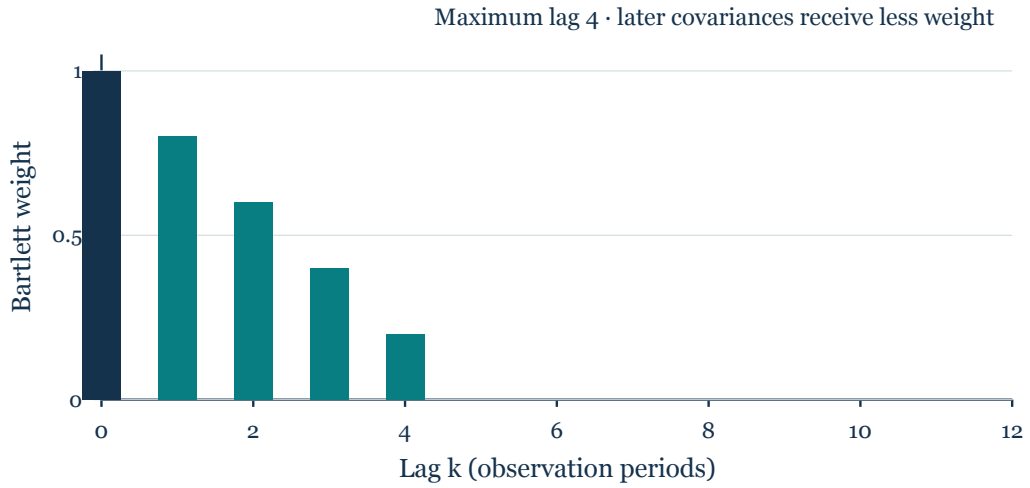
Choose between revising the mean, estimator and covariance

> Revise the conditional mean when the economic model needs additional dynamics.

> Compare model-based error transformations with HAC covariance under their stated assumptions.

Use HAC to account for unequal spread and time dependence

Newey–West keeps the OLS coefficients and changes their estimated uncertainty.



Declining lag weights only. The slider does not recalculate the Stata standard errors.

- $a_t = rX_t\hat{u}_t$, where rX_t is the regressor remainder after the other inputs.
- $w_h = 1 - h/(L+1)$ for $h = 1, \ldots, L$. For $L = 4$: $0.8, 0.6, 0.4, 0.2$.
- Variance estimate: $\dfrac{\sum_t a_t^2 + 2\sum_{h=1}^{L} w_h \sum_{t=h+1}^{T} a_t a_{t-h}}{(\sum_t rX_t^2)^2}$.

Predictable error variance is different from a predictable error mean

ARCH describes volatility clustering, not necessarily directional error predictability.

```
. estat archlm, lags(1)
LM test for autoregressive conditional heteroskedasticity (ARCH)
------------------------------------------------------------------------------
    lags(p)  |          chi2               df                 Prob > chi2
-------------+----------------------------------------------------------------
       1     |          0.350               1                   0.5540
------------------------------------------------------------------------------
         H0: no ARCH effects      vs.  H1: ARCH(p) disturbance
```

Actual StataMP output. Reproduce with overleaf_technical.do.

$E[u_t \mid I_{t-1}] = 0, E[u_t^2 \mid I_{t-1}] = \alpha_0 + \alpha_1 u_{t-1}^2, \alpha_0 > 0, \alpha_1 \geq 0.$

Conventional diagnostic: regress residual squares on a constant and their selected lags.

Actual monthly diagnostic: $\chi_1^2 \approx 0.350$, $p \approx 0.554$.

Run the conventional regression

Find the law coefficient and understand the labels before comparing SEs.

```
. regress logacc trend i.month unem spdlaw beltlaw wkends

------------------------------------------------------------------------------
      logacc | Coefficient  Std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
       trend |   .0011011   .0002579     4.27   0.000     .0005889    .0016133
             |
        unem |  -.0212173   .0033974    -6.25   0.000    -.0279659   -.0144688
      spdlaw |  -.0537593   .0126036    -4.27   0.000    -.0787948   -.0287238
     beltlaw |   .0954526    .014235     6.71   0.000     .0671765    .1237288
      wkends |   .0033333   .0037761     0.88   0.380    -.0041675    .0108342
       _cons |   10.63986   .0630859   168.66   0.000     10.51455    10.76518
------------------------------------------------------------------------------
```

Calendar-indicator coefficient rows are omitted here. The output download contains the full numerical results.

Run heteroskedasticity-robust inference

vce(robust) allows unequal error variance, while preserving the OLS coefficients.

```
. regress logacc trend i.month unem spdlaw beltlaw wkends, vce(robust)

------------------------------------------------------------------------------
             |               Robust
      logacc | Coefficient  std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
       trend |   .0011011   .0002676     4.12   0.000     .0005696    .0016326
             |
        unem |  -.0212173   .0038238    -5.55   0.000    -.0288128   -.0136218
      spdlaw |  -.0537593   .0120059    -4.48   0.000    -.0776076    -.029911
     beltlaw |   .0954526   .0123682     7.72   0.000     .0708847    .1200206
      wkends |   .0033333   .0039443     0.85   0.400    -.0045016    .0111682
       _cons |   10.63986   .0681238   156.18   0.000     10.50454    10.77518
------------------------------------------------------------------------------
```

Calendar-indicator coefficient rows are omitted here. The output download contains the full numerical results.

