



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*



ardl: Estimating autoregressive distributed lag and equilibrium correction models

Sebastian Kripfganz
University of Exeter Business School
Exeter, U.K.
S.Kripfganz@exeter.ac.uk

Daniel C. Schneider
Max Planck Institute for Demographic Research
Rostock, Germany
schneider@demogr.mpg.de

Abstract. We present a command, `ardl`, for the estimation of autoregressive distributed lag (ARDL) models in a time-series context. The `ardl` command can be used to fit an ARDL model with the optimal number of autoregressive and distributed lags based on the Akaike or Bayesian (Schwarz) information criterion. The regression results can be displayed in the ARDL levels form or in the error-correction representation of the model. The latter separates long-run and short-run effects and is available in two different parameterizations of the long-run (cointegrating) relationship. The popular bounds-testing procedure for the existence of a long-run levels relationship is implemented as a postestimation feature. Comprehensive critical values and approximate p -values obtained from response-surface regressions facilitate statistical inference.

Keywords: `st0734`, `ardl`, `ardl` postestimation, autoregressive distributed lag model, error-correction model, bounds test, long-run relationship, cointegration, time-series data

1 Introduction

Real-world phenomena are often characterized by complex relationships. Some observed variables might exhibit erratic behavior in the short run but tend to comove in a stable and predictable way over longer time horizons. Attempting to empirically uncover such long-run equilibrium relationships is tantamount to separating them from the overlaid short-run dynamics. This separation allows one to find evidence for or against an equilibrium relationship, which is often at the heart of a research question. It also allows analysis of the short-term fluctuations around the equilibrium, which can be valuable in its own right, for example, when conducting forecasting exercises or dynamic simulations.

When we observe the variables of interest over a sufficiently long stretch of consecutive time periods, multiequation vector autoregressive (VAR) and vector error-correction (VEC) models are commonly used to assess their dynamic relationships. When we have

reasons to assume that there is a natural ordering of the variables such that there is no contemporaneous feedback from a response variable to the other variables in the system, a single-equation autoregressive distributed lag (ARDL) model can simplify the analysis and facilitate more efficient inference.¹

ARDL models have many possible applications. They are extensively used in studies analyzing linkages of pollution and energy consumption to economic growth (Fatai, Oxley, and Scrimgeour [2004]; Narayan and Smyth [2005]; Wolde-Rufael [2006]; Ang [2007]; Halicioglu [2009]; Jalil and Mahmud [2009]; Zhang et al. [2015]; Ntanos et al. [2018]; Bekun, Emir, and Sarkodie [2019]; Kirikkaleli, Güngör, and Adebayo [2022]; and many more). Relationships with economic growth have also been investigated for foreign direct investment and trade (Oteng-Abayie and Frimpong 2006; Belloumi 2014), infrastructure (Fedderke, Perkins, and Luiz 2006), immigration (Morley 2006), tourism (Katircioglu 2009; Wang 2009; Song et al. 2011), stock market development (Enisan and Olufisayo 2009), and health expenditures (Murthy and Okunade 2016).

Other examples include the nexus between viral infections and meteorological factors (He et al. 2017; Doğan et al. 2020), childcare availability, fertility, and female labor force participation (Lee and Lee 2014), wages, productivity, and unemployment (Pesaran, Shin, and Smith 2001), savings and investment (Narayan 2005), exchange rates and trade (Bahmani-Oskooee and Brooks 1999; De Vita and Abbott 2004), exchange rates and monetary policy (Frankel, Schmukler, and Servén 2004; Shambaugh 2004; Obstfeld, Shambaugh, and Taylor 2005), financial development and inequality (Ang 2010), bank lending and property prices (Davis and Zhu 2011), financial reforms and credit growth (Adeleye et al. 2018), stock market efficiency and fiscal policy (Stoian and Iorgulescu 2020), democracy and the shadow economy (Esaku 2022), and the interdependencies among stock price indices and commodity prices (Narayan, Smyth, and Nandha 2004; Sari, Hammoudeh, and Soytas 2010; Büyüksahin and Robe 2014), as well as cryptocurrencies (Ciaian, Rajcaniova, and Kancs 2016, 2018), to list only a few.

