

Multivariate time series: series that move together

Sales, prices and volumes, or GDP, consumption and investment, do not move independently. Each responds to its own past and to the others'. A vector autoregression (VAR) models that web in a few formulas: which series lead, how a shock spreads, and where they are all heading, with honest bands around the forecast.

Functions VAR: [fx.RiskVARSelect](#) · [VARFit](#) · [VARSummary](#) · [VARStability](#) · [Granger](#) · [VARIRF](#) · [VARFEVD](#) · [VARForecast](#) · [VARSim](#) · [VARBacktest](#) · [VARPortmanteau](#) · [CCF](#). Cointegration: [Johansen](#) · [VECMFit](#) · [VECMToVAR](#)

In XL Edge XL Edge tab > Productivity Tools > Stat Tools > Insert Time Series > Multivariate - VAR and cointegration

Example *XL Edge Example - Multivariate Time Series.xlsx*: US real GDP, consumption and investment, 1959-2009 (public domain, FRED). Every number in this guide is read from it.

1 What a VAR is

2 The workflow at a glance

3 Using the ribbon

4 Levels or growth rates?

5 How many lags?

6 Fit the VAR and read it

7 Is the model sound?

8 How shocks travel: IRF and FEVD

9 Forecasts, paths and backtests

10 Leads and lags: cross-correlation

11 Cointegration and the VECM

12 Pitfalls and good practice

13 Function reference

1 What a VAR is

A VAR with p lags explains each series by the last p values of **every** series, plus a constant:

$$\begin{aligned} \text{gdp}(t) &= c_1 + a_{11} \text{gdp}(t-1) + a_{12} \text{cons}(t-1) + a_{13} \text{inv}(t-1) + \dots + e_1(t) \\ \text{cons}(t) &= c_2 + a_{21} \text{gdp}(t-1) + a_{22} \text{cons}(t-1) + a_{23} \text{inv}(t-1) + \dots + e_2(t) \\ \text{inv}(t) &= c_3 + a_{31} \text{gdp}(t-1) + a_{32} \text{cons}(t-1) + a_{33} \text{inv}(t-1) + \dots + e_3(t) \end{aligned}$$

Each equation is an ordinary regression on the same lagged values, so it is fitted by least squares one equation at a time. The shocks $e(t)$ are correlated across equations; that correlation, the residual covariance Σ , is part of the model. From the fitted coefficients come the questions people actually ask: *Does investment help predict GDP?* (Granger causality) · *If consumption jumps, what happens to GDP over the next two years?* (impulse responses) · *Where are all three heading, and how uncertain is that?* (forecasts and simulated paths).

2 The workflow at a glance

Step	Question	Function	Example sheet
1	Do the series wander (levels) or return to a mean (growth rates)?	<code>fx.RiskUnitRootλ</code>	1 Stationarity
1b	If they wander, are they tied together in the long run?	<code>fx.RiskJohansenλ</code>	8 Cointegration
2	How many lags?	<code>fx.RiskVARSelectλ</code>	2 Lags
3	Fit, then read the coefficients	<code>fx.RiskVARFitλ</code> , <code>fx.RiskVARSummaryλ</code>	3 Fit
4	Is the model stable, and is anything left in the residuals? Who leads whom?	<code>fx.RiskVARStabilityλ</code> , <code>fx.RiskVARPortmanteauλ</code> , <code>fx.RiskGrangerλ</code>	4 Diagnostics
5	How does a shock spread?	<code>fx.RiskVARIRFλ</code> , <code>fx.RiskVARFEVDλ</code>	5 Impulse responses
6	Forecast, simulate, backtest	<code>fx.RiskVARForecastλ</code> , <code>fx.RiskVARSimλ</code> , <code>fx.RiskVARBacktestλ</code>	6 Forecast

Everything after step 3 reads the **Model block** that `fx.RiskVARFitλ` writes, so you fit once and point every later function at that block. Change the data, and the whole chain recalculates.

3 Using the ribbon

XL Edge tab > Productivity Tools > Stat Tools (arrow) > Insert Time Series >
Multivariate - VAR and cointegration

Menu items	XL Edge asks for
Select VAR Lags, Fit VAR (Model Block)	The data table: one column per series, oldest row first, names in the header row above. Then the options (maximum lags, or lags; deterministic terms).
VAR Coefficient Summary, VAR Residuals, Granger Causality, Residual Portmanteau Test, VAR Forecast, VAR Simulated Paths	The data table and the Model block (select its top-left cell, or the whole block), then any options.
VAR Stability, Impulse Responses, Variance Decomposition	The Model block only.
VAR Backtest	The data table, the holdout length, the lags and the series to score.
Johansen Cointegration Test, Fit VECM, VECM Forecast	The table of <i>levels</i> ; the deterministic terms and lagged differences; for the VECM, the rank.
Cross-Correlation (CCF)	Two series and the number of leads and lags.