Run Newey–West on the same model

newey with lag(4) allows covariance contributions through four months.

```
. newey logacc trend i.month unem spdlaw beltlaw wkends, lag(4)

------------------------------------------------------------------------------
             |             Newey–West
      logacc | Coefficient  std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
       trend |   .0011011   .0003897     2.83   0.006      .000327    .0018752
             |
        unem |  -.0212173   .0051302    -4.14   0.000    -.0314079   -.0110268
      spdlaw |  -.0537593   .0170449    -3.15   0.002    -.0876169   -.0199018
     beltlaw |   .0954526   .0188985     5.05   0.000     .0579131    .1329922
      wkends |   .0033333   .0028746     1.16   0.249    -.0023768    .0090434
       _cons |   10.63986   .0635823   167.34   0.000     10.51356    10.76616
------------------------------------------------------------------------------
```

Calendar-indicator coefficient rows are omitted here. The output download contains the full numerical results.

Convert the fitted log difference into percent

Use 100[exp(b)−1] for the exact contrast between exponentiated fitted values.

```
. nlcom 100*(exp(_b[beltlaw])-1)

------------------------------------------------------------------------------
      logacc | Coefficient  Std. err.      z    P>|z|     [95% conf. interval]
-------------+----------------------------------------------------------------
       _nl_1 |   10.01567    2.07913     4.82   0.000     5.940653    14.09069
------------------------------------------------------------------------------
```

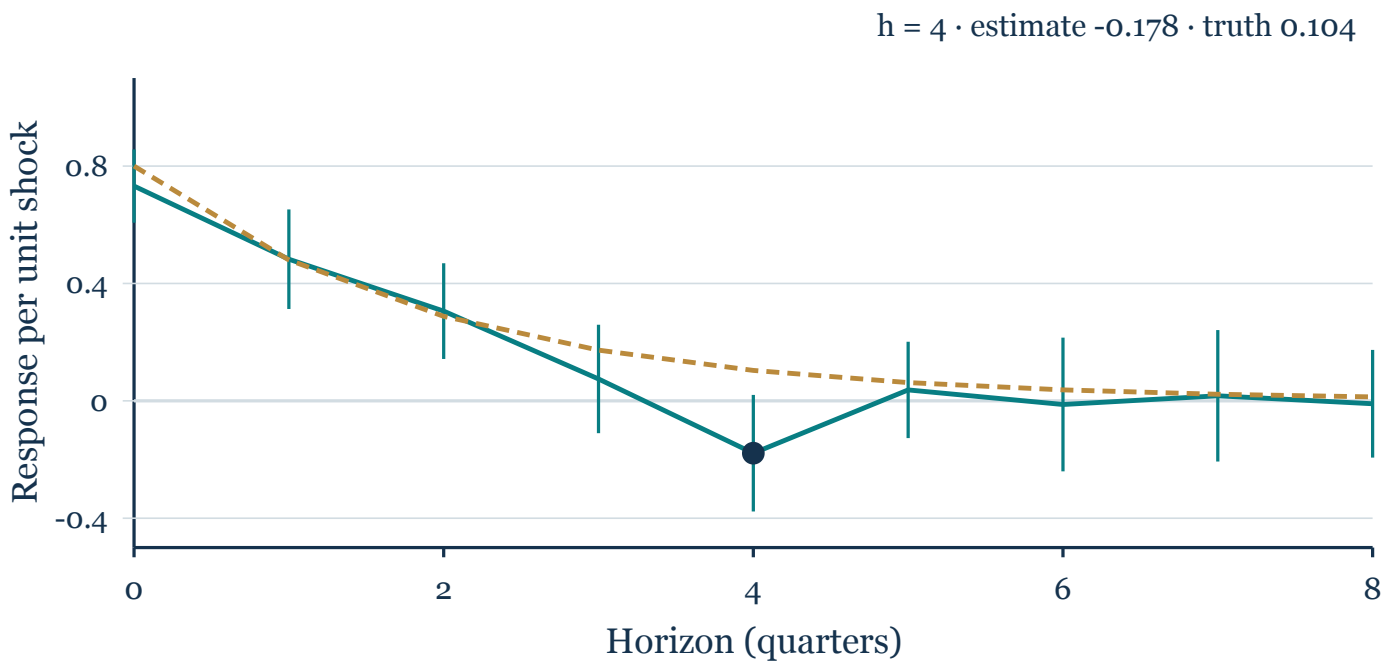
$$g(b) = 100\{\exp(b) - 1\}$$

$$g'(b) = 100\exp(b) \approx 110.02$$

$$SE\{g(\hat{b})\} \approx |g'(\hat{b})| SE(\hat{b}) \approx 2.08 \text{ percentage units}$$

Which change addresses the problem we found?

