

lpirf — Local-projection impulse–response functions

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`lpirf` estimates local-projection impulse–response functions (IRFs). Dynamic multipliers on exogenous variables can also be computed.

Quick start

Local-projection IRFs for dependent variable `y1` with the default of 8 steps

```
lpirf y1
```

Same as above, but for dependent variables `y1` and `y2` with 12 steps

```
lpirf y1 y2, step(12)
```

Same as above, but use lags 1 through 4 instead of the default lags 1 and 2

```
lpirf y1 y2, step(12) lags(1/4)
```

Local-projection IRFs for `y1` and `y2` and dynamic multipliers for exogenous variable `x`

```
lpirf y1 y2, exog(x)
```

Same as above, but use robust standard errors, and make a small-sample degrees-of-freedom adjustment

```
lpirf y1 y2, exog(x) vce(robust) dfk
```

Menu

Statistics > Multivariate time series > Local-projection IRFs

Syntax

```
lpirf depvarlist [if] [in] [, options]
```

<i>options</i>	Description
Model	
<code>lags(<i>numlist</i>)</code>	include specified lags of dependent variables
<code>step(#)</code>	set forecast horizon to # steps; default is <code>step(8)</code>
<code>exog(<i>varlist</i>)</code>	include exogenous variables
<code>dfk</code>	include small-sample degrees-of-freedom adjustment
<code>small</code>	report small-sample <i>t</i> statistics
SE/Robust	
<code>vce(<i>vcetype</i>)</code>	<i>vcetype</i> may be <code>ols</code> , <code>robust</code> , or <code>hac</code> <i>kernel</i>
Reporting	
<code>level(#)</code>	set confidence level; default is <code>level(95)</code>
<code>display_options</code>	control columns and column formats and row spacing
<code>coeflegend</code>	display legend instead of statistics

You must `tsset` your data before using `lpirf`; see [TS] `tsset`.

`depvarlist` and `varlist` may contain time-series operators; see [U] 11.4.4 Time-series varlists.

`collect` is allowed; see [U] 11.1.10 Prefix commands.

`coeflegend` does not appear in the dialog box.

See [U] 20 Estimation and postestimation commands for more capabilities of estimation commands.

Options

Model

`lags(numlist)` specifies the lags of the dependent variables to be included in the model. The default is `lags(1 2)`. The first-lag coefficient is the local-projection coefficient and must be included. Additional lags may be included as controls. Lags may be skipped; for example, `lags(1 3)` would include lags 1 and 3 but not lag 2.

`step(#)` specifies the step (forecast) horizon; the default is eight periods.

`exog(varlist)` specifies a list of exogenous variables to be included in the local projections. Coefficients on contemporaneous exogenous variables are interpreted as dynamic multiplier coefficients. Other exogenous variables such as lags are included as controls.

`dfk` specifies that a small-sample degrees-of-freedom adjustment be used when estimating the residual covariance matrix. Specifically, $1/(T - m)$ is used instead of the large-sample divisor $1/T$, where m is the number of right-hand-side variables used in each local projection.

`small` causes `lpirf` to report small-sample *t* statistics instead of large-sample normal statistics.

SE/Robust

`vce(vctype)` specifies the type of standard error reported, which includes types that are derived from asymptotic theory (`ols`) and that are robust to some kinds of misspecification (`robust`); see [R] [vce_option](#).

`vce(ols)`, the default, uses the standard variance estimator for ordinary least-squares regression.

`lpirf` also allows the following:

`vce(hac kernel [#])` specifies that a heteroskedasticity- and autocorrelation-consistent (HAC) variance estimate be used. HAC refers to the general form for combining weighted matrices to form the variance estimate. *kernel* may be one of the following (you can specify the name used by statisticians or the name more familiar to economists):

`nwest` or `bartlett` requests the Newey–West (Bartlett) kernel.

`gallant` or `parzen` requests the Gallant (Parzen) kernel.

Reporting

`level(#)`; see [R] [Estimation options](#).

`display_options`: `noci`, `nopvalues`, `vsquish`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and `nostretch`; see [R] [Estimation options](#).

The following option is available with `lpirf` but is not shown in the dialog box:

`coeflegend`; see [R] [Estimation options](#).