As everywhere in XL Edge, each item writes one live formula after telling you its size, and installs the functions it needs if the workbook does not have them yet. *VAR Simulated Paths* output feeds the existing *Fan Chart (Time Series)*.

4 Step 1. Levels or growth rates? `fx.RiskUnitRootA`

A VAR needs **stationary** series: series that keep returning to an average level rather than wandering off. Most levels (GDP, prices, sales) wander; their growth rates usually do not. Test both:

```
around a trend
=fx.RiskUnitRootA(LN(Data!B5:B207)*100, "ct", , , "log real GDP")
=fx.RiskUnitRootA(Data!F6:F207, , , , "GDP growth")
```

Test (what it assumes)	log GDP: p	verdict	GDP growth: p	verdict
ADF (unit root until rejected)	0.389	Unit root	< 0.001	Stationary
Phillips-Perron (same)	0.767	Unit root	< 0.001	Stationary
KPSS (stationary until rejected)	0.013	Not stationary	>0.10	Stationary

All three tests agree: log GDP has a unit root, and its growth rate is stationary. So the VAR uses **growth rates**: $100 \times$ the change in LN, about the percentage change per quarter (the example computes them on the Data sheet). If several series wander *together*, differencing throws away their long-run tie; section 11 covers that case.

5 Step 2. How many lags? `fx.RiskVARSelectA`

```
=fx.RiskVARSelectA(Data!F6:H207, 8)
```

Fits 1 to 8 lags, all on the same quarters so the comparison is fair, and reports four information criteria per order. Each trades fit against the number of coefficients, and the order with the lowest value wins (the last row lists each criterion's choice). AIC and FPE tend to choose longer lags, SC (Schwarz, also called BIC) shorter ones, and HQ sits between. Here all four agree: **1 lag** (AIC 1, HQ 1, SC 1, FPE 1). The fit on sheet 3 takes the AIC's choice straight from this block, so if the data change, the lag order follows.

6 Step 3. Fit the VAR and read it: `fx.RiskVARFitλ`, `fx.RiskVARSummaryλ`

```
data, lags, Type (const by default), Season, Exog, Names
=fx.RiskVARFitλ(Data!F6:H207, '2 Lags'!B16, , , , Data!F4:H4)
=fx.RiskVARSummaryλ(Data!F6:H207, A6#)
```

The Model block ((15 + m + K) rows × (K + 1), where K is the number of series and m the terms per equation) holds the specification (Variables, Lags, Type, Season, Exogenous, Observations, Sample), the fit statistics (LogLik, AIC, HQ, SC, FPE), the coefficients (one column per equation, rows `gdp.l1 ... inv.lp, const`) and the residual covariance Σ . It is plain cells: you can read it, chart it, or type one in by hand with coefficients from elsewhere.

The summary gives each equation's coefficients with standard errors, t-statistics and p-values, its R^2 , and the residual covariance and correlation matrices. Here it fits 201 quarters with 1 lag. The R^2 values are 0.229 (GDP), 0.118 (consumption) and 0.265 (investment). Growth rates are hard to predict, so modest R^2 is normal.

Reading the GDP equation. Last quarter's consumption growth carries a coefficient of 0.746 ($p < 0.001$): strong consumption growth tends to be followed by strong GDP growth. Investment growth's coefficient is 0.058 (p 0.023). Investment swings far more than GDP, so a small coefficient still matters. The residual correlations matter just as much: GDP's surprises correlate 0.759 with investment's and 0.607 with consumption's. They happen together, which is why the next steps have to decide which came first.

7 Step 4. Is the model sound?

Stability: `fx.RiskVARStabilityλ`

A VAR is stable when all the roots (eigenvalues of its companion matrix) lie inside the unit circle. Then shocks die out and forecasts settle at the series' averages. Here the verdict is **Stable**, with largest root modulus 0.459. A modulus at or above 1 means non-stationary data (difference them, or use a VECM) or too many lags for too little data.

