

Cointegration Testing and Dynamic Simulations of Autoregressive Distributed Lag Models^{*}

Soren Jordan[†]
Andrew Q. Philips[‡]

December 7, 2018
Forthcoming at *Stata Journal*

Abstract

In this paper we introduce `dynamac`, a suite of Stata programs designed to assist users in modeling and visualizing the effects of autoregressive distributed lag models, as well as testing for cointegration. We discuss the bounds cointegration test proposed by Pesaran, Shin and Smith (2001), which we have adapted into a Stata program. Since the resulting models can be dynamically complex, we follow the advice of Philips (2018) by introducing a flexible program designed to dynamically simulate and plot a variety of types of autoregressive distributed lag models, including error-correction models.

Word Count: 91 (abstract); approx. 6,900 (manuscript)

Keywords: cointegration; dynamic modeling; autoregressive distributed lag; error-correction

^{*}Source code for the programs discussed in this paper can be found at <http://andyphilips.github.io/dynamac/>. We wish to thank Lorena Barberia, Natália Moreira, Paul Kellstedt, Guy Whitten, Joe Ura, Eric Guntermann, the editor, and a reviewer for thoughtful comments and suggestions on various versions of these programs. Of course, any errors remain our own.

[†]sorenjordanpols@gmail.com. Assistant Professor, Department of Political Science, Auburn University, Auburn, AL 36849.

[‡]andrew.philips@colorado.edu. Assistant Professor, Department of Political Science, University of Colorado Boulder, UCB 333, Boulder, CO 80309.

Introduction

Time series models employing an autoregressive distributed lag (ARDL) are commonplace in the social sciences. Whether the dependent variable is estimated in levels (e.g., $y_t = \alpha_0 + \dots$) or in first differences (e.g., $\Delta y_t = \alpha_0 + \dots$), these models are able to test a host of theoretically-important theories, from the effect of public opinion on government response (Jennings and John 2009) to the effect of domestic and international factors on defense expenditures (Whitten and Williams 2011) or tax rates (Swank and Steinmo 2002), to analyzing dynamic changes in partisan responsiveness over time (Ura and Ellis 2008).

When employing an error-correction-style ARDL model, it becomes necessary to test for cointegration.¹ Philips (2018) shows that in small samples common in the social sciences—typically, when the number of time points is 80 or less—the ARDL bounds test for cointegration proposed by Pesaran, Shin and Smith (2001) tends to be more conservative (i.e., does not conclude cointegration when it does not exist) than either the popular Engle-Granger “two-step” (Engle and Granger 1987) or the Johansen (1991, 1995) approaches to cointegration testing. However, there is no standard implementation of this cointegration test in common statistical software.

In addition to its error-correction form, ARDL models in general may have complex dynamic specifications, including multiple lags, first-differences, and lagged first-differences. This makes it more difficult to interpret the effects of changes—especially short- and longer-run changes—in the independent variable. To mitigate this we introduce a flexible program that allows users to dynamically simulate a variety of ARDL models, including the error-correction model. Dynamic simulations offer an alternative to hypothesis testing of model coefficients by instead conveying the substantive significance of the results through meaningful counterfactual scenarios. Such an approach has been gaining popularity in the social sciences (e.g., Tomz, Wittenberg and King 2003; Imai, King and Lau 2009; Williams and Whitten 2011, 2012; Philips, Rutherford and Whitten 2016b).

¹In general, error-correction models regress the first-difference of the dependent variable on a constant, its own lag in levels, and the contemporaneous first-difference and lagged levels of each of the independent variables. For instance, in a model with a single independent variable x , we might estimate $\Delta y_t = \alpha_0 + \theta_0 y_{t-1} + \theta_1 x_{t-1} + \Delta x_t$.

Below, we offer a brief discussion of the ARDL-bounds approach to cointegration testing. We then present `dynamac`, a suite of two Stata commands for **dynamic ARDL modeling and cointegration testing**. This includes a command to provide the necessary critical values for the Pesaran, Shin, and Smith cointegration test (`pssbounds`), as well as a command to produce dynamic simulations of a multitude of ARDL-style models (`dynardl`).

The ARDL-Bounds Cointegration Test

The concept of cointegration has been around for several decades. To understand cointegration, we briefly discuss integrated versus stationary series. Time series may have “full-memory,” such that current realizations are fully a function of all previous stochastic shocks, plus some new innovation. Such series are said to be integrated of order one (or $I(1)$), a form of non-stationarity.² For instance, the series in Equation 1 is $I(1)$, since values at time t are a function of the prior value of y at time $t - 1$, plus innovation ϵ_t .³

$$y_t = y_{t-1} + \epsilon_t \tag{1}$$

Most of the time, $I(1)$ series cannot be included in standard regression models, since the nature of these data makes it more likely that we find statistically significant relationships simply due to random chance alone; this is often referred to as the “spurious regression problem in time series” (Yule 1926; Grant and Lebo 2016). Despite this, two or more $I(1)$ series may still have short-run, as well as long-run—or equilibrium—relationships between one another. That is to say, while short-run perturbations may move the series apart, over time this disequilibrium is corrected as the series move back towards a stable long-run relationship. Such series are said to be cointegrating.

Since not all relationships between $I(1)$ series are cointegrating, it becomes necessary

²Another term for an $I(1)$ series is that it contains a unit-root. Series that are stationary are said to be $I(0)$.

³This is commonly re-written as $(y_t - y_{t-1}) = \Delta y_t = \epsilon_t$.

to test for cointegration. The earliest test comes from Engle and Granger (1987), who show that cointegration between k I(1) regressors, $x_{1t}, x_{2t}, \dots, x_{kt}$, and an I(1) regressand, y_t , exists if the resulting residuals—from a regression of these variables entering into the equation in levels—is stationary:

$$y_t = \kappa_0 + \kappa_1 x_{1t} + \kappa_2 x_{2t} + \dots + \kappa_k x_{kt} + z_t \quad (2)$$

A number of other cointegration tests have since been proposed. Phillips (2018) offers an in-depth discussion of how to apply the Pesaran, Shin and Smith (2001) ARDL-bounds test for cointegration. Others include tests by Johansen (1991) and Phillips and Ouliaris (1990). However, the ARDL-bounds test offers several advantages. Chief among them is that users do not have to make the sharp I(0)/I(1) distinction for the regressors.⁴ Below, we briefly summarize this approach.⁵

First, the analyst must ensure that the dependent variable is I(1). There are many unit-root tests that can be used to determine the order of integration of a series, including the Dickey-Fuller, Phillips-Perron, Elliott-Rothenberg-Stock, and Kwiatkowski-Phillips-Schmidt-Shin tests, among others. Only an I(1) dependent variable is a potential candidate for cointegration.