Remarks and examples

[stata.com](https://www.stata.com)

An IRF measures how an endogenous variable responds in the future to a present-day shock to itself or another endogenous variable in the model. In economics, for example, these functions are used to study how inflation shocks today affect the evolution of prices and GDP in the months to come. Local-projections estimate the IRF directly by running multistep regressions of response variables on impulse variables. `lpirf` estimates these regressions jointly, allowing for tests of hypotheses involving multiple impulse–response coefficients.

Local projections were introduced by Jordà (2005) as an alternative to traditional IRF estimation based on vector autoregressive (VAR) models (see `var`). Local-projection estimation is not constrained by a model and thus provides more flexible impulse–response coefficients. Local projections also greatly simplify IRF estimation for multiple endogenous variables and horizons, making inference and hypothesis testing easier. See [example 3](#) in [TS] [lpirf postestimation](#) for an example. Moreover, confidence intervals provided by local projections can have better small-sample coverage than those based on the VAR’s asymptotic distribution and the delta method. The IRF coefficients are highly non-linear functions of the parameters in the VAR and thus the delta-method approximation can perform poorly. See Kilian and Kim (2011) and Montiel Olea and Plagborg-Møller (2021) for simulation evidence comparing the local-projection estimator against vector autoregression IRFs.

Because local projections are a direct multistep method, they have the disadvantage of having a reduced sample size compared with an equivalent VAR. Both local projections and VARs condition on p initial values. However, local projections also condition on $H - 1$ trailing values, where $H - 1$ is the maximum desired impulse–response step (horizon). Conditioning on the trailing observations can severely reduce the sample size when attempting to estimate long horizon responses using a short sample size.

The local-projection estimator conditions on past lags of the model variables. A series of regressions is run for each dependent variable and each horizon $h = 1, 2, \dots, H$

$$y_{i,t+h-1} = \theta_{ijh}y_{j,t-1} + \mathbf{z}_t\boldsymbol{\delta} + u_{t+h-1}$$

where y_i is the response variable and y_j is the impulse variable. The parameter of interest is the impulse–response coefficient θ_{ijh} . Additional controls \mathbf{z}_t dated t or earlier, such as further lags of the endogenous variables, may be included; their associated coefficients $\boldsymbol{\delta}$ are nuisance parameters.

Exogenous variables may also be included. The effect of the k th exogenous variable $x_{k,t}$ on the i th dependent variable h steps ahead $y_{i,t+h}$ is found by estimating the local projection

$$y_{i,t+h} = \phi_{ikh}x_{k,t} + \mathbf{z}_t\boldsymbol{\delta} + u_{t+h}$$

for $h = 0, 1, 2, \dots, H - 1$. The coefficient of interest is the dynamic multiplier coefficient ϕ_{ikh} . Additional controls \mathbf{z}_t dated t or earlier, such as further lags of the endogenous variables, may be included; their associated coefficients $\boldsymbol{\delta}$ are nuisance parameters.

`lpirf` estimates all the impulse–response and dynamic multiplier coefficients simultaneously, allowing for joint inference across arbitrary combinations of impulse variables, response variables, and horizons.

► **Example 1: Univariate local-projection model**

Consider a model with one endogenous variable, the interest rate, and no exogenous variables. In this univariate model, there is only one possible shock, so the IRF measures the effect of a shock to the interest rate on itself. We use data on the federal funds rate from 1955 to 2010 to estimate the local-projection IRF.

```
. use https://www.stata-press.com/data/r18/usmacro
(Federal Reserve Economic Data - St. Louis Fed)
. lpirf fedfunds
Local-projection impulse-responses
Sample: 1955q1 thru 2009q1
Number of obs      = 217
Number of impulses = 1
Number of responses = 1
Number of controls = 1
```

	IRF				
	coefficient	Std. err.	z	P> z	[95% conf. interval]
fedfunds					
F1.	1.221569	.0653797	18.68	0.000	1.093427 1.34971
F2.	1.080856	.1049245	10.30	0.000	.8752076 1.286504
F3.	1.118528	.1267763	8.82	0.000	.8700512 1.367005
F4.	1.166434	.145672	8.01	0.000	.8809218 1.451945
F5.	1.255471	.1620139	7.75	0.000	.9379294 1.573012
F6.	1.139327	.1803676	6.32	0.000	.7858127 1.492841
F7.	.8283938	.1959921	4.23	0.000	.4442565 1.212531
F8.	.7639486	.2038356	3.75	0.000	.3644381 1.163459

```
Impulses: fedfunds
Responses: fedfunds
Controls:  L2.fedfunds
```

The header of the output reports the time span of the estimation sample and the numbers of observations, impulse and response variables, and controls. The estimation table provides the impulse–response coefficients, ordered by impulse, response, and step. In this example, there is only one impulse and one response variable. By default, impulse–response coefficients are reported up to horizon 8. The footer contains information on the impulse and response variables; it also provides a full list of controls used in each local projection.