Granger causality: `fx.RiskGrangerλ`

```
does gdp help predict cons and inv?
=fx.RiskGrangerλ(Data!F6:H207, '3 Fit'!A6#, "gdp")
=fx.RiskGrangerλ(Data!F6:H207, '3 Fit'!A6#, "inv")
```

"X Granger-causes Y" means X's *past* improves forecasts of Y beyond Y's own past and the other series. That is a statement about prediction, not about cause. The output gives the F and Wald χ^2 tests, and a test of **instantaneous** causality (correlated surprises in the same quarter):

Cause → the others	F	p	Verdict (5%)	Instantaneous p
gdp → cons, inv	3.22	0.041	Granger-causes	< 0.001
inv → gdp, cons	2.88	0.057	No Granger causality	< 0.001

GDP's past helps predict the other two at the 5% level; investment's falls just short. Both are strongly linked *within* the quarter, as the residual correlations showed.

Anything left in the residuals? `fx.RiskVARPortmanteau`

The multivariate portmanteau test asks whether any autocorrelation or cross-correlation is left in the residuals up to a chosen lag (10 here). With 1 lag the adjusted statistic is 110.9 (p 0.015):


Serial correlation. With two lags it is still significant (p 0.023). The information criteria prefer the simpler model, but some structure remains: quarterly macro data carry features (the 1970s and 2008 shocks, changing volatility) that no short VAR captures fully. The fit is still usable for the analyses below, but read their bands as approximate, and test longer lags before relying on fine detail. For a single series, `fx.RiskACF` on a column of `fx.RiskVARResids` shows where the remaining correlation sits.

8 Step 5. How shocks travel: `fx.RiskVARIRF`, `fx.RiskVARFEVD`

```
all 9 shock → response pairs, quarters 0 - 12
=fx.RiskVARIRF('3 Fit'!A6#, 12)
what explains GDP's forecast errors
=fx.RiskVARFEVD('3 Fit'!A6#, 12, "gdp")
```

An **impulse response** traces how each series responds, quarter by quarter, to a one-standard-deviation shock in one series. Because the shocks are correlated, they are first made independent with a Cholesky factor of Σ , in an **order**: series earlier in the order may move later ones within the same quarter, not the reverse. The default order is the column order (gdp, cons, inv). Change it with the `Order` argument, and check whether your conclusions survive.

Response of GDP growth to a one-s.d. shock



Responses of GDP growth (percentage points per quarter) to a one-s.d. shock in each series. Sheet 5 *Impulse responses*.

Reading it. A GDP shock raises GDP growth by 0.773 points on impact and fades within about a year. In the same quarter it raises investment growth by 3.062 points: investment swings several times as hard as GDP. A consumption or investment shock cannot move GDP on impact, because both come after GDP in the order. Their effects show from the next quarter (investment: 0.120 points). With a different order, those impact effects would change. The **Cumulative** option adds the responses up, which for growth rates gives the effect on the *level*.


The **variance decomposition** asks what share of a series' forecast uncertainty each shock explains. Three years out, GDP's own shocks explain 84.9% of its forecast-error variance, consumption's 12.7% and investment's 2.5%. Each row sums to 100%, and the shares depend on the same Cholesky order.

9 Step 6. Forecasts, paths and backtests

```
8 quarters, every series, P10 / P90 band; labels for the future quarters
=fx.RiskVARForecastλ(Data!F6:H207, '3 Fit'!A6#, 8, , , , , TRANSPOSE(A5:H5))
10,000 simulated 8-quarter paths of GDP growth (correlated shocks in all three series)
=fx.RiskVARSimλ(Data!F6:H207, '3 Fit'!A6#, 8, "gdp", , 10000)
```

The forecast table gives, for every series, the point forecast and a P10-P90 band from the model's forecast-error variance (change **Lower** and **Upper** for other bands). For 2009 Q4 onward it forecasts GDP growth of 0.79% for the first quarter, between -0.21% and 1.78% with 80% probability, settling toward 0.76% by the eighth. The band is wide because quarterly GDP growth is noisy, and the model knows it.

GDP growth, % per quarter



XL Edge's fan chart of 10,000 simulated paths: the P10-P90 band, the median (solid), and the point forecast (dashed, passed to `fx.RiskChartFanλ` as its Plan). Sheet 6 Forecast.

`fx.RiskVARSimλ` draws correlated shocks for **all** the series and lets them feed through the model. Call it once per series with the same VarID and Seed and the paths belong together: row r of each is the same simulated future. That makes sums, spreads and ratios across series valid, and lets the paths drive a Monte Carlo model. Its percentiles reproduce the forecast band (the median's first quarter is 0.78%).