Compare improving dynamics, changing an error model and estimating a future response.



Constructed quarterly process; actual local-projection estimates with approximate pointwise 95% intervals. The dashed path is the known response.

Transform a serial-error equation step by step

Subtracting the error carry-over also transforms the regressors.

$$Y_t = a + \beta X_t + u_t, \quad u_t = \rho u_{t-1} + \varepsilon_t$$

$$Y_t - \rho Y_{t-1} = a(1-\rho) + \beta(X_t - \rho X_{t-1}) + \varepsilon_t$$

Make the transformation feasible

Prais–Winsten estimates the carry-over and retains the first observation.

> For $\rho = 0.5$: $12 - 0.5(10) = 7$, $3 - 0.5(2) = 2$

> First-observation weight: $\sqrt{1 - \rho^2}$

> Estimate rho, transform all columns, and fit the transformed equation.

Prais–Winsten models an AR(1) error process

Its additional structure can change the fit.

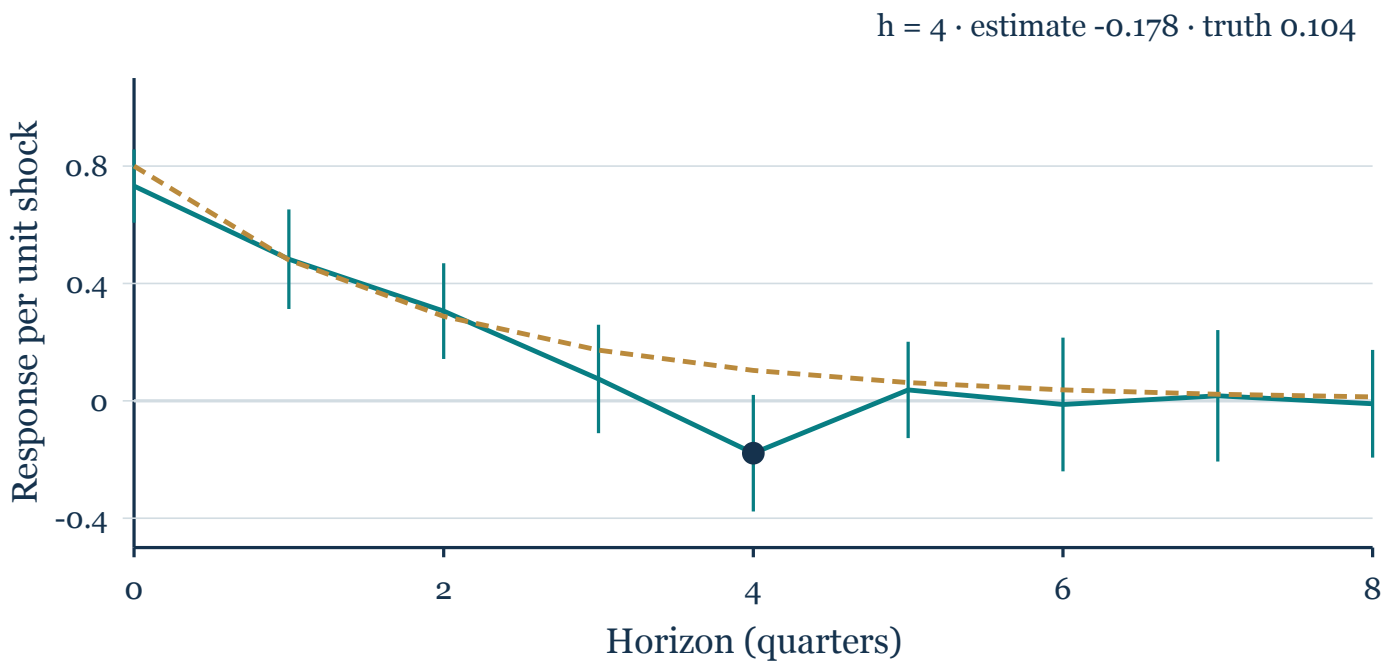
```
. prais logacc trend i.month unem spdlaw beltlaw wkends

------------------------------------------------------------------------------
      logacc | Coefficient  Std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
       trend |   .0009183    .000368     2.50   0.014     .0001874    .0016493
             |
        unem |   -.019975   .0047144    -4.24   0.000    -.0293397   -.0106104
      spdlaw |   -.045256   .0179614    -2.52   0.013     -.080934   -.0095779
     beltlaw |   .1008944   .0201055     5.02   0.000     .0609572    .1408315
      wkends |   .0042874   .0028062     1.53   0.130    -.0012867    .0098615
       _cons |   10.61796   .0628171   169.03   0.000     10.49318    10.74274
-------------+----------------------------------------------------------------
```

Calendar-indicator coefficient rows are omitted here. The output download contains the full numerical results.