The IRF coefficients in the table are interpreted as follows. For a 1 percentage point shock to the interest rate, we expect the interest rate rises by 1.22 percentage points in the following period. The interest rate remains elevated for all steps 1 through 8. Eight periods after the shock, the interest rate remains 0.76 percentage points above its long-run value. By convention, the response at time 0 is set to 1 and is not reported.



► Example 2: Model with two endogenous variables

Next we estimate the impulse–response coefficients of a model with two endogenous variables, the output gap and the federal funds rate. We use the `lags()` option to specify lags 1 through 4 of both endogenous variables. Lags 2 through 4 are controls, while coefficients on the first lags are the impulse–responses of interest.

6 lpirf — Local-projection impulse–response functions

. lpirf ogap fedfunds, lags(1/4)

Local-projection impulse-responses

Sample: 1955q3 thru 2009q1

Number of obs = 215

Number of impulses = 2

Number of responses = 2

Number of controls = 6

		IRF				
		coefficient	Std. err.	z	P> z	[95% conf. interval]
ogap						
ogap						
	F1.	1.195888	.0703241	17.01	0.000	1.058055 1.333721
	F2.	1.357788	.1113556	12.19	0.000	1.139535 1.576041
	F3.	1.258217	.1461423	8.61	0.000	.971783 1.54465
	F4.	1.154532	.1703871	6.78	0.000	.8205797 1.488485
	F5.	.9490696	.1886837	5.03	0.000	.5792564 1.318883
	F6.	.8409692	.2005783	4.19	0.000	.447843 1.234095
	F7.	.7290706	.2099352	3.47	0.001	.3176052 1.140536
	F8.	.5371741	.2168984	2.48	0.013	.112061 .9622872
fedfunds						
	F1.	.340385	.0755028	4.51	0.000	.1924023 .4883677
	F2.	.6155252	.1219423	5.05	0.000	.3765228 .8545277
	F3.	.6645369	.1471894	4.51	0.000	.376051 .9530228
	F4.	.5518402	.1727784	3.19	0.001	.2132007 .8904797
	F5.	.6452311	.1949615	3.31	0.001	.2631136 1.027349
	F6.	.7062151	.2179209	3.24	0.001	.279098 1.133332
	F7.	.6356759	.2388224	2.66	0.008	.1675926 1.103759
	F8.	.3931728	.2502723	1.57	0.116	-.097352 .8836975
fedfunds						
ogap						
	F1.	.0706023	.0642572	1.10	0.272	-.0553395 .196544
	F2.	-.1944006	.1017489	-1.91	0.056	-.3938248 .0050235
	F3.	-.2381567	.1335345	-1.78	0.075	-.4998795 .0235661
	F4.	-.3092885	.1556876	-1.99	0.047	-.6144307 -.0041464
	F5.	-.4381839	.1724058	-2.54	0.011	-.776093 -.1002748
	F6.	-.5121501	.1832742	-2.79	0.005	-.871361 -.1529392
	F7.	-.5796563	.1918239	-3.02	0.003	-.9556242 -.2036884
	F8.	-.6646151	.1981864	-3.35	0.001	-1.053053 -.2761769
fedfunds						
	F1.	1.170208	.0689891	16.96	0.000	1.034992 1.305424
	F2.	.9238313	.1114222	8.29	0.000	.7054478 1.142215
	F3.	.9727738	.1344912	7.23	0.000	.7091758 1.236372
	F4.	.9796306	.1578727	6.21	0.000	.6702059 1.289055
	F5.	1.044107	.178142	5.86	0.000	.694955 1.393259
	F6.	1.019702	.1991207	5.12	0.000	.6294331 1.409972
	F7.	.7324305	.218219	3.36	0.001	.3047292 1.160132
	F8.	.6598259	.2286811	2.89	0.004	.2116192 1.108033

Impulses: ogap fedfunds

Responses: ogap fedfunds

Controls: L2.fedfunds L2.ogap L3.fedfunds L3.ogap L4.fedfunds L4.ogap

The estimation table has two blocks, one for each impulse variable. Responses are grouped by impulse and then by response.

