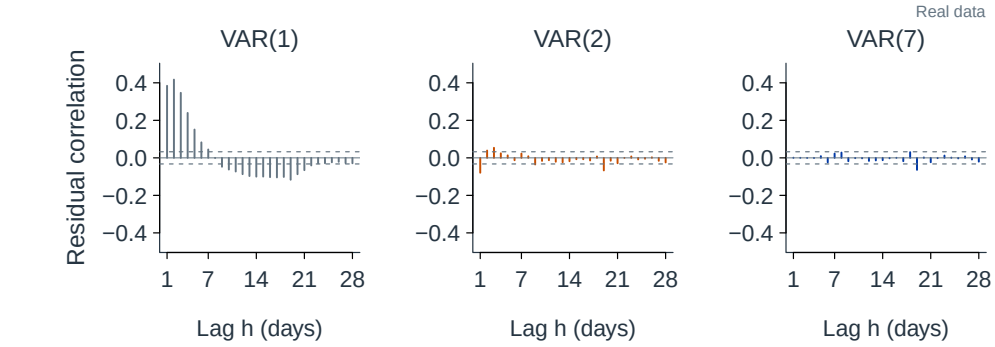


The displayed correlation peaks at two days. The forecasts concern flow levels.

Training data only. Positive h: Keokuk's change occurs h days after Clinton's.

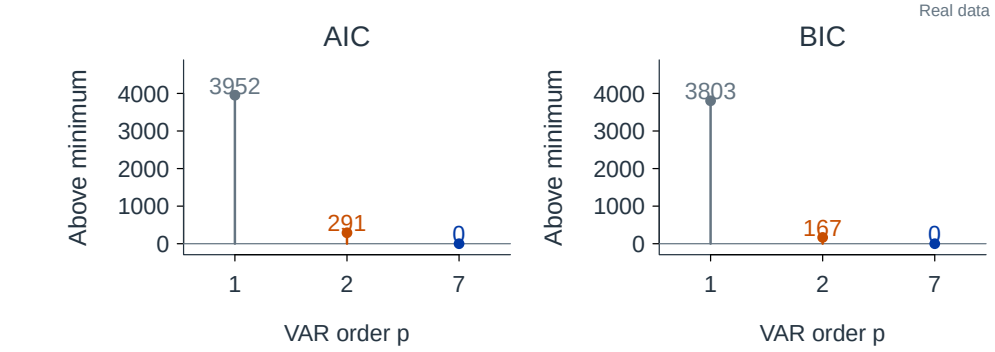
Additional lags reduce the river residual pattern

Joint flow VARs with an intercept and annual and half-year sine–cosine terms.



Later Keokuk errors versus earlier Clinton errors. All three Q_{28} checks reject.
For VAR(7): $Q_{28} = 140.3$, df 84, $p \approx 0.000114$.

Both criteria select seven river lags



Twenty more lag coefficients than VAR(2). Does the forecast gain justify them?

$k = 4p + 10 + 3$. Common initial response window: $T_* = 3{,}645$.

The same questions apply to a cointegrated model

For the price–rent model, examine:

- The fitted innovations after removing the equilibrium gap and lagged changes.
- The parameter cost of the chosen lag specification.
- Forecasts made before their target observations arrive.

The equilibrium gap and the fitted innovations are different residuals.

Is there a rent–home price equilibrium?



Ben Felix, 2019



Humphrey Yang, 2023

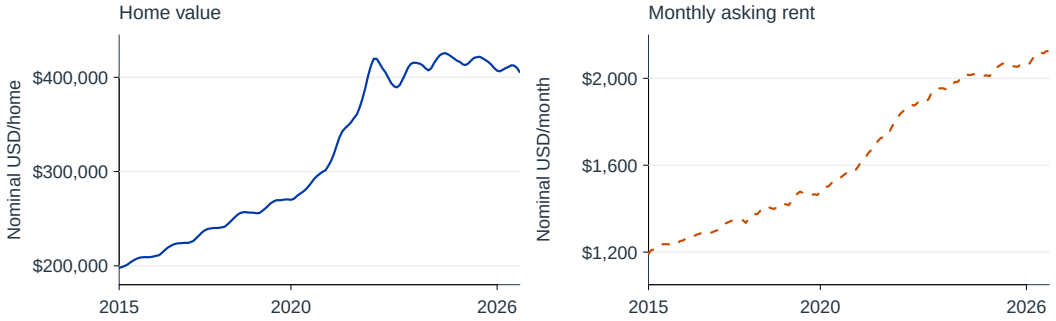
Renting Is $1,066 Cheaper Per Month Than Buying and Investing It Pays Off (August Rent Report)

Zillow Research, September 2026

Key question: Do rents follow house prices, do house prices follow rents, or do both adjust?

