Local projections and fiscal multipliers

Tyler Sotomayor

2024

These solutions develop the quarterly sample calculations and covariance formulas used in the local-projections exercise. The IV derivation explains why the structural residual matters for inference. A final example works through a test of equal responses across unemployment states.

Based on the 2024 BSE course taught by Luca Gambetti and Konstantin Boss. I have revised the classroom code and added the accompanying explanations. The notes identify the papers, data, and assumptions used in each application. These are my course notes, not an official BSE publication. A classroom replication should not be read as a replication of every result in its underlying paper.

1 Quarterly samples

Let $q(y, Q) = 4y + (Q - 1)$ denote the integer code for quarter Q of year y. The number of consecutive quarters from a through b, including both endpoints, is $q(b) - q(a) + 1$. The impact sample therefore contains

$$[4(2015) + 3] - [4(1891) + 0] + 1 = 500$$

origins. At horizon h, the last usable origin must satisfy $q(t) + h \leq q(2015\ \mathrm{Q4})$. With no interior gaps and a fixed first origin, the final origin moves back h quarters and $n_h = 500 - h$.

Horizon	First origin	Last origin	Observations
0	1891 Q1	2015 Q4	500
4	1891 Q1	2014 Q4	496
20	1891 Q1	2010 Q4	480

The last column counts regression origins, not rows in the workbook. For example, 2010 Q4 is the final horizon-twenty origin because its outcome is observed five years later, in 2015 Q4. A horizon of twenty does not mean twenty observations in the estimation sample.

Now let m be the integer code for the missing outcome in 1950 Q2. The level outcome at horizon four is unavailable only when $q(t) + 4 = m$. The affected origin is therefore 1949 Q2. A missing intermediate quarter does not, by itself, invalidate an observed level endpoint.

The cumulative outcome is $y_t + y_{t+1} + y_{t+2} + y_{t+3} + y_{t+4}$. It is unavailable whenever

$$q(t) \leq m \leq q(t) + 4 \quad \Longleftrightarrow \quad m - 4 \leq q(t) \leq m.$$

Thus the five affected origins are 1949 Q2, 1949 Q3, 1949 Q4, 1950 Q1, and 1950 Q2. Each corresponding window includes the missing quarter. Ignoring that value while summing would change a five-quarter cumulative outcome into a four-quarter sum. Substituting the next available row would instead shift a dated outcome to the wrong horizon.

This calculation holds the regressors fixed. Missing news or an outcome that also enters the lag controls can remove further origins. The actual estimator requires both the relevant outcome window and the entire design row to be observed; it does not infer the sample size from the endpoints alone when there are interior gaps.

2 OLS and the HAC covariance

Write $x = (-3, -1, 1, 3)'$ for the second column of X. Its entries sum to zero, while $\sum_t x_t^2 = 9 + 1 + 1 + 9 = 20$. The two cross-products with the outcome are

$$\sum_t y_t = 1.5 + 0.5 + 1.5 + 4.5 = 8,$$

$$\sum_t x_t y_t = -4.5 - 0.5 + 1.5 + 13.5 = 10.$$

The centered second regressor makes the cross-product matrix diagonal:

$$X'X = \begin{pmatrix} 4 & 0 \\ 0 & 20 \end{pmatrix}, \qquad X'y = \begin{pmatrix} 8 \\ 10 \end{pmatrix}.$$

Solving $4\widehat{a} = 8$ and $20\widehat{b} = 10$ gives an intercept of 2 and slope of 0.5. The fitted values and residuals can be checked row by row:

t	x_t	y_t	$2 + 0.5x_t$	$\widehat{u}_t$
1	-3	1.5	0.5	1
2	-1	0.5	1.5	-1
3	1	1.5	2.5	-1
4	3	4.5	3.5	1

The intercept orthogonality condition is $1 - 1 - 1 + 1 = 0$. For the second regressor it is $-3 + 1 - 1 + 3 = 0$. Hence $X'\widehat{u} = 0$, as required by the OLS first-order conditions.

Every squared residual equals one. Therefore the HC0 meat is $X'X$, and

$$\widehat{V}_{HC0} = (X'X)^{-1}(X'X)(X'X)^{-1} = \begin{pmatrix} 1/4 & 0 \\ 0 & 1/20 \end{pmatrix}.$$

The intercept standard error is $\sqrt{1/4} = 0.5$; the slope standard error is $\sqrt{1/20} \simeq 0.223607$. There is no degrees-of-freedom multiplier. An HC1 adjustment would multiply this covariance by $n/(n-k) = 4/(4-2) = 2$, and multiply each standard error by $\sqrt{2}$. It would not answer the stated HC0 question.