The coefficients of the table are interpreted as follows. The effect of an impulse to `ogap` on `ogap` itself is 1.196 at horizon 1, rises to 1.358 at horizon 2, and declines to 0.537 at horizon 8. The effect of an impulse to `ogap` on `fedfunds` is 0.340 at horizon 1, rises to 0.706 at horizon 6, and declines to 0.393 at horizon 8. A shock to `fedfunds` is associated with a rise in `ogap` of 0.071 at horizon 1, and the response at every other horizon is negative. The response of `ogap` at horizon 8 is -0.665 . ◀

▷ Example 3: Estimating dynamic multipliers

The local-projection estimator allows for the inclusion of exogenous variables. The responses of endogenous variables to shocks to exogenous variables are called dynamic multipliers. In this example, we fit a three-equation model of inflation, the output gap, and the interest rate, treating the interest rate as exogenous.

We include the `step()` option to compute the responses for four steps, instead of the default eight. We use the `exog()` option to specify `fedfunds` as the exogenous variable.

8 lpirf — Local-projection impulse–response functions

```
. lpirf ogap inflation, lags(1/4) step(4) exog(fedfunds)
```

Local-projection impulse-responses

Sample: 1956q3 thru 2010q1

Number of obs = 215

Number of impulses = 3

Number of responses = 2

Number of controls = 6

		IRF				
		coefficient	Std. err.	z	P> z	[95% conf. interval]
ogap						
ogap						
F1.		1.183248	.0696613	16.99	0.000	1.046714 1.319782
F2.		1.298995	.1084817	11.97	0.000	1.086374 1.511615
F3.		1.240221	.1408233	8.81	0.000	.9642124 1.51623
F4.		1.163917	.1641595	7.09	0.000	.8421708 1.485664
inflation						
F1.		.065811	.0546801	1.20	0.229	-.04136 .172982
F2.		.1562728	.0875892	1.78	0.074	-.0153988 .3279444
F3.		.2659506	.1150475	2.31	0.021	.0404616 .4914396
F4.		.4547818	.1445588	3.15	0.002	.1714517 .7381119
inflation						
ogap						
F1.		-.1348397	.0849333	-1.59	0.112	-.301306 .0316265
F2.		-.2116188	.1322643	-1.60	0.110	-.4708521 .0476145
F3.		-.2891468	.1716962	-1.68	0.092	-.6256653 .0473716
F4.		-.4686839	.2001485	-2.34	0.019	-.8609677 -.0764001
inflation						
F1.		1.205986	.0666677	18.09	0.000	1.075319 1.336652
F2.		1.209659	.1067915	11.33	0.000	1.000351 1.418966
F3.		1.396595	.1402696	9.96	0.000	1.121671 1.671518
F4.		.9968414	.1762507	5.66	0.000	.6513963 1.342286
fedfunds						
ogap						
--.		.0289642	.0270094	1.07	0.284	-.0239732 .0819016
F1.		-.0017852	.0420609	-0.04	0.966	-.0842231 .0806527
F2.		-.0506758	.0546005	-0.93	0.353	-.1576909 .0563393
F3.		-.0776979	.0636485	-1.22	0.222	-.2024467 .0470509
inflation						
--.		.073647	.0212008	3.47	0.001	.0320943 .1151997
F1.		.115983	.0339604	3.42	0.001	.0494218 .1825442
F2.		.120239	.0446067	2.70	0.007	.0328115 .2076664
F3.		.1092225	.0560489	1.95	0.051	-.0006313 .2190763

Note: IRF coefficients for exogenous variables are dynamic multipliers.

Impulses: ogap inflation fedfunds

Responses: ogap inflation

Controls: L2.inflation L2.ogap L3.inflation L3.ogap L4.inflation L4.ogap

The estimation table has three blocks, one for each impulse. The first block reports the responses of the output gap and inflation to an impulse to the output gap. The second block reports the responses of the output gap and inflation to an impulse to inflation. The third block reports the dynamic multipliers of the output gap and inflation to an impulse to the interest rate.