Recently, the ARDL methodological toolkit was used extensively to analyze adjustment processes during the COVID-19 pandemic, including tourism demand forecasts (Zhang et al. 2021) and the effects on macroeconomic activity (Varona and Gonzales 2021) or energy consumption (Aruga, Islam, and Jannat 2020).

The ARDL model can be conveniently reparameterized in so-called error-correction (EC) form, which disentangles the long-run relationship from the short-run dynamics. When the variables are nonstationary—to be precise, integrated of order 1—the long-run relationship embedded in an EC model corresponds to a cointegrating relationship (Engle and Granger 1987; Hassler and Wolters 2006). Testing for cointegration in such a setup therefore equals testing for the existence of a long-run relationship. However, the latter concept retains its relevance when some of or all the variables are stationary.

Pesaran and Shin (1998) and Hassler and Wolters (2006) highlight some advantages of the ARDL approach over alternative strategies for cointegration analysis—such as the Engle and Granger (1987) two-step procedure implemented in the community-

1. Occasionally, the abbreviation ADL is used in the literature instead of ARDL.

contributed command **egranger** (Schaffer 2010) or the Phillips and Hansen (1990) fully modified ordinary least-squares approach implemented in **cointreg** (Wang 2012). First of all, it can accommodate a mixture of stationary and nonstationary variables without the need for pretesting the order of integration. Moreover, the short-run and long-run coefficients can be consistently estimated in one step, and the estimator's asymptotic normality eases statistical inference.²

Compared with a system-based Johansen (1995) cointegration analysis, which is implemented in Stata's **vec** command suite, the single-equation approach can be more efficient if the focus is on one outcome variable, in addition to the aforementioned flexibility regarding the integration orders. However, in the ARDL framework, the outcome variable is not allowed to simultaneously determine the long-run equilibrium of other explanatory variables, which would cause an endogeneity problem. The VAR or VEC approach can be more suitable for impulse-response analysis or dynamic forecasts because feedback from the dependent variable to the weakly exogenous variables is explicitly modeled.

Despite its advantages, testing for the existence of a long-run (cointegrating) relationship with the ARDL framework still requires a bit of effort. The test statistic has a nonstandard distribution that depends on various characteristics of the model and the data, including the integration order of the variables. Pesaran, Shin, and Smith (2001) propose a “bounds test”, which involves comparing the values of conventional F and t statistics with pairs of critical values (CV). Outside these bounds, the test either conclusively rejects or does not reject the null hypothesis. Within the bounds, the test is inconclusive.

This bounds test is implemented as a postestimation feature in our **ardl** package for the estimation of single-equation ARDL and EC models. Improved CV bounds and approximate p -values have been obtained by Kripfganz and Schneider (2020) with response-surface regressions using billions of simulated test statistics. These CVs are more precise and exhaustive than earlier ones tabulated by Pesaran, Shin, and Smith (2001) and Narayan (2005). A key feature of **ardl** is the automatic selection of the optimal lag order with the Akaike information criterion (AIC) or Bayesian (Schwarz) information criterion (BIC). With an increasing number of independent variables, the number of candidate models—which are characterized by all possible combinations of lag orders—is quickly in the tens or even hundreds of thousands. A computationally efficient implementation of this procedure ensures that the optimal model is still found within seconds.

Closely related, Jordan and Philips (2018) recently introduced the **dynardl** command for dynamic simulations of ARDL models. Their **pssbounds** command also provides an interface to display the original Pesaran, Shin, and Smith (2001) and Narayan (2005) asymptotic and finite-sample CVs for the bounds test. As argued above, those

2. Shin, Yu, and Greenwood-Nimmo (2014) extend the ARDL framework by introducing nonlinearities that allow for asymmetric long-run effects. Such a nonlinear ARDL model can be fit in Stata using the command **nardl**, implemented by M. Sunder (<http://www.marco-sunder.de/stata/>). Here we restrict our attention to the symmetric case.

CVs are now largely superseded. Moreover, their commands do not perform an automatic lag-order selection, which is a key feature of our `ardl` command. Once the optimal model specification is obtained with the `ardl` command, the `dynardl` command can still be a useful complement if a visualization of the dynamic effects is desired.

This article is concerned only with time-series data. For the estimation of ARDL models in a large- T panel-data context, see the community-contributed commands `xtpmg` (Blackburne and Frank 2007), `xtdcce2` (Ditzen 2018, 2021), and `xtivdfreg` (Kripfganz and Sarafidis 2021). The command `xtwest` (Persyn and Westerlund 2008) enables cointegration tests based on panel-data EC models.