For the separate scalar-score example, the zero-lag sum is $1^2 + 2^2 + 3^2 = 14$. With maximum lag $L = 1$, the Bartlett weight is $w_1 = 1 - 1/(1+1) = 1/2$. Consecutive dates give two lag-one pairs, so

$$\sum_t s_t s_{t-1} = 2 \cdot 1 + 3 \cdot 2 = 8, \qquad \widehat{S} = 14 + \tfrac{1}{2}(8+8) = 22.$$

When the dates are 1, 3, and 4, the scores 1 and 2 are two quarters apart. They contribute nothing at lag one. Only the scores at quarters 3 and 4 are paired, giving

$$\widehat{S} = 14 + \tfrac{1}{2}(6+6) = 20.$$

The zero-lag term is unchanged: all three scores remain in the sample. Only the set of calendar pairs changes. These last two answers are HAC meats, not coefficient variances; a regression covariance would also require the appropriate bread on both sides.

3 Structural residuals in IV

Let X contain the structural regressors, including the endogenous cumulative-spending variable. Let Z contain the excluded instrument and the included controls. Both matrices refer to the same n observations. Assume Z has full column rank and $X'P_Z X$ is nonsingular. Define

$$P_Z = Z(Z'Z)^{-1}Z', \qquad \widehat{X} = P_Z X.$$

The projection satisfies $P_Z' = P_Z$ and $P_Z^2 = P_Z$. Therefore

$$\widehat{X}'\widehat{X} = X'P_Z'P_Z X = X'P_Z X, \qquad \widehat{X}'Y = X'P_Z Y.$$

Regressing Y on $\widehat{X}$ consequently yields

$$(\widehat{X}'\widehat{X})^{-1}\widehat{X}'Y = (X'P_Z X)^{-1}X'P_Z Y,$$

which is the 2SLS coefficient vector. The implementation uses QR factorizations and linear solves to evaluate this expression, rather than forming the displayed inverses. The equality concerns the point estimate; it does not identify the disturbance in the structural equation.

The two residuals are $\widehat{e} = Y - X\widehat{\theta}$ and $\widehat{r} = Y - \widehat{X}\widehat{\theta}$. Their difference is $\widehat{r} - \widehat{e} = (X - \widehat{X})\widehat{\theta}$. The latter term is generally nonzero when instruments explain only part of the endogenous regressor. Included exogenous controls lie in the instrument span, so they project onto themselves. In this application, the difference comes from the part of cumulative spending that the instruments do not explain, multiplied by its estimated structural coefficient.

To make the covariance distinction explicit, write $\mathcal{H}_L(s)$ for the dated Bartlett score sum used in the lecture notes, and let $A = \widehat{X}'\widehat{X}$. The two calculations are

$$\widehat{V}_{IV} = A^{-1}\mathcal{H}_L(\widehat{x}_t\widehat{e}_t)A^{-1},$$

$$\widehat{V}_{fitted} = A^{-1}\mathcal{H}_L(\widehat{x}_t\widehat{r}_t)A^{-1}.$$

The bread is the same. The meat differs because its scores use different disturbance estimates. Ordinary robust or HAC inference from the second regression uses the second formula; structural IV inference requires the first. Changing only the label on the second-stage regression does not change which residual the software uses.

The production line `fit.residual = y - X * fit.beta` uses the actual regressors. Those residuals are multiplied by `projectedX` to form the IV scores. `classroomSE` is stored separately for comparison and is not used as the preferred IV standard error.

There is no general ordering of the two standard errors. If $d_t = \widehat{r}_t - \widehat{e}_t$, the fitted-regression score is the sum of the structural score and $\widehat{x}_t d_t$. Its HAC meat contains the score covariance of this extra term and cross-products between the extra term and the structural score, including their lagged cross-products. Those cross-products can have either sign. A smaller corrected standard error is therefore not, by itself, evidence of an implementation error; neither is a larger one.

If $Z = X$ and X has full column rank, then $P_Z X = X$. The two residuals coincide and 2SLS reduces to OLS:

$$\widehat{\theta} = (X'X)^{-1}X'Y, \qquad \widehat{V}_{IV} = \widehat{V}_{fitted} = \widehat{V}_{OLS}.$$

This equality requires the same observations, calendar, bandwidth, and finite-sample convention. The test suite checks equality of the full covariance matrices under those conditions, not only one coefficient's standard error.

4 A cumulative multiplier

Using the rounded values in the exercise,

$$0.680 \pm 1.96(0.098) = 0.680 \pm 0.19208,$$

so the displayed interval is $[0.48792, 0.87208]$. Both zero and one lie outside it. Under the conventional normal approximation, the corresponding two-sided pointwise statistics are

$$z_0 = \frac{0.680 - 0}{0.098} \simeq 6.939, \qquad z_1 = \frac{0.680 - 1}{0.098} \simeq -3.265.$$

Each absolute value exceeds 1.96. Thus the conventional tests reject a zero multiplier and a unit multiplier at the five percent level in this single-horizon specification. These are statements about the stated inference procedure, not independent confirmation of the identifying assumptions or of its finite-sample coverage.