Backtest. `fx.RiskVARBacktestλ` refits without the last h quarters, forecasts them, and scores the result. Holding out the last 8 quarters, which here are 2007 Q4 to 2009 Q3, the deep recession, GDP growth was missed by an RMSE of 1.37 points (MAE 1.07), and only 5 of the 8 actual quarters fell inside the 80% band. No model of normal times foresees a financial crisis. Backtest over calmer windows too before judging a model.

10 Leads and lags: `fx.RiskCCFA`

```
=fx.RiskCCFA(Data!G6:G207, Data!H6:H207, 6, , {"cons","inv"})
```

The cross-correlation at lag k is the correlation of $X(t + k)$ with $Y(t)$, following R's convention: **a spike at negative k means X leads Y** . Consumption and investment growth correlate 0.278 in the same quarter, but **0.489** at $k = -1$, consumption one quarter ahead of investment, against 0.239 at $k = +1$. Anything beyond ± 0.138 is significant at 5%. Consumption tends to lead investment by a quarter. Trends and strong autocorrelation inflate cross-correlations, so use growth rates, not levels.

11 Cointegration and the VECM

Log consumption and log investment both trend upward with unit roots, yet their gap does not drift away: when investment falls far behind consumption it catches up. Series like this are **cointegrated**. A VAR in growth rates ignores the gap; a **vector error-correction model (VECM)** keeps it:

$$\Delta y(t) = \alpha \beta' y(t-1) + \Gamma_1 \Delta y(t-1) + \dots + \text{constant} + e(t)$$

$\beta'y$ is the long-run relation (the gap), and α says how hard each series is pulled back when the gap opens.

How many ties? `fx.RiskJohansenλ`

```
levels; constant, 1 lagged difference (defaults)
=fx.RiskJohansenλ(Data!J5:K207)
```

Hypothesis	Trace	95% critical	Max-eigen	95% critical
$r \leq 0$ (no tie)	18.42	15.49	14.99	14.26
$r \leq 1$	3.43	3.84	3.43	3.84

Read down and stop at the first row whose statistic is below its critical value. "No tie" is rejected (statistic above critical); "at most one" is not. So the **rank is 1**, by both tests, as the block's title says. Rank 0 would mean: difference and use a VAR. Rank equal to the number of series would mean the levels are already stationary.

Fit and forecast: `fx.RiskVECMFitλ`, `fx.RiskVECMToVARλ`

```
rank 1, 1 lagged difference
=fx.RiskVECMFitλ(Data!J5:K207, 1, 1, , Data!J4:K4)
forecast the levels
=fx.RiskVARForecastλ(Data!J5:K207, fx.RiskVECMToVARλ(A13#), 8)
```

The VECM block holds α (one column per relation), β (normalized so its first entry is 1), Γ and Σ . Here β gives the relation **lcons** $-0.830 \times \text{linv}$, the long-run gap. α is -0.0008 for consumption, essentially zero (consumption does not adjust), and **0.1469 for investment**: each quarter investment closes about 15% of the gap. `fx.RiskVECMToVARλ` rewrites the VECM as the equivalent VAR in levels, so every VAR function, forecast, paths, impulse responses and stability, works on it unchanged. From the 2009 Q3 trough, the VECM forecasts investment up 5.1% in the next quarter (log level 730.41 $\rightarrow$ 735.53) against 0.8% for consumption, as the gap the recession opened starts to close.

DETERMINISTIC TERMS

Johansen's `DetOrder`: -1 none, 0 constant (default), 1 linear trend. The VECM's `Deterministic`: "co" (constant outside the relation, default), "ci" (inside), "lo" / "li" (trend outside/inside), "colo", "cili" or "n". The results follow statsmodels (coint_johansen, VECM), with critical values from MacKinnon, Haug and Michelis (1999) for up to 12 series.

12 Pitfalls and good practice

- **Make the data stationary first**, or model the levels with a VECM if they are cointegrated. A VAR on wandering levels gives spurious coefficients and nonsense impulse responses.
- **Keep K and p small**. Each equation has $K \times p + 1$ coefficients; 6 series with 4 lags is 25 per equation. With a few hundred observations, three or four series and a handful of lags is typical.
- **Check the residuals** (portmanteau, ACF, normality with `fx.RiskNormalityλ` on `fx.RiskVARResidsλ`) and stability before believing the dynamics.
- **The Cholesky order is an assumption**. Put the slowest-reacting series first, and report results under alternative orders.
- **Granger causality is about prediction**, not cause, and only among the series in the model.
- **The bands leave out parameter uncertainty**, as vars and statsmodels do by default, so they are a little too narrow, more so for short samples.
- **Seasonal data**: add `Season` (e.g. 4 or 12) for centered seasonal dummies, or use seasonally adjusted series (`fx.RiskDecomposeλ`'s Adjusted column). **Outside drivers** (a policy rate, oil prices): pass them as `Exog` and give their future values to the forecast's `FutureExog`.