In section 2, we outline the econometric background for the ARDL approach to the analysis of long-run equilibrium relationships, and we provide guidance for the model specification and bounds-testing procedure. In sections 3 and 4, we describe the syntax and options for the `ardl` package. In section 5, we illustrate the approach with an empirical example from the realm of cryptocurrencies. Section 5 concludes.

2 Econometric model and methods

2.1 ARDL model

Suppose we expect the existence of an equilibrium relationship between an outcome variable y_t and a set of K explanatory variables $\mathbf{x}_t = (x_{1t}, x_{2t}, \dots, x_{Kt})'$:

$$y_t = b_0 + b_1 t + \mathbf{x}_t' \boldsymbol{\theta} + e_t \quad (1)$$

b_0 is the intercept of the regression line, and b_1 is the slope coefficient of a linear time trend. The data are observed at consecutive time points $t = 1, 2, \dots, T$. Estimating the regression coefficients in such a static model by ordinary least squares (OLS) might result in spuriously large coefficient estimates even if there is no underlying relationship among the variables. This is known to happen when the error term e_t is nonstationary because of the nonstationarity of y_t and \mathbf{x}_t (after accounting for the possibility of a deterministic time trend).

Equation (1) remains a valid regression model if y_t and some of or all the variables \mathbf{x}_t are cointegrated, that is, when y_t and \mathbf{x}_t are individually integrated of order 1, $I(1)$, but there exists a linear combination among them such that e_t is integrated of order 0, $I(0)$. Equation (1) reflects a conditional long-run equilibrium relationship—if it exists—to which a process reverts over time. In the short run, the process might divert from this equilibrium, but the above equation is silent about the dynamic evolution of the process when it is off the equilibrium path. Such deviations are transitory, and the elements in the data-generating process (DGP) governing them are therefore $I(0)$. These neglected $I(0)$ components in the DGP affect the finite-sample (and possibly the asymptotic) distributions of test statistics and thus invalidate conventional hypothesis tests and regression diagnostics.³

To circumvent the problems associated with fitting a static model, we can augment the regression equation with lags of the dependent and independent variables. We can even include another set of L exogenous variables \mathbf{z}_t , which may have predictive power to explain the short-term fluctuations of y_t but do not affect its equilibrium path. We assume that all variables in \mathbf{z}_t are (trend) stationary. Augmenting the model in this way aims at obtaining a dynamically complete model in which the regression error term u_t is free of serial correlation:

$$y_t = c_0 + c_1 t + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=0}^q \beta'_i \mathbf{x}_{t-i} + \gamma' \mathbf{z}_t + u_t \quad (2)$$

$t = 1 + p^*, \dots, T$. Leaving aside the variables \mathbf{z}_t , this is a general ARDL (p, q, \dots, q) model with intercept c_0 , linear trend $c_1 t$, and lag orders $p \in [1, p^*]$ and $q \in [0, p^*]$.⁴ To ensure that there are enough degrees of freedom available to fit the model's coefficients with sufficient precision, we may need to choose the maximum admissible lag order p^* conservatively. This is especially relevant when the number of observations in the dataset (T) is relatively small, the number of variables in \mathbf{x}_t (K) is relatively large, or both.⁵

Given the initial observations y_1, y_2, \dots, y_{p^*} and the time paths of \mathbf{x}_t and \mathbf{z}_t , (2) describes the dynamic evolution of y_t over time, irrespective of whether an equilibrium relationship—as postulated in (1)—exists. The intercept c_0 and the linear time trend $c_1 t$ may or may not be included in the model, depending on the nature of the variables under consideration.⁶ We assume that enough lags have been included in the ARDL model (2) to purge the error term from any remaining serial correlation and to ensure that the

3. We provide an extended introduction to these topics in the working-paper version of this article (Kripfganz and Schneider 2022). For further background reading, interested readers are referred to any textbook on time-series econometrics of their choice.

4. Allowing for different lag orders among the components of \mathbf{x}_t is straightforward and can be treated as a special case of the general model by restricting some coefficients to be zero.