Specify the known quarterly data-generating rule

A stable process gives a transparent response benchmark.



Constructed quarterly process; actual local-projection estimates with approximate pointwise 95% intervals. The dashed path is the known response.

- $\pi_t = 0.6\pi_{t-1} + 0.8s_t + \varepsilon_t$
- Hold other external draws fixed; allow inflation to propagate.
- Known intervention response: $0.8(0.6)^h$

A local projection asks about a specified future outcome

Define the shock and horizon before the coefficient.

> $$\pi_{t+h} = a_h + b_h s_t + c_h \pi_{t-1} + u_{t+h}^{(h)}$$

> h: quarters ahead; b_h: response per current shock unit.

The horizon coefficient has outcome-per-shock units

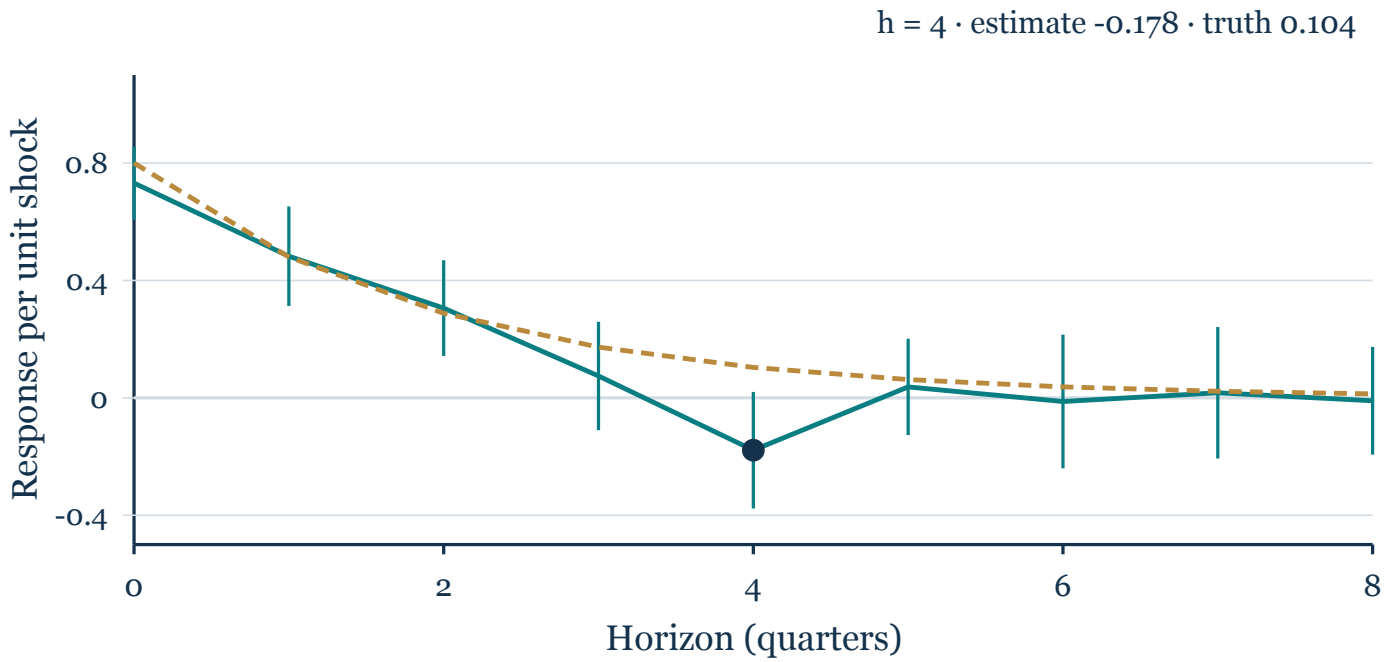
Do not confuse a future value with a cumulative sum.

> At horizon four, the coefficient concerns inflation four quarters after the shock.

> Write the economic sentence before calculating: a one-unit shock is associated with a fitted difference in the outcome at a specified future date, conditional on the included information.

Inspect the estimated path without forcing it to match the truth

Finite-sample response estimates can be noisy.



Constructed quarterly process; actual local-projection estimates with approximate pointwise 95% intervals. The dashed path is the known response.

Read the horizon-four output honestly

A negative sample estimate can coexist with a positive true response.

```
. newey F4.inflation shock L.inflation, lag(8)

------------------------------------------------------------------------------
             |             Newey–West
F4.inflation | Coefficient  std. err.      t    P>|t|     [95% conf. interval]
-------------+----------------------------------------------------------------
       shock |    -.17841   .1013114    -1.76   0.080    -.3780179    .0211979
             |
   inflation |
          L1. |   .0652363   .0710399     0.92   0.359    -.0747295    .2052021
             |
       _cons |  -.2784281   .1669179    -1.67   0.097    -.6072968    .0504406
------------------------------------------------------------------------------
```

Future-outcome errors can overlap across dates

A sensible mean model can still leave serial dependence.