13 Function reference

Data is $n \times K$ (one column per series, oldest first); Model is the block from `fx.RiskVARFitλ` (or `fx.RiskVECMToVARλ`). Optional arguments in brackets; omit them with an empty argument, e.g. `(Data, 2, , , , Names)`.

Function	Output	Arguments and defaults
<code>VARSelectλ</code>	$(\text{MaxLags} + 3) \times 5$	<code>(Data, [MaxLags]=10, [Type]="const", [Season], [Exog])</code>
<code>VARFitλ</code>	$(15 + m + K) \times (K + 1)$	<code>(Data, [Lags]=1, [Type]="const" "trend" "both" "none", [Season], [Exog], [Names])</code>
<code>VARSummaryλ</code>	per equation + Σ	<code>(Data, Model, [Exog])</code>
<code>VARResidsλ</code>	$(T + 1) \times K$	<code>(Data, Model, [Exog])</code>
<code>VARStabilityλ</code>	$(K p + 2) \times 4$	<code>(Model)</code>
<code>Grangerλ</code>	5×6	<code>(Data, Model, Cause, [Effect]=all others, [Alpha]=0.05, [Exog])</code>
<code>VARIRFλ</code>	$(H + 3) \times (1 + K^2)$	<code>(Model, [Horizon]=10, [Impulse], [Orthogonal]=TRUE, [Cumulative]=FALSE, [Order])</code>
<code>VARFEVDλ</code>	$(K H + 2) \times (K + 2)$	<code>(Model, [Horizon]=10, [Response], [Order])</code>
<code>VARForecastλ</code>	$(H + 1) \times (1 + 3K)$	<code>(Data, Model, Horizon, [Variable], [Lower]=0.1, [Upper]=0.9, [FutureExog], [Labels])</code>
<code>VARSimλ</code>	Trials $\times$ Periods	<code>(Data, Model, Periods, [Variable]=1, [DegFree], [Trials]=10000, [VarID]=1, [Seed]=0, [FutureExog])</code>
<code>VARBacktestλ</code>	$(\text{Holdout} + 6) \times 6$	<code>(Data, Holdout, [Lags]=1, [Type], [Variable]=1, [Lower], [Upper], [Season], [Names])</code>
<code>VARPortmanteauλ</code>	4×6	<code>(Data, Model, [Lags]=10, [Alpha]=0.05, [Exog])</code>
<code>CCFλ</code>	$(2L + 3) \times 5$	<code>(X, Y, [Lags]=INT(10 LOG10(n/2)), [Alpha]=0.05, [Names])</code>
<code>Johansenλ</code>	$(K + 2) \times 10$	<code>(Data, [DetOrder]=0, [LagDiff]=1, [Alpha]=0.05: 0.10, 0.05 or 0.01)</code>
<code>VECMFitλ</code>	VECM block	<code>(Data, Rank, [LagDiff]=1, [Deterministic]="co", [Names])</code>
<code>VECMToVARλ</code>	VAR Model block	<code>(Model): the VECM as a VAR in levels, for every VAR function.</code>
<code>Eigenvaluesλ</code>	$(n + 1) \times 3$	<code>(Matrix): any real square matrix. Real, Imaginary, Modulus.</code>

All names start with `fx.Risk`, e.g. `fx.RiskVARFitλ`.

References: Sims (1980), *Econometrica* 48(1); Granger (1969), *Econometrica* 37(3); Johansen (1988, 1991); Lütkepohl (2005), *New Introduction to Multiple Time Series Analysis*; Hosking (1980), *JASA* 75; Zivot and Wang, *Modeling Financial Time Series with S-PLUS*, ch. 11; Woodward, Sadler and Robertson (2022), *Time Series for Data Science*, ch. 10. Data: statsmodels macrodata, from FRED (Federal Reserve Bank of St. Louis), public domain. XL Edge and the `fx.Risk...` functions: © 2026 W Phillips, edgewisedata.com, MIT License.