The contemporaneous effect of the exogenous variable is not constrained to be 1. These effects are displayed in the third block of the table. The contemporaneous effect, also known as the impact effect, of a shock to the interest rate on the output gap is 0.029 with a standard error of 0.027. The responses of the output gap at horizons 1, 2, and 3 are all negative. The effect of a shock to the interest rate on inflation is 0.074 on impact. Inflation continues to rise in the first three periods after a shock, reaching a value of 0.109 in the third period after the shock.

◀

► Example 4: Robust inference

Several options control how standard errors are calculated and reported. The `dfk` option provides a small-sample degrees-of-freedom adjustment for estimating the residual covariance matrix. The `small` option requests that small-sample t statistics be reported instead of large-sample z statistics.

We revisit the simple univariate model of the interest rate, but this time we specify the `dfk` and `small` options.

```
. lpirf fedfunds, dfk small
Local-projection impulse-responses
Sample: 1955q1 thru 2009q1
                                     Number of obs      = 217
                                     Number of impulses = 1
                                     Number of responses = 1
                                     Number of controls  = 1
```

	IRF					
	coefficient	Std. err.	t	P> t	[95% conf. interval]	
fedfunds						
F1.	1.221569	.0658364	18.55	0.000	1.091798	1.351339
F2.	1.080856	.1056574	10.23	0.000	.8725934	1.289118
F3.	1.118528	.1276618	8.76	0.000	.8668925	1.370164
F4.	1.166434	.1466895	7.95	0.000	.8772924	1.455575
F5.	1.255471	.1631456	7.70	0.000	.9338928	1.577049
F6.	1.139327	.1816275	6.27	0.000	.7813188	1.497335
F7.	.8283938	.1973611	4.20	0.000	.4393732	1.217414
F8.	.7639486	.2052594	3.72	0.000	.3593595	1.168538

Note: Small-sample degrees-of-freedom adjustment applied when estimating covariance matrix of residuals.

```
Impulses: fedfunds
Responses: fedfunds
Controls:  L2.fedfunds
```

Confidence intervals here are somewhat wider than in [example 1](#).

Montiel Olea and Plagborg-Møller (2021) provide a thorough discussion of inference in local-projection models. They suggest two additional adjustments. First, if the true model is believed to be a VAR of order p , then $p + 1$ lags should be included in the local projections. Second, White-corrected standard errors should be reported.

We use the `lags(1/3)` option to add one further lag to the controls. Additionally, we specify the `vce(robust)` option to provide White-corrected standard errors. With all of these options, the output of `lpirf` changes to the following:

```
. lpirf fedfunds, lags(1/3) dfk small vce(robust)
Local-projection impulse-responses
Sample: 1955q2 thru 2009q1
Number of obs      = 216
Number of impulses = 1
Number of responses = 1
Number of controls = 2
```

	IRF coefficient	Robust std. err.	t	P> t	[95% conf. interval]	
fedfunds						
F1.	1.264674	.1123728	11.25	0.000	1.043163	1.486185
F2.	1.087209	.1681724	6.46	0.000	.755705	1.418714
F3.	1.098138	.1927819	5.70	0.000	.7181227	1.478153
F4.	1.105892	.2030137	5.45	0.000	.705708	1.506076
F5.	1.216923	.2099187	5.80	0.000	.8031281	1.630719
F6.	1.168236	.2192774	5.33	0.000	.7359926	1.600479
F7.	.8415941	.2382936	3.53	0.001	.3718657	1.311323
F8.	.7565939	.2355971	3.21	0.002	.2921808	1.221007

Note: Small-sample degrees-of-freedom adjustment applied when estimating covariance matrix of residuals.

Impulses: fedfunds
 Responses: fedfunds
 Controls: L2.fedfunds L3.fedfunds

The use of three lags indicates a belief that the underlying VAR is of order two. Thus, three lags provide the appropriate adjustment. We also adjust for the possibility of heteroskedasticity by specifying `vce(robust)`, which is equivalent to the standard errors that would be obtained with `regress, vce(robust)`.