Second, the analyst must ensure that the regressors are not of an order of integration higher than I(1). While this means that the analyst does not have to make the potentially-difficult I(0)/I(1) decision, they must ensure that all regressors are not explosive or contain seasonal unit roots.

Third, the analyst estimates an ARDL model in error-correction form. The model appears as follows:

⁴Users must ensure, however, that regressors are not of order I(2) or more, and that seasonality has been removed from the series.

⁵A more in-depth discussion can be found in Phillips (2018).

$$\Delta y_t = \alpha_0 + \theta_0 y_{t-1} + \theta_1 x_{1,t-1} + \dots + \theta_k x_{k,t-1} + \sum_{i=1}^p \alpha_i \Delta y_{t-1} + \sum_{j=0}^{q_1} \beta_{1j} \Delta x_{1,t-j} + \dots + \sum_{j=0}^{q_k} \beta_{kj} \Delta x_{k,t-j} + \varepsilon_t \quad (3)$$

Where the change in the dependent variable is a function of a constant, its value at $t - 1$ (appearing in levels), values at $t - 1$ of all regressors appearing in levels, as well as up to p and q_k lags of the first difference of the dependent variable and regressors, respectively. These may enter into Equation 3 for theoretical reasons, but also have the added benefit of helping to ensure white-noise residuals. While it is ideal to have well-behaved residuals in all models, this is a crucial step before running the ARDL-bounds test for cointegration. Information criteria such as SBIC and AIC, as well as autocorrelation and heteroskedasticity tests (e.g., Breusch-Godfrey, Durbin's Alternative, Cook-Weisberg, or Cumby-Huizinga tests) can be used to check for white-noise residuals. Many of these are available as canned Stata procedures.

While the ARDL-bounds test may be relatively easy to implement, the test uses special critical values. These critical values are available, but cumbersome for users to map onto the test; Pesaran, Shin and Smith (2001) provide asymptotic critical values, and Narayan (2005) provides finite samples critical values. In order to make the bounds test more accessible to users, below we introduce `pssbounds`, which provides the necessary critical values through an easy-to-use command.

pssbounds Syntax

```
pssbounds, observations( ) k( ) fstat( ) [tstat( ) case( )]
```

Options

`observations()` is the number of observations from the estimated ARDL model in error correction form. This is a required option.⁶

`k()` is the number of regressors, k , modeled in levels in the estimated ARDL model *not* including the lagged dependent variable.⁷ The bounds F-test is a test that the k parameters on the regressors appearing in levels (plus the coefficient on the lagged dependent variable, θ_0) are jointly equal to zero: $H_0 = \theta_0 + \theta_1 + \dots + \theta_k = 0$. This option is required, since critical values associated with the test differ based on the number of regressors.

`fstat()` is the value of the F-statistic from the test that all parameters on the regressors appearing in levels, plus the coefficient on the lagged dependent variable, are jointly equal to zero: $H_0 = \theta_0 + \theta_1 + \dots + \theta_k = 0$. After running the ARDL model in error-correction form, users should use Stata's `test` command to obtain the F-statistic.⁸ This option is required.

`case()` specifies the potential restrictions on the constant and trend terms, which in turn lead to different critical values for the bounds test. This option is not required, however. The default if this option is not specified—by far the most common—is an unrestricted intercept with no trend term: `case(3)`. Pesaran, Shin and Smith (2001) refer to this as “Case 3.”⁹ Other cases that are supported are:

- Case 1: No intercept and no trend, `case(1)`.
- Case 2: Restricted intercept and no trend, `case(2)`.

⁶As discussed below, if first using `dynardl` in error-correction form, users can simply run `pssbounds` after without having to specify the required options, since the latter program obtains the necessary stored values from the former.

⁷The lagged dependent variable is not included in the count of k , even though it is included in the restriction tested. The reason is that the lagged dependent variable is *always* present in the ARDL-bounds procedure. This is why special critical values are required. For instance if the ARDL model was: $\Delta y_t = \beta_0 - \theta_0 y_{t-1} + \beta_1 \Delta x_{1t} + \theta_1 x_{1,t-1} + \beta_2 \Delta x_{2t} + \theta_2 x_{2,t-1}$, then $k = 2$, since $x_{1,t-1}$ and $x_{2,t-1}$ appear in lagged levels.

⁸For example, `test 1.y 1.x1 1.x2... = 0`.

⁹These “cases” are alternatively referred to in both Arabic and Roman numerals in the above literature. We use Arabic numerals here, although `pssbounds` accepts both.

- Case 4: Unrestricted intercept and restricted trend, `case(4)`.
- Case 5: Unrestricted intercept and unrestricted trend, `case(5)`.

`tstat()` is the value of the t-statistic for the coefficient on the lagged dependent variable. This serves as a one-sided auxiliary test to the bounds F-test, and should supplement—but not replace—the conclusions offered by the F-statistic test above. If the tests are in conflict, or are inconclusive, users should follow the full procedure outlined in Philips (2018, 235). Only asymptotic critical values are available for the t-statistic, so this option is not required. Note that critical values do not currently exist for this test for Cases 2 and 4 (see below).

Examples

Two example applications of `pssbounds` are shown below. For the first example, we will use the Lutkepohl West German quarterly macroeconomic dataset available in Stata:

```
. webuse lutkepohl2
(Quarterly SA West German macro data, Bil DM, from Lutkepohl 1993 Table E.1)
. tsset
      time variable:  qtr, 1960q1 to 1982q4
              delta:  1 quarter
```

Following Philips (2018), the first step is to assess whether any of the three variables—log investment, log income, and log consumption—contain a unit root. Both the Phillips-Perron and Dickey-Fuller GLS unit root tests fail to reject the null hypothesis of a unit root for all series, as shown in Table 1:

```
. pperron ln_inv, lags(3)
. dfghs ln_inv, maxlag(4)

. pperron ln_inc, lags(3)
```

```
. dfgls ln_inc, maxlag(4)

. pperron ln_consump, lags(3)

. dfgls ln_consump, maxlag(4)
```

Table 1: Unit Root Tests Indicate I(1) Series

| | ln(Investment) | ln(Income) | ln(Consumption) |
|-----------------|----------------|------------|-----------------|
| Phillips-Perron | -1.25 | -2.01 | -1.54 |
| DF-GLS | -2.05 | -1.29 | -1.60 |

Note: * $p < 0.05$. Three augmenting lags included for all tests.