Do rents catch up with house prices?

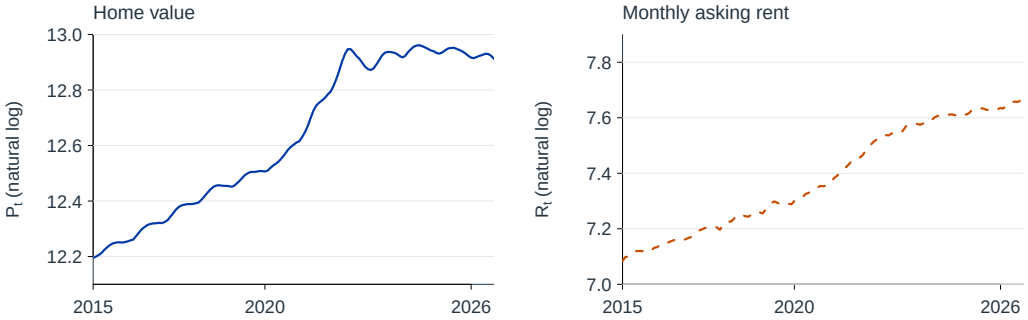
Durham metro: monthly single-family home values and asking rents.



Real data: Zillow ZHVI and ZORI, Jan. 2015–Sept. 2026, $T = 141$.

The same series after taking logs

$P_t = \log(\text{home value}_t)$ and $R_t = \log(\text{monthly rent}_t)$.

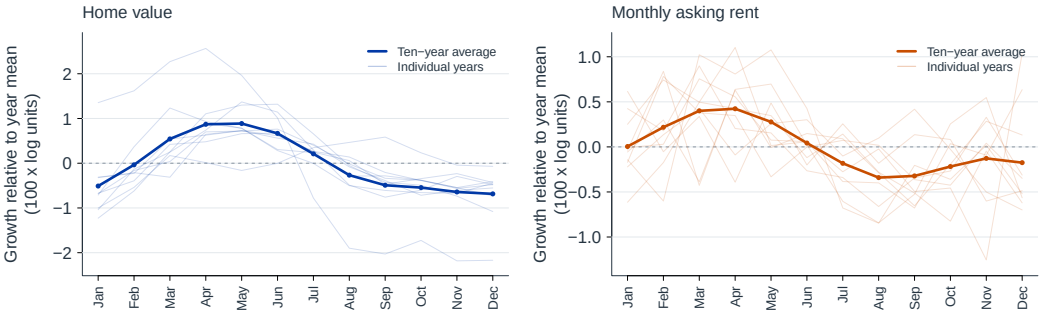


Do these log series also have a seasonal pattern?

Natural logs of the same nominal dollar series. Zillow ZHVI and ZORI, Jan. 2015–Sept. 2026, $T = 141$.

Seasonal patterns in monthly growth

Each year's curve shows monthly log growth minus that year's average.



Faint lines: individual years. Bold line: average across the ten complete years, 2016–2025.
Growth is $100\nabla P_t$ or $100\nabla R_t$, approximately monthly percentage growth.

Cointegration Evidence

For each log level, $X_t = a + bt + \delta_{m(t)} + u_t$:
KPSS null: u_t is $I(0)$.
ADF null: u_t has a unit root.

Test	P_t	R_t	5% rejection
KPSS, $L = 4$	0.344	0.313	Both
KPSS, $L = 13$	0.144	0.127	Neither
ADF (BIC)	-0.577	-1.333	Neither

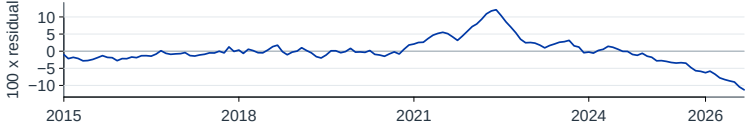
OLS with monthly effects:

$$\widehat{U}_t = P_t - 1.404 R_t - 2.264 - \widehat{\delta}_{m(t)}.$$

Residual KPSS ($L = 4$): 0.466.