> Illustration: $w_t = v_{t+1} + v_{t+2}$

> $w_{t+1} = v_{t+2} + v_{t+3}$: shared v_{t+2}

Pointwise intervals do not cover an entire response path at the same rate

Multiple horizons change the inferential question.

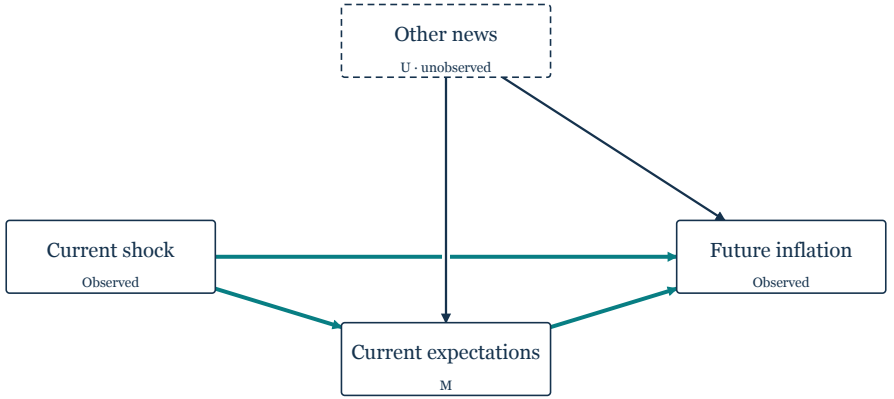
> Pointwise coverage: one stated horizon.

> Simultaneous coverage: the complete stated collection of horizons.

> Illustration only, with independent events: $0.95^9 \approx 0.63$.

Controlling a mechanism can change the response being studied

A current mediator is not just another lag control.



2 open directed routes · 0 open alternatives · gold: conditioned on

Assumed causal structure · dashed boxes: unobserved · gold boxes: conditioned on

> New application: expectations mediate a shock and also respond to other news.

> Post-shock controls can change the target and create selection.

Distinguish an estimated shock from a known innovation

Measurement and first-step uncertainty can matter.

> Known simulated shock: observed by construction.

> Estimated empirical shock: identification and first-step uncertainty need separate arguments.

Choose among the three changes using the application

The preferred response follows the diagnosed failure.

> For omitted predictable dynamics, reconsider the conditional-mean equation and its timing assumptions.

> None automatically solves policy endogeneity or shock mismeasurement.

Compare a known response with a noisy estimate

Reproduce the evidence, then explain its economic meaning.

Practice

Then use actual traffic.csv. Add one lag of log accidents to the monthly mean model, retain the calendar and policy controls, and use four-month HAC inference. Report the law coefficient and the stable-model long-run contrast. Explain why this observed application has no known historical causal truth supplied by the simulation. Keep your commands, log, output and interpretation.

Compare a known response with a noisy estimate: explanation

A reproducible result needs both a calculation and an interpretation.

> Constructed quarterly response: known truth and realised estimate differ.

> Real monthly fit: lag coefficient 0.332658; law impact 0.065212.

> Stable-model long-run log contrast 0.097719, approximate SE 0.019357.

From time-series methods to repeated units and instruments

> From time-series methods to repeated units and instruments

HAC estimates covariance of regressor-error scores across dates

Autocorrelation of errors and covariance of the coefficient estimates are connected by the regressors.

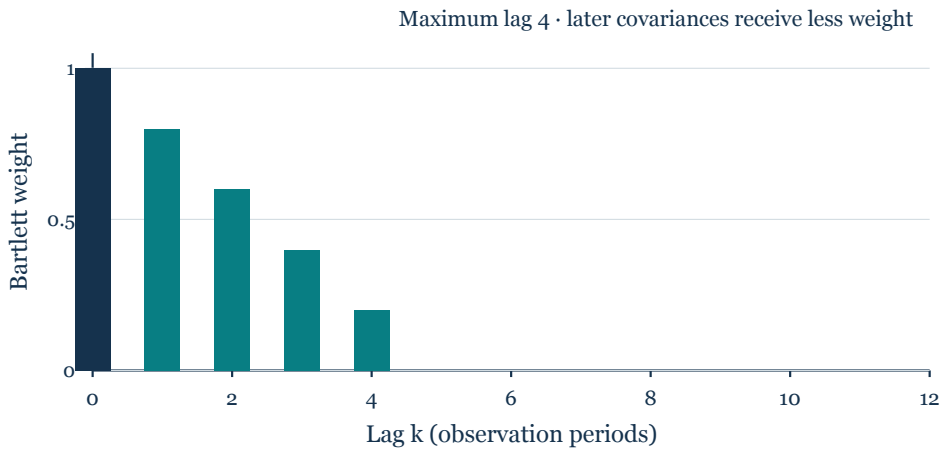
> $s_t = x_t u_t$, $S = \sum_{h=-\infty}^{\infty} E[s_t s_{t-h}']$, when the sum and limits are justified.