Next, we estimate the following ARDL-bounds model in error-correction form:¹⁰

$$\Delta \ln(\text{Investment})_t = \alpha_0 + \theta_0 \ln(\text{Investment})_{t-1} + \beta_1 \Delta \ln(\text{Income})_t + \theta_1 \ln(\text{Income})_{t-1} + \beta_2 \Delta \ln(\text{Consumption})_t + \theta_2 \ln(\text{Consumption})_{t-1} \quad (4)$$

In other words, we believe that investment is a function of income and consumption, and that there is a cointegrating relationship between investment and the two regressors. The results are shown in Model 1 in Table 2.

```
. regress d.ln_inv l.ln_inv d.ln_inc l.ln_inc d.ln_consump l.ln_consump
```

We then run an F-test that the coefficients on the variables appearing in lagged levels (i.e., \ln_inv_{t-1} , \ln_inc_{t-1} , and $\ln_consump_{t-1}$) are jointly equal to zero:

```
. test l.ln_inv l.ln_inc l.ln_consump

( 1)  L.ln_inv = 0
( 2)  L.ln_inc = 0
( 3)  L.ln_consump = 0

      F( 3, 85) = 2.60

      Prob > F = 0.0573
```

¹⁰Estimation is not the first step of the ARDL-bounds approach; we would need to conduct unit root test on the dependent variable, and ensure that the independent variables were I(1) or less (Philips 2018). We would also need to ensure that the resulting residuals from the model are white-noise. This is a simply a stylized example used to showcase the command using readily-available Stata datasets.

Table 2: Lutkepohl Example

| | (1) | (2) |
|-------------------------|---------------------|---------------------|
| | $\Delta \ln_inv_t$ | $\Delta \ln_inv_t$ |
| \ln_inv_{t-1} | -0.140* (0.059) | -0.152** (0.056) |
| $\Delta \ln_inc_t$ | -0.201 (0.459) | -0.0985 (0.423) |
| \ln_inc_{t-1} | -0.209 (0.334) | |
| $\Delta \ln_consump_t$ | 1.548** (0.538) | 1.395** (0.459) |
| $\ln_consump_{t-1}$ | 0.336 (0.333) | 0.131** (0.048) |
| Constant | -0.008 (0.071) | |
| N | 91 | 91 |
| R^2 | 0.17 | 0.27 |

Dependent variable is $\Delta \ln_inv_t$. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Recall that the critical values are non-standard, so we only need the value of the F-statistic, which is 2.60. In order to test for cointegration, we use `pssbounds`. For the required option `fstat`, we input the F-statistic from above. From the estimated model, we tell `pssbounds` that the number of observations is 91, the case is case 3 (i.e., unrestricted intercept with no trend, as shown in Model 1), and that there are two regressors appearing in levels ($k = 2$). The resulting output appears as follows:

```
. pssbounds, fstat(2.60) obs(91) case(3) k(2)
PESARAN, SHIN AND SMITH (2001) COINTEGRATION TEST
Obs: 91
No. Regressors (k): 2
Case: 3
-----
                                F-test
-----
<----- I(0) ----- I(1) ----->
10% critical value      3.170      4.140
```

| | | |
|-------------------|-------|-------|
| 5% critical value | 3.790 | 4.850 |
| 1% critical value | 5.150 | 6.360 |
| F-stat. = | 2.600 | |

F-statistic note: Asymptotic critical values used.

For this model, since the F-statistic of 2.60 is below the I(0) critical value—even at the 10 percent level—we can conclude that there is no cointegration, and that all regressors appearing in levels are stationary. As discussed in Philips (2018), we would next want to exclude one of the regressors from appearing in lagged levels—and thus, removing it from the potentially cointegrating equation—to see if either `ln_inc` or `ln_consump` on its own was in a cointegrating relationship with `ln_inv`.

Purely for illustrative purposes, let us now assume that we wanted to estimate a model without a constant. These results are shown in Model 2, Table 2. As with Model 1, we next run an F-test of all variables appearing in levels. For this example, note that the model also has `ln_inc` appearing in first differences, but not in levels; thus, $k = 1$:

```
. regress d.ln_inv l.ln_inv d.ln_inc d.ln_consump l.ln_consump, noconstant

. test l.ln_inv l.ln_consump
( 1)  L.ln_inv = 0
( 2)  L.ln_consump = 0

          F( 2, 87) = 3.83
          Prob > F = 0.0254
```

We can account for a restricted constant by specifying the `case(1)` option (i.e., no intercept and no trend). As an additional option, we add `tstat(-2.73)`, which is the t-statistic on the coefficient on the lagged dependent variable.

```
. pssbounds, fstat(3.83) obs(91) case(1) k(1) tstat(-2.73)
PESARAN, SHIN AND SMITH (2001) COINTEGRATION TEST
```

Obs: 91

No. Regressors (k): 1

Case: 1

F-test

| | <----- I(0) -----> | I(1) -----> |
|--------------------|--------------------|-------------|
| 10% critical value | 2.440 | 3.280 |
| 5% critical value | 3.150 | 4.110 |
| 1% critical value | 4.810 | 6.020 |

F-stat. = 3.830

t-test

| | <----- I(0) -----> | I(1) -----> |
|--------------------|--------------------|-------------|
| 10% critical value | -1.620 | -2.280 |
| 5% critical value | -1.950 | -2.600 |
| 1% critical value | -2.580 | -3.220 |

t-stat. = -2.730

F-statistic note: Asymptotic critical values used.

t-statistic note: Asymptotic critical values used.

The output of `pssbounds` now contains critical values for both the F-test and one sided t-test. Based on the F-statistic, we can conclude cointegration at the 10 percent level since the F-statistic of 3.83 is above the I(1) critical threshold of 3.28. However, there is not strong enough evidence to support cointegration at the five percent level. For the t-test, the t-statistic of -2.73 falls below the critical I(1) threshold of -2.60, supporting the earlier conclusion of cointegration. Also note that for both tests, `pssbounds` issued a

warning that asymptotic critical values are used. For all cases, only asymptotic critical values from Pesaran, Shin and Smith (2001) are provided for the t-statistic test.¹¹ Thus, interpreting the results of this test should be done with caution in small samples. Small sample critical values for the F-statistic are not available for case 1. Recall, though, that `pssbounds` only implements the procedure; when the results are inconclusive, users are encouraged to follow the process outlined in Philips (2018).