For comparison, the results of `sureg` for the first two response horizons would be given by the following:

```
. sureg (f(0/1).fedfunds = 1(1/3).fedfunds) if e(sample), vce(robust) dfk small
Seemingly unrelated regression
```

Equation	Obs	Params	RMSE	"R-squared"	F	P>F
fedfunds	216	3	.8940584	0.9278	412.58	0.0000
F_fedfunds	216	3	1.453377	0.8102	129.56	0.0000

	Coefficient	Robust std. err.	t	P> t	[95% conf. interval]	
fedfunds						
fedfunds						
L1.	1.264674	.1113275	11.36	0.000	1.045852	1.483497
L2.	-.4670027	.1513732	-3.09	0.002	-.7645381	-.1694674
L3.	.1604177	.121733	1.32	0.188	-.0788575	.3996929
_cons	.2333601	.1568402	1.49	0.138	-.0749211	.5416413
F_fedfunds						
fedfunds						
L1.	1.087209	.166608	6.53	0.000	.7597289	1.41469
L2.	-.2141638	.242216	-0.88	0.377	-.6902574	.2619298
L3.	.023556	.1679403	0.14	0.889	-.3065431	.3536551
_cons	.5746647	.2276758	2.52	0.012	.1271509	1.022178

Note: Small-sample degrees-of-freedom adjustment applied when estimating covariance matrix of residuals.

The regression of the current interest rate on its first three lags is identical to the underlying projection for the first horizon. The coefficient estimate on `L.fedfunds` is 1.265, identical to the coefficient on `F1.fedfunds` in the local-projection estimation table. The standard errors are also identical, 0.111.

The regression of the one-period-ahead interest rate on the first three lags of the interest rate in the `sureg` output is identical to the underlying projection for the second horizon of `lpirf`. The coefficient estimate on `L.fedfunds` is 1.087, identical to the coefficient on `F2.fedfunds` in the local-projection estimation table. The standard errors are also identical, 0.167.

Newey–West standard errors may be requested by specifying `vce(hac nwest #)`, where `#` represents the maximum lag order of autocorrelation. If a number is not specified, the number of lags used is equal to the maximum step.

```
. lpirf fedfunds, lags(1/3) dfk small vce(hac nwest)
```

Local-projection impulse-responses

Sample: 1955q2 thru 2009q1

Number of obs = 216

Number of impulses = 1

HAC kernel: Newey–West with 8 lags

Number of responses = 1

Number of controls = 2

	IRF	HAC				
	coefficient	std. err.	t	P> t	[95% conf. interval]	
fedfunds						
F1.	1.264674	.1367966	9.24	0.000	.9950182	1.53433
F2.	1.087209	.1731642	6.28	0.000	.7458652	1.428554
F3.	1.098138	.1677261	6.55	0.000	.767513	1.428762
F4.	1.105892	.2446402	4.52	0.000	.6236531	1.588131
F5.	1.216923	.2595175	4.69	0.000	.7053581	1.728489
F6.	1.168236	.2757997	4.24	0.000	.6245749	1.711897
F7.	.8415941	.2968191	2.84	0.005	.2564993	1.426689
F8.	.7565939	.3054195	2.48	0.014	.1545458	1.358642

Note: Small-sample degrees-of-freedom adjustment applied when estimating covariance matrix of residuals.

Impulses: fedfunds

Responses: fedfunds

Controls: L2.fedfunds L3.fedfunds

The Newey–West standard errors reported here correct for both heteroskedasticity and autocorrelation and are often recommended for local-projection IRFs.

Stored results

lpirf stores the following in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(k)</code>	number of parameters
<code>e(k_endog)</code>	number of dependent variables
<code>e(k_exog)</code>	number of exogenous variables
<code>e(k_impulses)</code>	number of impulse variables
<code>e(k_responses)</code>	number of response variables
<code>e(k_controls)</code>	number of control variables
<code>e(df_r)</code>	residual degrees of freedom (small only)
<code>e(tmin)</code>	minimum time in sample
<code>e(tmax)</code>	maximum time in sample
<code>e(step)</code>	maximum step
<code>e(rank)</code>	rank of <code>e(V)</code>