> Approximate coefficient covariance: $Q^{-1} S Q^{-1} / T$.

> Newey–West estimates a weighted, truncated version using fitted residuals.

Interpret the lag in calendar time

Interpret the selected maximum lag using months or quarters.



Declining lag weights only. The slider does not recalculate the Stata standard errors.

$$w_k = 1 - \frac{k}{L+1}, \quad k = 0, \ldots, L$$

Bartlett weights taper the contribution of later error covariances.

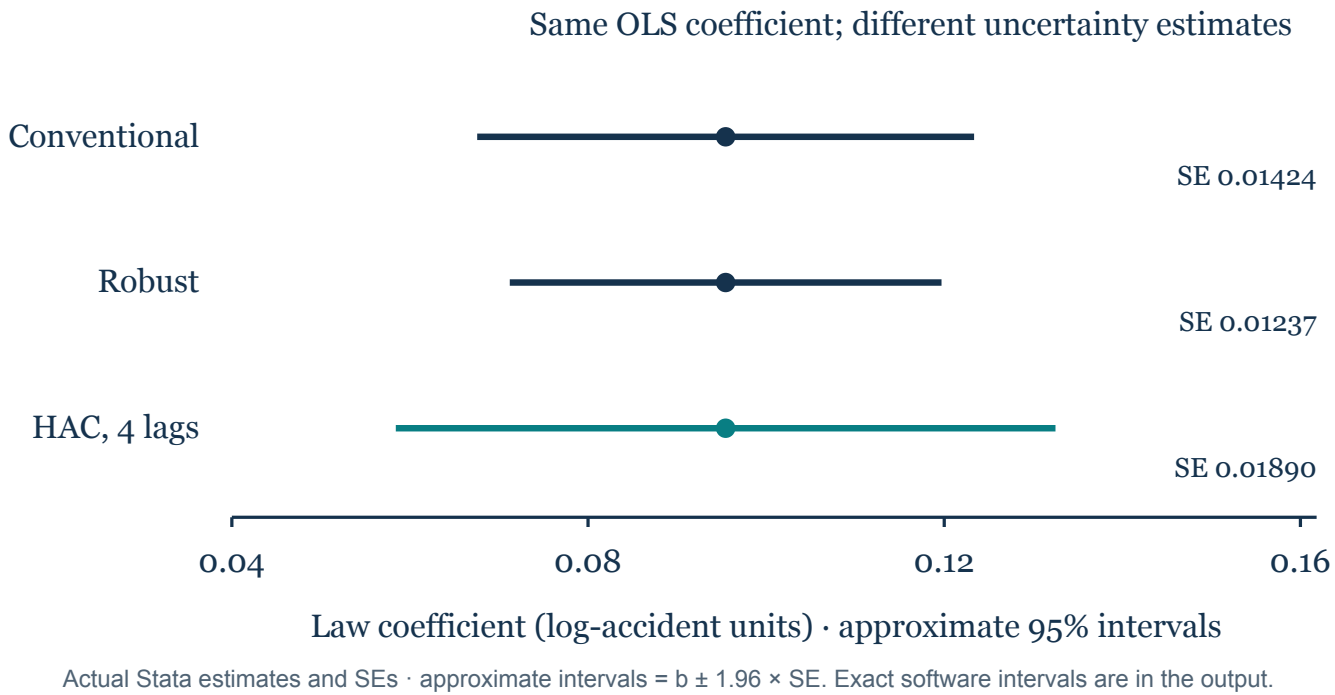
Practice

For quarterly data, what calendar window is lag(12)? Does that option add twelve regressors?

Explanation: It reaches twelve quarters, or three years. It adjusts the covariance calculation; it adds no conditional-mean regressors.

What stays fixed in the three fits?