As a second example, we use data from Ura (2014), who examines public mood liberalism in the US. Ura argues that in the short-run, there will be a public “backlash” in response to liberal Supreme Court decisions, but that in the long-run the sentiments of the public tend to follow closely to those of the Court. Using the same dataset, Philips (2018, see supplemental materials p. 110) finds that the dependent variable, public mood liberalism, is $I(1)$, and that Ura’s ARDL model estimated in error-correction form with an additional lagged first-difference of unemployment produces the white-noise residuals needed to conduct the ARDL-bounds test. The model is shown in Table 3.

```
. use "supreme court mood replication.dta", clear
. tsset
      time variable:  year, 1955 to 2009
              delta:  1 unit

. regress d.mood l.mood d.policy l.policy d.unemployment dl.unemployment ///
l.unemployment d.inflation l.inflation d.caselaw l.caselaw
```

Next we run an F-test that all variables appearing in lagged levels ($mood_{t-1}$, $policy_{t-1}$, $unemployment_{t-1}$, $inflation_{t-1}$, and $caselaw_{t-1}$) are jointly equal to zero:

```
. test l.mood l.policy l.unemployment l.inflation l.caselaw
( 1)  L.mood = 0
( 2)  L.policy = 0
( 3)  L.unemployment = 0
```

¹¹Nor do critical values for the t-statistic test exist for cases 2 and 4.

Table 3: Ura (2014) Example

| | (1) |
|-----------------------------------|-----------------------|
| | Δmood_t |
| mood_{t-1} | -0.241** (0.076) |
| Δpolicy_t | 0.051 (0.069) |
| policy_{t-1} | -0.073*** (0.020) |
| $\Delta\text{unemployment}_t$ | -0.106 (0.265) |
| $\Delta\text{unemployment}_{t-1}$ | -0.538* (0.242) |
| $\text{unemployment}_{t-1}$ | -0.024 (0.198) |
| $\Delta\text{inflation}_t$ | -0.306* (0.123) |
| inflation_{t-1} | -0.299* (0.120) |
| $\Delta\text{caselaw}_t$ | -0.093* (0.037) |
| caselaw_{t-1} | 0.027* (0.011) |
| Constant | 15.65** (5.050) |
| N | 53 |
| R^2 | 0.47 |

Dependent variable is Δmood_t .

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

(4) L.inflation = 0

(5) L.caselaw = 0

F(5, 42) = 5.15
Prob > F = 0.0009

We then run `pssbounds` on the resulting F-statistic of 5.15, where $k = 4$ (i.e., four regressors: policy, unemployment, inflation, and caselaw), along with the value of the t-statistic of the lagged dependent variable, which is -3.19 :

```
. pssbounds, fstat(5.15) obs(53) case(3) k(4) tstat(-3.19)
```

PESARAN, SHIN AND SMITH (2001) COINTEGRATION TEST

Obs: 53

No. Regressors (k): 4

Case: 3

```
-----  
                                F-test  
-----  
                                <----- I(0) ----- I(1) ----->  
10% critical value           2.578           3.710  
5% critical value            3.068           4.334  
1% critical value            4.244           5.726  
  
F-stat. =           5.150
```

```
-----  
                                t-test  
-----  
                                <----- I(0) ----- I(1) ----->  
10% critical value           -2.570           -3.660  
5% critical value            -2.860           -3.990  
1% critical value            -3.430           -4.600
```

t-stat. = -3.190

F-statistic note:

t-statistic note: Small-sample critical values not provided
for Case III. Asymptotic critical values used.

We can conclude evidence of cointegration at the 5 percent level for the F-test, since the F-statistic of 5.15 is above the I(1) critical value of 4.334. However, note that we fail to clear the critical I(1) threshold for the ARDL-bounds t-test. However, the t-test values of this auxiliary test are asymptotic (`pssbounds` issues a note at the bottom of the output that warns the user of this) and not precisely tailored to small samples. Based on the small-sample F-statistics, we have relatively strong evidence of cointegration.

Dynamic Simulations of ARDL Models

ARDL models may have a fairly complex lag structure, with lags, contemporaneous values, first differences, and lagged first differences of the independent (and sometimes the dependent) variable appearing in the model specification. While interpreting short- and long-run effects may be simple in something like an ARDL(1,1) model (i.e., one lag of the dependent variable, and contemporaneous and one-period lags of all independent variables), understanding the short-, medium-, and long-run effects becomes difficult as the model specification grows in complexity.

To better interpret the substantive significance of our results, below we introduce `dynardl`, a command to dynamically simulate a variety of different ARDL models. `dynardl` estimates, simulates, stores the results from and automatically plots substantively interesting predictions from ARDL models. Users can even run `pssbounds` afterwards as a post-estimation command if simulating an error-correction model.

The output in `dynardl` helps us visualize the effect of a counterfactual change in

one regressor at a single point in time, holding all else equal, using stochastic simulation techniques. Dynamic simulation approaches are gaining in popularity as a straightforward way to show the substantive results of time series models, whose coefficients often have non-intuitive or “hidden” interpretations (Breunig and Busemeyer 2012; Williams and Whitten 2011; Philips, Rutherford and Whitten 2016*a,b*; Gandrud, Williams and Whitten 2016).¹² Before using the command, it is assumed that the user has already determined the order of integration of the variables through unit-root testing, diagnosed and addressed other issues such as seasonal unit roots, and used information criteria (and theory) to identify the best fitting lagged-difference structure, which is used to purge autocorrelation and to ensure the residuals are white noise. If an error-correction model is estimated, users should use the ARDL-bounds test to determine if there is cointegration, and if there is not, adjust the model accordingly.¹³

`dynardl` first runs a regression using OLS. Then, using a self-contained procedure similar to the popular Clarify program for Stata (Tomz, Wittenberg and King 2003), it takes 1000 draws (or however many simulations a user desires) of the vector of parameters from a multivariate normal distribution. These distributions are assumed to have means equal to the estimated parameters from the regression. The variance of these distributions is equal to the estimated variance-covariance matrix from the regression. In order to re-introduce stochastic uncertainty back into the model when creating predicted values, `dynardl` simulates $\hat{\sigma}^{2*}$ by taking draws from a scaled inverse χ^2 distribution. The distribution is scaled by the residual degrees of freedom (n-k), as well as the estimated $\hat{\sigma}^2$ from the regression (Gelman et al. 2014, pp. 43, 581), which ensures that draws of $\hat{\sigma}^{2*}$ are bounded by zero and one. Simulated parameters and sigma-squared values are then used to create predicted values of the dependent variable over time, \hat{Y}_t , for each of the simulations, by setting all covariates to certain values (typically means). Stochastic

¹²For instance, regressors in a stationary ARDL(1,0) model have both a contemporaneous effect (given by the coefficient on the regressor, β) as well as a long-run or cumulative effect, given by $\frac{\beta}{1-\theta_0}$.