Macros

<code>e(cmd)</code>	lpirf
<code>e(cmdline)</code>	command as typed
<code>e(depvar)</code>	names of dependent variables
<code>e(title)</code>	title in estimation output
<code>e(allcontrols)</code>	list of control variables used in each projection
<code>e(tsfmt)</code>	format of the time variable
<code>e(tvar)</code>	variable denoting time within groups
<code>e(tmins)</code>	formatted minimum time
<code>e(tmaxs)</code>	formatted maximum time
<code>e(lags)</code>	lags used in controls
<code>e(dfk)</code>	dfk, if specified
<code>e(small)</code>	small, if specified
<code>e(endog)</code>	names of endogenous variables
<code>e(exog)</code>	names of exogenous variables
<code>e(vce)</code>	<i>vctype</i> specified in <code>vce()</code>
<code>e(vctype)</code>	title used to label Std. err.
<code>e(properties)</code>	b V

Matrices

<code>e(b)</code>	coefficient vector
<code>e(V)</code>	variance–covariance matrix of the estimators

Functions

<code>e(sample)</code>	marks estimation sample
------------------------	-------------------------

In addition to the above, the following is stored in `r()`:

Matrices

<code>r(table)</code>	matrix containing the coefficients with their standard errors, test statistics, <i>p</i> -values, and confidence intervals
-----------------------	--

Note that results stored in `r()` are updated when the command is replayed and will be replaced when any *r*-class command is run after the estimation command.

Methods and formulas

lpirf is equivalent to running `regress`, equation by equation, for each possible combination of impulse variable, response variable, and horizon (step). lpirf performs these regressions jointly to allow for tests of hypotheses that involve multiple impulse–response coefficients.

Let y_t be the $k \times 1$ vector of endogenous variables in the model. For a given horizon h , consider the regression of response variables at time $t + h - 1$ on lags 1 through p of the dependent variables. Hence, for horizon 1, we use the current values of y on the left-hand side; for horizon 2, we use the

values of \mathbf{y} one period ahead; for horizon 3, we use the values of \mathbf{y} two periods ahead; and so on. In each regression, the right-hand-side variables remain the same, lags 1 through p of the dependent variables.

When the model does not include any exogenous variables, the system of regressions is

$$\mathbf{y}_{t+h-1} = \mathbf{B}_1^h \mathbf{y}_{t-1} + \mathbf{B}_2^h \mathbf{y}_{t-2} + \cdots + \mathbf{B}_p^h \mathbf{y}_{t-p} + \mathbf{u}_{t+h-1}$$

for $h = 1, \dots, H$, where H is the desired maximum horizon. In the system of regressions, \mathbf{B}_1^h is the $k \times k$ matrix of impulse–response coefficients at horizon h . The `lags()` option specifies which additional lags to use as controls in the local projections. Missing lags are allowed.

Contemporaneous exogenous variables are used to estimate dynamic multipliers. Here the timing convention is simpler: a regression of the endogenous variables at time $t+h$, \mathbf{y}_{t+h} , on the exogenous variable today, \mathbf{x}_t , gives the dynamic multiplier at horizon h :

$$\mathbf{y}_{t+h} = \mathbf{D}^h \mathbf{x}_t + \mathbf{\Gamma} \mathbf{z}_t + \mathbf{u}_{t+h}$$

Matrix \mathbf{D}^h contains the dynamic multipliers at horizon h . Additional controls dated t or earlier, \mathbf{z} , can be specified using the `exog()` option.

OLS and robust standard errors are calculated as in `sureg`. The `dfk`, `small`, and `vce(robust)` options perform in `lpirf` exactly as they do in `sureg`; see [R] [sureg](#). The `vce(hac nwest #)` and `vce(hac gallant #)` options perform exactly as they do in [R] [gmm](#).

References

- Jordà, Ò. 2005. Estimation and inference of impulse responses by local projections. *American Economic Review* 95: 161–182. <https://doi.org/10.1257/0002828053828518>.
- Kilian, L., and Y. J. Kim. 2011. How reliable are local projection estimators of impulse responses? *Review of Economics and Statistics* 93: 1460–1466. https://doi.org/10.1162/REST_a_00143.
- Montiel Olea, J. L., and M. Plagborg-Møller. 2021. Local projection inference is simpler and more robust than you think. *Econometrica* 89: 1789–1823. <https://doi.org/10.3982/ECTA18756>.

Also see

- [TS] [lpirf postestimation](#) — Postestimation tools for `lpirf`
- [TS] [tsset](#) — Declare data to be time-series data
- [TS] [var](#) — Vector autoregressive models
- [R] [regress](#) — Linear regression
- [U] [20 Estimation and postestimation commands](#)