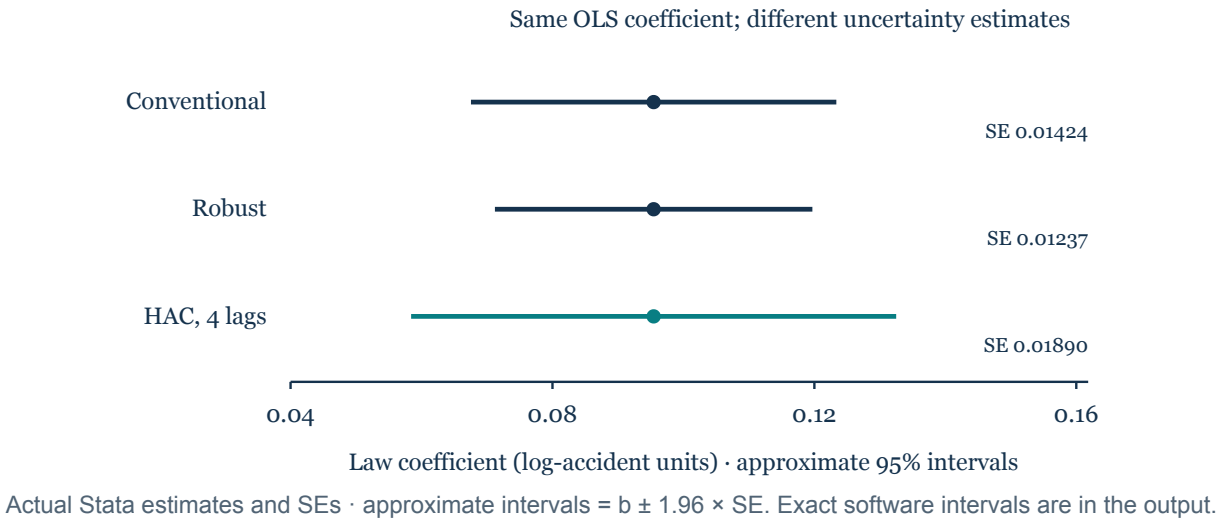
Hold the fitted model fixed and compare its three SE calculations.



Actual Stata estimates and SEs · approximate intervals = b ± 1.96 × SE. Exact software intervals are in the output.

Calculate, then explain the comparison

Calculate the test and interval, then report what they describe.



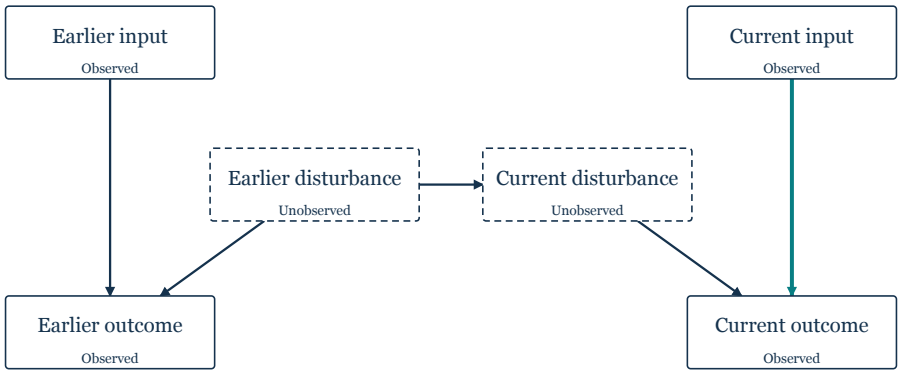
Practice

Use b=0.09545264 and HAC4 SE=0.01889849 to calculate the approximate t ratio and 95% interval for b.

Explanation: t≈5.05. The large-sample interval is 0.09545264±1.96×0.01889849≈[0.05841, 0.13249], on the log-outcome coefficient scale. This is uncertainty about the fitted coefficient under the inference assumptions; it is not automatically an interval for a causal law effect.

DAG practice: persistent disturbances and HAC

Trace the routes before choosing a control.



1 open directed route · 0 open alternatives · gold: conditioned on

Assumed causal structure · dashed boxes: unobserved · gold boxes: conditioned on

Practice

The earlier disturbance affects the current disturbance, but the input nodes have no common-cause connection to these disturbances in the stated graph. Does disturbance persistence alone create regressor endogeneity? What does HAC change, and which arrows does it remove?

DAG practice: persistent disturbances and HAC: explanation

Explain what conditioning changes in this stated graph.

> Persistence alone need not create endogeneity in this proposed model. It can create dependent outcomes and errors while the relevant input–disturbance mean conditions remain valid. Those are separate questions. HAC changes the estimated covariance of the coefficient under suitable exogeneity, dependence and sample conditions. It removes no causal arrows and keeps OLS coefficients fixed on an identical specification and sample. If the actual input selection process introduces a common-cause route to the disturbance, HAC does not close that new route.

Choose an uncertainty statement for repeated months

Reproduce the evidence, then explain its economic meaning.

Practice

Change the Newey–West bandwidth from four to eight months without changing the conditional-mean model or observations. Verify the identical coefficient and compare the SEs. Keep your commands, log, output and interpretation.

Choose an uncertainty statement for repeated months: explanation

A reproducible result needs both a calculation and an interpretation.

> Coefficient unchanged: 0.095453.

> HAC SE, four versus eight months: 0.018898 versus 0.021580.

> The bandwidth changes estimated uncertainty, not the conditional-mean regressors.

Exit check: connect the model, estimate and SE

Explain the three choices as responses to different empirical concerns.

> Same equation/sample → same OLS coefficient.

> Different justified covariance calculations → potentially different intervals.

> A policy effect still needs an assignment argument.

Practice

In one paragraph, compare adding a lag, using Newey–West and fitting Prais–Winsten. Interpret lag(4) for this dataset.