At horizon four, the dependent variable sums normalized output over $t, t+1, t+2, t+3, t+4$. The endogenous regressor sums normalized purchases over the same five quarters. The coefficient relates those two cumulative quantities using current news as an instrument and the specified lagged controls. Under the source exercise's normalization and identifying assumptions, 0.680 is read as approximately 68 cents of cumulative output per dollar of cumulative government purchases. It is not GDP growth of 68 percent, nor is it the output response in the fourth quarter alone.

It is a pointwise, asymptotic normal interval at horizon four, not a simultaneous band across horizons. It does not test instrument exclusion, rule out omitted confounding events, or provide weak-instrument-robust coverage. In particular, conventional normal approximations can be unreliable when the first stage is weak. The coefficient is learned from variation in cumulative purchases explained by news, conditional on the controls. A calculated standard error and the factor 1.96 do not repair the problem when that variation is weak.

The reported relevance measure is the HAC Wald statistic for the excluded instrument, divided by the number of excluded restrictions. With one excluded instrument, it is the squared HAC coefficient statistic from that first stage. It is not a Kleibergen–Paap statistic and is not used with an automatic cutoff to declare the instrument strong. Partial R^2 records its incremental explanatory power relative to the included controls. Neither measure tests the instrument's exclusion restriction.

Pointwise coverage concerns a fixed horizon chosen for the question. Repeating the calculation at 21 horizons does not make the probability that all intervals cover their respective parameters equal to 95 percent. A simultaneous band requires a procedure that accounts for the joint estimation uncertainty across horizons. The current plot does not supply that procedure.

The code calculates with unrounded coefficients and errors. Small differences between this hand interval and exported bounds are therefore expected.

5 Initial unemployment and state contrasts

Let $I_{t-1} = 1$ when last quarter's unemployment rate is at least 6.5 percent, and zero otherwise. With x_t denoting the complete linear-LP regressor vector, including its intercept, the state equation is

$$y_{t+h} = I_{t-1} x_t' \theta_h^H + (1 - I_{t-1}) x_t' \theta_h^L + u_{t+h}.$$

The two state-intercept columns sum to the ordinary constant: $I_{t-1} + (1 - I_{t-1}) = 1$. Adding another intercept would make the design rank deficient. Interacting the intercept and all controls allows the conditional mean to differ across initial states, not only the news slope.

Moving the state from $t-1$ to t makes it contemporaneous with the shock. If the shock can affect unemployment in t, it can also affect the state to which the observation is assigned. The interpretation is no longer unambiguously conditional on pre-shock unemployment. Lagging the state avoids that particular timing problem; it does not prove exogeneity of the news. Nor does the specification require unemployment to remain in the initial state throughout the response horizon.

For the equality test, let j_H and j_L be the locations of `high_news` and `low_news` in the full coefficient vector. Define c to be zero except for $c_{j_H} = 1$ and $c_{j_L} = -1$. The estimated difference is $\widehat{d} = c'\widehat{\theta} = \widehat{b}_h^H - \widehat{b}_h^L$. Its estimated variance is

$$\widehat{V}_d = c'\widehat{V}c = V_{HH} + V_{LL} - 2V_{HL}.$$

Provided $\widehat{V}_d > 0$, the pointwise normal statistic for $H_0 : d = 0$ is $z = \widehat{d}/\sqrt{\widehat{V}_d}$, with interval $\widehat{d} \pm 1.96\sqrt{\widehat{V}_d}$. Equivalently, the one-restriction Wald statistic is z^2, compared asymptotically with a χ_1^2 distribution. These approximations inherit the assumptions and limitations of the underlying HAC covariance.

An observation cannot contribute to both state blocks at the same date: $I_{t-1}(1 - I_{t-1}) = 0$. That removes contemporaneous cross-block score products. HAC also uses products across dates, however. A high-state origin and a nearby low-state origin can form a lagged score pair. Their cross-product need not vanish. Treating the two coefficients as independent discards precisely these terms from the contrast variance.

For the stated teaching example, each marginal standard error is $\sqrt{0.09} = 0.3$. The separate intervals are

$$0.8 \pm 1.96(0.3) = [0.212,\ 1.388],$$

$$0.4 \pm 1.96(0.3) = [-0.188,\ 0.988].$$

They overlap substantially. But the difference is $0.8 - 0.4 = 0.4$, and

$$\widehat{V}_d = 0.09 + 0.09 - 2(0.08) = 0.02, \qquad \text{se}(\widehat{d}) = \sqrt{0.02} \simeq 0.141421.$$

Its interval is

$$0.4 \pm 1.96\sqrt{0.02} \simeq [0.122814,\ 0.677186].$$

This interval excludes zero; $z \simeq 2.828$ exceeds 1.96. The strong positive covariance makes the difference more precisely estimated than either coefficient separately. The example shows why overlapping marginal intervals do not imply failure to reject equality. Non-overlap is not the general testing rule either: the formal test uses the variance of the difference. Testing a collection of such differences across horizons would require a further decision about joint inference.