¹³If the test is inconclusive, “Each regressor should be tested for a unit root. Only I(1) variables can appear in levels in the error correction model. Stationary variables may still appear in first differences...If the resulting statistic is still inconclusive, combinations of variables appearing in levels may need to be tested.” (Philips 2018, pp. 13-14). See Philips (2018) for a step-by-step example of this process for the ARDL model in general.

uncertainty is introduced into the prediction by taking a draw from a multivariate normal distribution with mean zero and variance $\hat{\sigma}^{2*}$. The program then averages across the simulations, creating \hat{Y}_t^* (the predicted values plus stochastic uncertainty) as well as percentile confidence intervals of the distribution of simulated values at a particular point in time. These are then saved, allowing a user to make a table or (more commonly) a graph of the results over time.

dynardl Syntax

`dynardl depvar indepvars, lags() shockvar() shockval() [options]`

The options below are required:

- `lags(numlist)` is a numeric list of the number of lags to include for each variable, separated by a comma. The number of desired lags is listed in exactly the same order in which the variables `depvar` and `indepvars` appear. For instance, the command:

$$\text{dynardl } y \text{ } x1 \text{ } x2, \text{ lags}(1, 2, 3)\dots$$

would lag `y` by $t-1$, `x1` by $t-2$, and `x2` by $t-3$. Note that the lag on `depvar` (always the first entry in `lags()`) must always be specified. To estimate a model without a lag for a particular variable, simply replace the number with a “.”; for instance, if we did not want a lag on the first regressor, and wanted a lag of $t-1$ on the second regressor, we type: `lags(1, ., 1)`. `dynardl` can accommodate consecutive lags by specifying the minimum lag, a forward slash, followed by the maximum lag. For instance, `lags(1/3, ., .)` will introduce lags of y_t at $t-1$, $t-2$, and $t-3$ into the model. The program can also add non-consecutive lags. For instance, to add a single lag of y_t at $t-1$ and $t-3$, specify `lags(1 3, ., .)`.

- `shockvar(varname)` is a single independent variable from the list of `indepvars` that is to be shocked. It will experience a counterfactual shock of size `shockval(#)` at time `time(#)`.

- `shockval(#)` is the amount to shock `shockvar(varname)` by. A common shock value is a +/- one standard deviation shock, although any shock value can be used.

The following options are not required:

- `diffs(numlist)` is a numeric list of the number of contemporaneous first differences to include for each variable, separated by a comma. Note that the first entry (the placeholder for the `depvar`) will always be empty (denoted by “.”), since the first difference of the dependent variable cannot appear on the right-hand side of the model.¹⁴ Only first-differences can be taken using this option; for instance, `diffs(., 1, .)` would first difference only the first regressor in the equation.
- `lagdiffs(numlist)` is a numeric list of the number of lagged first differences to include for each variable, separated by a comma. The syntax is similar as that of `lags()`. For instance, to include a lagged first difference at $t - 2$ for `depvar` (i.e., $\Delta y_{t-2} = y_{t-2} - y_{t-3}$), a lagged first difference at $t - 1$ for the first regressor, and none for the second, specify `lagdiff(2, 1, .)`. To include an additional lagged first difference for both the first and second lags of `depvar`, specify `lagdiff(1/2, 1, .)`. Users can also include non-consecutive lagged first differences.¹⁵
- `levels(numlist)` is a numeric list of variables to appear in levels (i.e., not lagged or differenced but appearing contemporaneously at time t), separated by a comma.¹⁶ For example, `levels(., 1, .)` tells `dynard1` to include the first regressor contemporaneously at time t .
- `ec` if specified, `depvar` will be estimated in first differences. If estimating an error correction model, users will need to use this option.
- `trend` if specified, the program will add a deterministic linear trend to the model.

¹⁴It can however, appear in *lagged* first differences, as shown below.

¹⁵For example, `lagdiff(1 3/4, ., .)` would add a first difference at $t - 1$, $t - 3$, and $t - 4$ for the dependent variable.

¹⁶If both `levels()` and `ec` are specified, `dynard1` will issue a warning message. Of course, users may have a valid reason to include a variable in levels; for instance, a dummy variable.

- `noconstant` if specified, the constant will be suppressed.
- `range(#)` is the length of the scenario to simulate. By default, this is $t = 20$. Note that the range must be larger than `time()`.
- `sig(#)` specifies the significance level for the percentile confidence intervals. The default is for 95% confidence intervals.
- `time(#)` is the scenario time in which the shock occurs to `shockvar()`. The default time is $t = 10$.
- `saving(string)` specifies the name of the output file that contains the predictions. If no filename is specified, the program will save the results as “dynardl_results.dta”.
- `forceset(numlist)` by default, the program will estimate the ARDL model in equilibrium; all lagged variables and variables appearing in levels are set to their sample means. All first differences and any lagged first differences are set to zero. This option allows the user to change the setting of the lagged (or unlagged if using `levels()`) levels of the variables. This could be useful when estimating a dummy variable; for instance, when we wish to see the effect of a movement from zero to one.
- `sims(#)` is the number of simulations (default is 1000). If confidence intervals are particularly noisy, it may help to increase this number. Note that you may also need to increase the `matsize` in Stata.
- `burnin(#)` allows `dynardl` to run a number of early simulations that are dropped, resulting in more stable starting values. This option is rarely used. However, if using the option `forceset()`, the predicted values will not be in equilibrium at the start of the simulation, and will take some time to converge on stable values. To get around this, one can use the `burnin` option to specify a number of simulations to “throw away” at the start. By default, this is 20. Setting a burn-in does not change the simulation range or time; to simulate a range of 25 with a shock time at 10 and a burn-in of 30, specify: `burnin(30) range(25) time(10)`.