Explanation: Adding a lag changes the mean equation and may change the sample and coefficients. Newey–West keeps OLS coefficients on the same equation and sample, while changing their estimated covariance; lag(4) here reaches four months. Prais–Winsten uses an AR(1) error model to transform the fit and may change coefficients and SEs. None of these changes, by itself, creates an exogenous policy comparison.

Recall what HAC preserves

An uncertainty adjustment does not add predictors.

> Newey–West retains the OLS coefficient on the same equation and sample while changing its estimated covariance.

> Before comparing the results, say which concern motivates each operation.

The first observation needs its own GLS transformation

Quasi-differencing and ordinary first differences impose different error models.

> For $t \geq 2$: $Y_t - \rho Y_{t-1} = (1-\rho)\beta_0 + \beta_1(X_t - \rho X_{t-1}) + e_t$.

> Stationary first row: $(Y_1^\star, X_1^\star, 1^\star) = \sqrt{1-\rho^2}(Y_1, X_1, 1)$.

> Estimated rho and a misspecified error process require care.

Model or inference?

Try the comparison before opening the explanation.

Practice

Compare adding a lagged regressor, changing OLS SEs to HAC, and fitting Prais–Winsten. Which can change coefficients?

Model or inference?: explanation

Follow the calculation and interpret its units.

> Adding regressors changes the conditional-mean model and generally the OLS coefficients and sample. HAC on the identical specification and sample preserves OLS coefficients. Prais–Winsten transforms/weights the model under a modeled error process and can change coefficients as well as SEs; its assumptions must be justified.

A horizon

Try the comparison before opening the explanation.

Practice

In quarterly local projections $Y_{t+h} = a_h + b_h\, shock_t + \text{controls} + u_{t+h}$, what does b_4 describe?

A horizon: explanation

Follow the calculation and interpret its units.

> It describes the outcome association or response four quarters, one year, after a unit shock, conditional on the specification. A causal impulse-response interpretation requires valid shock identification and controls; the regression form alone does not provide it.

A HAC lag has calendar meaning

The same number describes different windows at different frequencies.

> The simulation uses quarterly data and Newey–West lag eight, allowing contributions through eight quarters, or two years.

> The appropriate window depends on the dependence structure and the asymptotic justification, and a finite-sample analysis should inspect plausible choices.

Mediation is not proven

Try the comparison before opening the explanation.

Practice

A shock coefficient becomes smaller after adding current expectations. Does this prove that expectations mediate the causal effect?

Mediation is not proven: explanation

Follow the calculation and interpret its units.

> No. It is consistent with a mediation story but may reflect confounding, post-treatment conditioning, endogeneity or other specification changes. State the target and the assumptions under which conditioning has a direct-effect interpretation.

Initial information and post-shock information have different roles

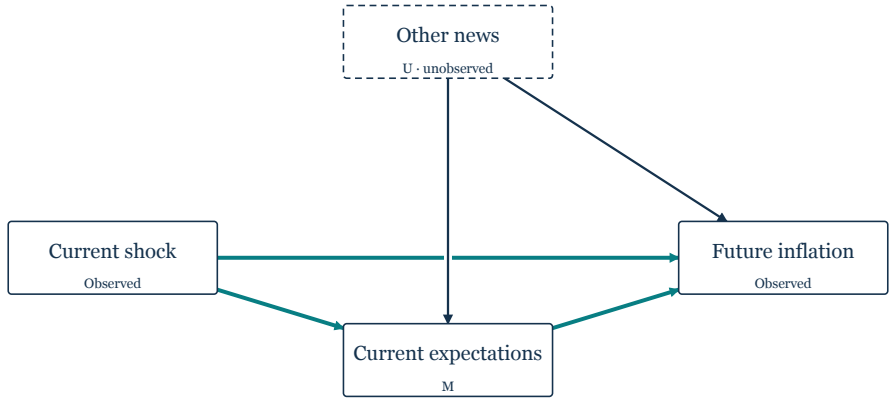
Timing determines what a control can do.

> A variable known before the shock can help account for the state of the economy that predicts both measured shocks and future outcomes.

> Write a time line for the shock, controls and outcome.

DAG practice: expectations after a shock

Trace the routes before choosing a control.



2 open directed routes · 0 open alternatives · gold: conditioned on

Assumed causal structure · dashed boxes: unobserved · gold boxes: conditioned on

Practice

A shock changes expectations and future inflation. Other news affects expectations and future inflation. Explain why adding current expectations can remove a mechanism and also create an alternative route, even if the shock initially had no common cause with the news.

DAG practice: expectations after a shock: explanation

Explain what conditioning changes in this stated graph.

> Expectations lie on $s_t \to M \to \pi_{t+h}$, so conditioning on them blocks part of the shock's response. They also form a collider on $s_t \to M \leftarrow U \to \pi_{t+h}$; conditioning can open the route through other news. The adjusted shock coefficient therefore does not automatically identify a direct response. An observed reduction in its size is compatible with several stories. State the target, the control's timing and the additional assumptions needed to separate mediation from selection or mediator–outcome confounding.