Conditional bootstrap: $p = 0.595$.

Assume cointegration for the fit.



$m(t)$: calendar month.
Table 5% cutoffs: KPSS > 0.146, ADF < -3.43. L: bandwidth. ADF lags: price 2, rent 5 based off BIC, 0–12.

Fitting the Cointegrated VARMA Model

For $V_t = (P_t, R_t)^\top$, a rank-one cointegrated VARMA(p, q) is

$$\nabla V_t = \alpha(\beta^\top V_{t-1} + \rho) + \sum_{k=1}^{p-1} B_k \nabla V_{t-k} + \boldsymbol{\delta}_{m(t)} + Z_t + \sum_{j=1}^{q} C_j Z_{t-j},$$

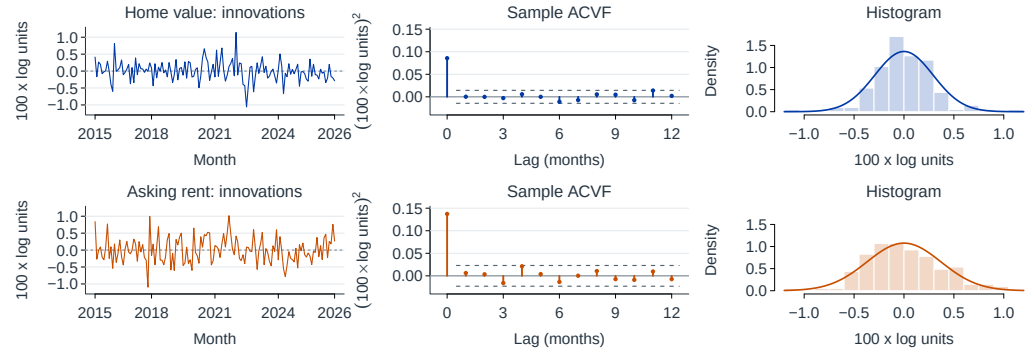
with monthly effects $\boldsymbol{\delta}_{m(t)} \in \mathbb{R}^2$ and $Z_t \overset{\text{iid}}{\sim} N_2(0, \Sigma)$.
Joint Cointegrated VARMA(4,0) fit selected using AIC:

$$\nabla V_t = \begin{pmatrix} 0.0011 \\ 0.0321 \end{pmatrix} \underbrace{(P_{t-1} - 1.385 R_{t-1} - 2.279)}_{\text{refitted level relation}} + \sum_{k=1}^{3} \widehat{B}_k \nabla V_{t-k} + \widehat{\boldsymbol{\delta}}_{m(t)} + \widehat{Z}_t.$$

Separate LR tests: $\alpha_P = 0$ (p-val 0.779), $\alpha_R = 0$ (p-val 1.16×10^{-9}), asymptotically valid under assumed iid normal residuals.

What remains after fitting the dynamics?

These are fitted innovations $\widehat{Z}_t$, after removing the gap, lagged changes, and monthly effects.



137 conditional innovations, May 2015–Sept. 2026. Histograms: fitted Gaussian density overlaid.
Dashed ACVF lines: approximate 95% pointwise iid reference for $h > 0$; ignores parameter estimation.

What does a 10% price–rent gap imply?

Let $H_t = e^{P_t}$ be home value and $M_t = e^{R_t}$ monthly rent, both in dollars.

$$H_t^\star = e^{2.279} M_t^{1.385}, \qquad g_t = \log\left(\frac{H_t}{H_t^\star}\right).$$

Here $H_t^\star$ is the home value on the fitted relation.

Along that relation, 1% higher rent corresponds to about 1.385% higher home value.

If last month's home value is 10% above its benchmark, $g_{t-1} = \log(1.10) \approx 0.0953$.