- **graph** although **dynardl** saves the means of the predicted values and user-specified confidence intervals in **saving**, users can use this option to automatically plot the dynamic results using a spikeplot. Two alternative plots are possible:
 - By adding the option **rarea**, the program will automatically create an area plot. Predicted means along with 75, 90, and 95 percent confidence intervals are shown with this option.
 - By adding the option **change**, predicted changes (from the sample mean) are shown across time, starting with the time at which the shock occurs; similar to an impulse response function.
- **expectedval** by default, **dynardl** will calculate predicted values of the dependent variable for a given number of simulations. For every simulation, the predicted value comes from a systematic component—which contains uncertainty surrounding the parameter estimates from the model—as well as a single draw from the stochastic component. With the **expectedval** option, the program instead calculates expected values of the dependent variable such that the average of 1000 stochastic draws now becomes the estimate of the stochastic component for each of the simulations. This effectively removes the stochastic uncertainty introduced in calculating \hat{Y}_t^* (Tomz, Wittenberg and King 2003). Because of this, predicted values are more conservative than expected values. Note that **dynardl** takes longer to run if calculating expected values. We recommend that unless the user has a specific theoretical or substantive justification for using expected values, they instead use the default predicted values that take into account the random uncertainty surrounding the predictions.

Examples

For the first example we will once again use the results from Model 1 in Table 2 using the Lutkepohl data. We estimate the following model:

$$\Delta \ln_inv_t = \ln_inv_{t-1} + \Delta \ln_inc_t + \ln_inc_{t-1} + \Delta \ln_consump_t + \ln_consump_{t-1} \quad (5)$$

Since `dynardl` uses Stata's matrix capabilities, we will increase the maximum matrix size as well:

```
. set matsize 5000
. webuse lutkepohl2
(Quarterly SA West German macro data, Bil DM, from Lutkepohl 1993 Table E.1
. tsset
      time variable:  qtr, 1960q1 to 1982q4
      delta:  1 quarter
```

To estimate the model shown in Equation 5 using `dynardl`, and see the effect of a -1 shock to `ln_inc` (about two standard deviations) we specify the command below. In Equation 5 we have two regressors and the lagged dependent variable, all of which are lagged one period. Therefore, we specify `lags(1, 1, 1)`. The first difference of all regressors appear, so we add `diff(., 1, 1)`. There are no lagged differences appearing in Equation 5.

```
dynardl ln_inv ln_inc ln_consump,   ///
lags(1, 1, 1) diffs(., 1, 1)   ///
shockvar(ln_inc) shockval(-1)   ///
time(10) range(30) graph ec
```

In the command, `lags(1, 1, 1)` tells `dynardl` to add lags (of $t - 1$) for each of the variables, while `diffs(., 1, 1)` means that the second and third variables (`ln_inc` and `ln_consump`) enter into the model as first-differences as well. In `shockvar()` we include the variable to be shocked, and specify the amount to shock it by using `shockval`. Additional options include the time at which the shock occurs, `time(10)`, the total range of the simulations, `range(30)`, and that the dependent variable is to be included in first-differences, `ec`. Last, since we specified the `graph` option, `dynardl` will produce a plot, which is shown in Figure 1. As is clear from the figure, a -1 shock at $t = 10$ produces a small increase that is not statistically significant in the short-run, which eventually increases to a predicted value of about 7.5 over the long-run, an increase that is statistically significant.

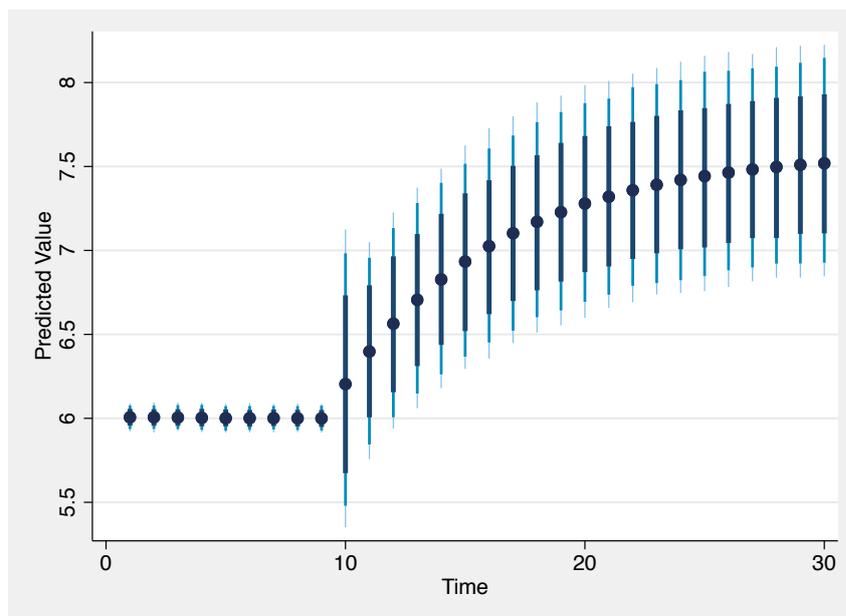


Figure 1: Plot Produced from `dynardl` Using the `graph` Option

Note: Dots show average predicted value. Shaded lines show (from darkest to lightest) the 75, 90, and 95 percent confidence intervals.

In addition to producing figures, `dynardl` also saved the prediction output, which can be used to create more customizable figures (e.g., colors, lines, labels) if users desire. Since we did not specify a filename to save as using the `saving()` option, the results are automatically saved as “`dynardl_results.dta`”.

More complex dynamic specifications are possible using `dynardl`. For instance, perhaps we wanted to estimate the following equation:

$$\begin{aligned} \Delta \ln_inv_t = & \ln_inv_{t-1} + \Delta \ln_inv_{t-1} + \ln_inc_{t-1} + \ln_inc_{t-2} + \ln_inc_{t-3} + \\ & \Delta \ln_consump_t + \ln_consump_{t-1} \end{aligned} \quad (6)$$

Now, the dependent variable appears not only in lagged-level form at $t - 1$, but also as a lagged first-difference. `ln_inc` no longer appears contemporaneously but at the first, second, and third lags, while `ln_consump` continues to appear in lagged and first-differenced forms. Equation 6 would appear as the following in `dynardl`:

```
dynardl ln_inv ln_inc ln_consump, ///
```

```
lags(1, 1/3, 1) diffs(., ., 1) lagdiffs(1, ., .)   ///
shockvar(ln_inc) shockval(-1)   ///
time(10) range(30) graph ec rarea sims(5000)
```

Note that, since the first through third lag of `ln_inc` are desired, we specify `lags(1, 1/3, 1)`.¹⁷ Since there is a lagged first-difference included for the dependent variable, we add `lagdiffs(1, ., .)`. The “.” are placeholders for the other variables and must be included. Last, we add the `rarea` option to produce an area plot, and increase the number of simulations to 5000 with `sims(5000)`. The resulting plot is shown in Figure 2. As is clear from the figure, as a result of a negative shock to `ln(Income)` at time $t = 10$, investment decreases over the next few periods, though this does not appear to be statistically significantly lower than the average predicted value (shown when $t < 10$). After about three periods, investment increases in response to the negative income shock, resulting in a new equilibrium prediction just above 10. Figure 2 is appealing since it shows 75, 90 and 95 percent confidence intervals, and since it is an area plot, its continuous style may allow users to see changes slightly easier than the plot style in Figure 1. Note too that since we increased the number of simulations to 5000, the confidence intervals in Figure 2 are more smooth from time point to time point than in Figure 1.

`dynardl` can also be used with `pssbounds` when estimating an error correction model. For instance, directly after estimating Equation 5, we can run `pssbounds` as a post-estimation command to test for cointegration, without having to specify any additional options; the program automatically obtains the necessary values for the F- and t- statistics, the number of observations, the case, and the number of regressors, k :

```
. dynardl ln_inv ln_inc ln_consump,           ///
lags(1, 1, 1) diffs(., 1, 1)                 ///
shockvar(ln_inc) shockval(-1)               ///
```

¹⁷Equivalently, we could specify `lags(1, 1 2 3, 1)` instead of `lags(1, 1/3, 1)`.

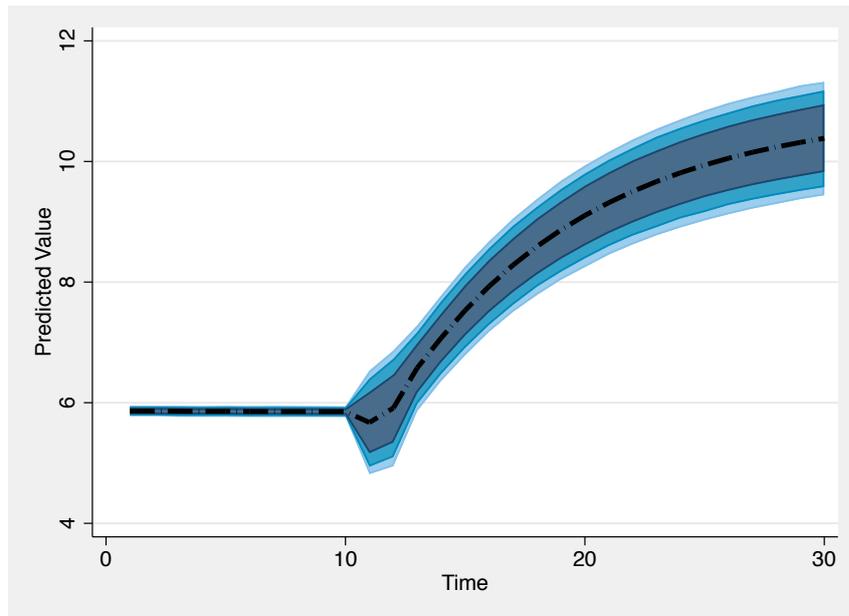


Figure 2: Plot Produced from `dynard1` Using the `rarea` Option

Note: Black dotted line shows average predicted value. Shaded area shows (from darkest to lightest) the 75, 90, and 95 percent confidence intervals.

```
time(10) range(30) ec
```

```
. pssbounds
```

```
PESARAN, SHIN AND SMITH (2001) COINTEGRATION TEST
```

```
Obs: 91
```

```
No. Regressors (k): 2
```

```
Case: 3
```

```
-----
```

F-test

```
-----
```

```

          <----- I(0) ----- I(1) ----->
10% critical value      3.170      4.140
5% critical value       3.790      4.850
1% critical value       5.150      6.360

```

```
F-stat. =      2.602
```

```
-----
```

t-test

```

-----
                <----- I(0) ----- I(1) ----->
10% critical value    -2.570          -3.210
5% critical value     -2.860          -3.530
1% critical value     -3.430          -4.100

t-stat. =           -2.372
-----

```

F-statistic note: Asymptotic critical values used.

t-statistic note: Asymptotic critical values used.

In addition to error-correction style models, `dynardl` can handle ARDL models where the dependent variable is estimated in levels.¹⁸ For instance, perhaps we wanted to estimate the following ARDL(1,1) model:

$$\ln_inv_t = \ln_inv_{t-1} + \ln_inc_t + \ln_inc_{t-1} + \ln_consump_t + \ln_consump_{t-1} \quad (7)$$

In `dynardl`, we add the `levels(., 1, 1)` to let the program know that the two independent variables are to appear contemporaneously in levels. If we wanted to see the effect of a change from `ln_inc=6` to `ln_inc=5`, while holding `ln_consump` constant at 7, we can also use the `forceset()` option to force the program to evaluate the simulations at these values, not the sample means (by default).¹⁹

```

dynardl ln_inv ln_inc ln_consump,   ///
lags(1, 1, 1) levels(., 1, 1) forceset(., 6, 7)   ///
shockvar(ln_inc) shockval(-1)   ///
time(10) range(30) graph change sims(5000)

```

¹⁸Although this is a stylized example, users should always first perform unit root tests in order to determine the most appropriate model specification.

¹⁹Note that while we can set some or all of the independent variables using `forceset`, the lagged dependent variable cannot be forced to a fixed value.

Since we added the `change` option, the resulting plot is akin to an impulse response function, as shown in Figure 3. In other words, we are looking not at the level of the predicted value, but the difference between the predictions at each point in time relative to the average predicted value before the shock. Figure 3 shows the change in predicted value, starting when the shock occurs.²⁰ As is clear from the figure, there is no statistically significant change in the predicted value in the short-run as a result of the shock. However, over the long run the change is statistically significant at the 90 percent level of confidence.