Series	Loading $\widehat{\alpha}_j$	Monthly growth contribution $100\widehat{\alpha}_j g_{t-1}$
Home value	0.0011	+0.010 percentage points
Rent	0.0321	+0.306 percentage points

Approximate percentage-point contributions; lagged changes and month held fixed.
The price loading is compatible with zero ($p = 0.779$); adjustment is mainly through rent growth.

Bonus slides: a joint residual check

Using the fitted innovations $e_t = \widehat{Z}_t$, test all auto- and cross-covariances:

$$H_0 : \quad \Gamma_Z(h) = 0, \qquad h = 1, \ldots, H.$$

Workbook 14's adjusted portmanteau statistic is

$$Q_H = T_*^2 \sum_{h=1}^{H} \frac{1}{T_* - h} \operatorname{tr}\left\{ \widehat{\Gamma}_e(h)^\top \widehat{\Gamma}_e(0)^{-1} \widehat{\Gamma}_e(h) \widehat{\Gamma}_e(0)^{-1} \right\}.$$

Maximum lag H	Q_H	Conditional bootstrap p
12 months	53.20	0.114
24 months	108.29	0.116

Neither check rejects at 5% under the fitted Gaussian VECM.

$T_* = 137$; 999 Gaussian simulations, with the full VECM refitted each time. Monte Carlo SE: about 0.010. Rank one, four level lags, monthly specification, and initial four observations held fixed.

Bonus slides: why can adjustment be slow?

Write the fitted model in levels and stack $S_t = (V_t^\top, V_{t-1}^\top, V_{t-2}^\top, V_{t-3}^\top)^\top$.

$$\widehat{F} = \begin{pmatrix} \widehat{A}_1 & \widehat{A}_2 & \widehat{A}_3 & \widehat{A}_4 \\ I_2 & 0 & 0 & 0 \\ 0 & I_2 & 0 & 0 \\ 0 & 0 & I_2 & 0 \end{pmatrix}.$$

One unit eigenvalue gives a permanent component.

The slowest decaying mode has

$$r = 0.993217, \qquad r^h \approx \tfrac{1}{2} \text{ at } h = 101.8.$$

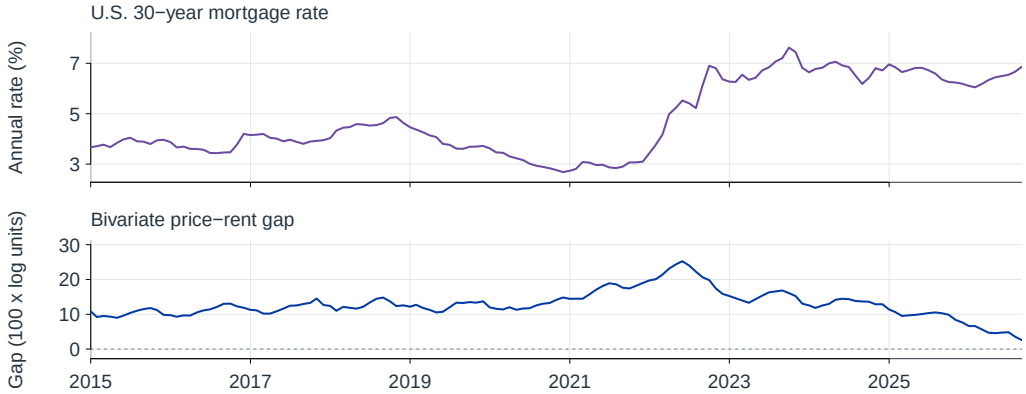
About 8.49 years to halve its amplitude.

$\widehat{A}_1 = I_2 + \widehat{\alpha}\widehat{\beta}^\top + \widehat{B}_1$, $\widehat{A}_2 = \widehat{B}_2 - \widehat{B}_1$, $\widehat{A}_3 = \widehat{B}_3 - \widehat{B}_2$, $\widehat{A}_4 = -\widehat{B}_3$.

Half-amplitude time describes one fitted mode; price/rent paths combine modes.

Bonus slides: rates and the price–rent gap

The original price–rent gap is $g_t = P_t - 1.385R_t - 2.279$.



Monthly data, Jan. 2015–Sept. 2026. Rate: Freddie Mac/FRED. Gap: original two-series VECM fit.