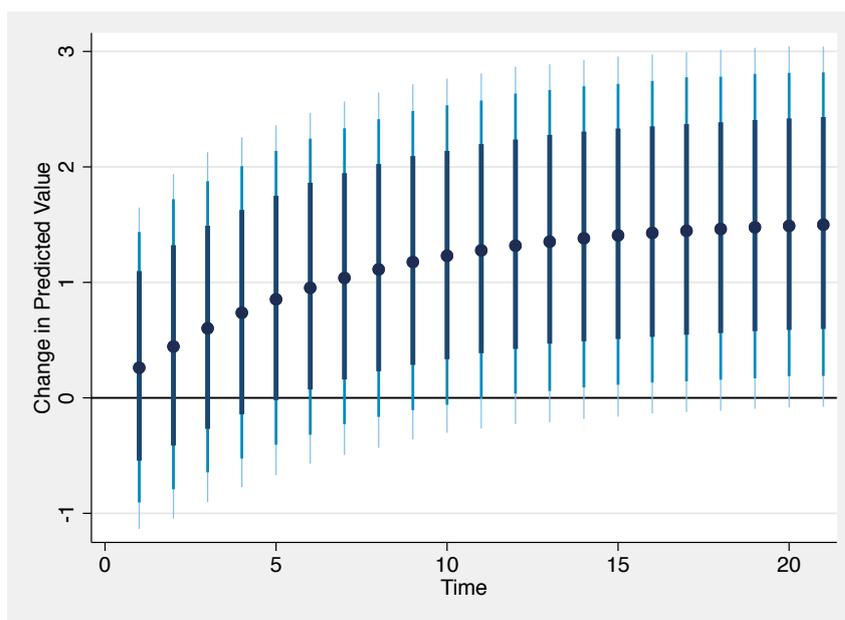


Figure 3: Plot Produced from `dynardl` Using the `change` Option

Note: Dots show mean change in predicted value from sample mean. Shaded area shows (from darkest to lightest) the 75, 90, and 95 percent confidence intervals.

Conclusion

In this paper we have introduced `dynamac`, a suite of programs for dynamic autoregressive distributed lag modeling and cointegration testing. This is comprised of two programs designed to assist time series analysts. `pssbounds` helps users test for cointegration by

²⁰In other words, an analyst might want to show the effect of a shock without showing the periods preceding it (like the first ten periods of Figure 2). In Figure 3, since the shock occurred at $t = 10$, and the total range of the simulation was $t = 30$, the graph shows a total range of $t = 20$.

providing critical values from Pesaran, Shin and Smith (2001) and Narayan (2005) automatically in a tabular format. The program `dynard1` helps users dynamically simulate a variety of autoregressive distributed lag models in order to gain a better understanding of the substantive significance of their results. Users can then graph or save their simulated predicted values for use elsewhere. Both programs make it easier for users to test and interpret their dynamic models.

References

- Breunig, Christian and Marius R Busemeyer. 2012. “Fiscal austerity and the trade-off between public investment and social spending.” *Journal of European Public Policy* 19(6):921–938.
- Engle, Robert F and Clive WJ Granger. 1987. “Co-integration and error correction: representation, estimation, and testing.” *Econometrica* 55(2):251–276.
- Gandrud, Christopher, Laron K Williams and Guy D Whitten. 2016. “dynsim: Dynamic simulations of autoregressive relationships.” *R package version 1.2.2* .
- Gelman, Andrew, John B Carlin, Hal S Stern and Donald B Rubin. 2014. *Bayesian data analysis*. Vol. 2 Chapman and Hall/CRC Boca Raton, FL, USA.
- Grant, Taylor and Matthew J. Lebo. 2016. “Error correction methods with political time series.” *Political Analysis* 24:3–30.
- Imai, Kosuke, Gary King and Olivia Lau. 2009. “Zelig: Everyone’s statistical software.” *R package version 3(5)*.
- Jennings, Will and Peter John. 2009. “The dynamics of political attention: Public opinion and the Queen’s Speech in the United Kingdom.” *American Journal of Political Science* 53(4):838–854.
- Johansen, Søren. 1991. “Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models.” *Econometrica: Journal of the Econometric Society* 59(6):1551–1580.
- Johansen, Soren. 1995. *Likelihood-based inference in cointegrated vector autoregressive models*. Oxford University Press.
- Narayan, Paresh Kumar. 2005. “The saving and investment nexus for China: Evidence from cointegration tests.” *Applied Economics* 37(17):1979–1990.

- Pesaran, M Hashem, Yongcheol Shin and Richard J Smith. 2001. "Bounds testing approaches to the analysis of level relationships." *Journal of Applied Econometrics* 16(3):289–326.
- Philips, Andrew Q. 2018. "Have your cake and eat it too? Cointegration and dynamic inference from autoregressive distributed lag models." *American Journal of Political Science* 62(1):230–244.
- Philips, Andrew Q, Amanda Rutherford and Guy D Whitten. 2016a. "Dynamic pie: A strategy for modeling trade-offs in compositional variables over time." *American Journal of Political Science* 60(1):268–283.
- Philips, Andrew Q, Amanda Rutherford and Guy D Whitten. 2016b. "dynsimpie: A command to examine dynamic compositional dependent variables." *Stata Journal* 16(3):662–677.
- Phillips, Peter CB and Sam Ouliaris. 1990. "Asymptotic properties of residual based tests for cointegration." *Econometrica: Journal of the Econometric Society* pp. 165–193.
- Swank, Duane and Sven Steinmo. 2002. "The new political economy of taxation in advanced capitalist democracies." *American Journal of Political Science* pp. 642–655.
- Tomz, Michael, Jason Wittenberg and Gary King. 2003. "CLARIFY: Software for interpreting and presenting statistical results." *Journal of Statistical Software* 8(1):1–30.
- Ura, Joseph Daniel. 2014. "Backlash and legitimation: Macro political responses to supreme court decisions." *American Journal of Political Science* 58(1):110–126.
- Ura, Joseph Daniel and Christopher R Ellis. 2008. "Income, preferences, and the dynamics of policy responsiveness." *PS: Political Science & Politics* 41(4):785–794.
- Whitten, Guy D and Laron K Williams. 2011. "Buttery guns and welfare hawks: The politics of defense spending in advanced industrial democracies." *American Journal of Political Science* 55(1):117–134.

- Williams, Laron K and Guy D Whitten. 2011. "Dynamic simulations of autoregressive relationships." *Stata Journal* 11(4):577–588.
- Williams, Laron K and Guy D Whitten. 2012. "But wait, there's more! Maximizing substantive inferences from TSCS models." *The Journal of Politics* 74(03):685–693.
- Yule, G Udny. 1926. "Why do we sometimes get nonsense-correlations between Time-Series?—a study in sampling and the nature of time-series." *Journal of the royal statistical society* 89(1):1–63.