5. The importance of the maximum lag order p^* is explained further below. In practice, the data frequency often guides this choice. For instance, it is customary to allow for up to 12 lags with monthly data and up to 4 or 8 lags with quarterly data.

6. In general, other deterministic model components—such as quadratic time trends or impulse dummy variables—can be added. We abstract from them here but note that their inclusion may affect the applicability of the CVs for the bounds test, which is presented further down in this article.

variables \mathbf{x}_t are weakly exogenous/long-run forcing—ruling out any contemporaneous feedback from y_t to \mathbf{x}_t . If there exists a stable long-run relationship, conventional asymptotic theory can be applied for statistical inference on any of the coefficients even if some of the variables are nonstationary (Pesaran and Shin 1998). This highlights the importance of testing for the existence of such a long-run relationship, which we consider in section 2.3.

While the inclusion of further lags improves the regression fit, this comes at the cost of a higher variance of the coefficient estimates. To balance this tradeoff, we can base a data-driven approach to optimal lag selection on the AIC or the BIC,

$$\begin{aligned}\text{AIC} &= -2 \ln(\mathcal{L}) + 2K^* \\ \text{BIC} &= -2 \ln(\mathcal{L}) + \ln(T^*)K^*\end{aligned}$$

where $\ln(\mathcal{L})$ is the value of the log-likelihood function from the fitted regression model, $T^* = T - p^*$ is the effective sample size, and $K^* = 2 + p + K(q + 1) + L$ is the number of estimated coefficients in (2). These criteria balance the desire for a better fit of the model—higher values of $\ln(\mathcal{L})$ —against the temptation of creating ever larger models. The BIC has a larger penalty term than the AIC (for $T^* \geq 8$) and therefore tends to select more parsimonious models. The optimal lag orders are then found by fitting model (2) for all possible combinations of p and q and choosing the model that minimizes the AIC or BIC.

For the comparability of the model-selection criteria, we must base all regressions on the same estimation sample. This is the reason for initially choosing a fixed maximum lag order p^* . When both p and q are smaller than p^* , the estimation of model (2) does not use all the available observations. This is the price we need to pay for consulting the model-selection criteria. Once the optimal lag orders p and q have been found, we can subsequently refit the model, utilizing all available observations by setting $p^* = \max(p, q)$.

2.2 Error-correction representation

To gain a better interpretability of the model's coefficients, we can reformulate the ARDL model in EC representation (Hassler and Wolters 2006):⁷

$$\Delta y_t = c_0 + c_1 t - \alpha(y_{t-1} - \theta \mathbf{x}_{t-1}) + \sum_{i=1}^{p-1} \psi_{yi} \Delta y_{t-i} + \omega' \Delta \mathbf{x}_t + \sum_{i=1}^{q-1} \psi'_{xi} \Delta \mathbf{x}_{t-i} + \gamma' \mathbf{z}_t + u_t \quad (3)$$

7. By convention, the summations evaluate to zero if the upper limit is smaller than the lower limit.

The coefficients in (3) can be mapped in a straightforward algebraic way to the coefficients in (2):

$$\alpha = 1 - \sum_{i=1}^p \phi_i, \quad \boldsymbol{\theta} = \frac{\sum_{j=0}^q \beta_j}{\alpha}$$

$$\psi_{yi} = - \sum_{j=i+1}^p \phi_j, \quad \boldsymbol{\omega} = \beta_0, \quad \psi_{xi} = - \sum_{j=i+1}^q \beta_j$$

Now recall the hypothesized long-run equilibrium relationship between y_t and \mathbf{x}_t in (1). Ignoring the intercept and linear time trend for the moment, we see the deviations from this equilibrium, $e_{t-1} = y_{t-1} - \boldsymbol{\theta}'\mathbf{x}_{t-1}$, can be found again in the EC model (3). Because of the nonlinear interaction between the coefficients α and $\boldsymbol{\theta}$, we cannot directly fit (3) with OLS. However, given the mapping above, we can recover consistent estimates of all coefficients from the ARDL model (2). Yet a computationally more convenient approach is to instead fit the following model:⁸

$$\Delta y_t = c_0 + c_1 t + \pi_y y_{t-1} + \boldsymbol{\pi}_x' \mathbf{x}_{t-1} + \sum_{i=1}^{p-1} \psi_{yi} \Delta y_{t-i} + \boldsymbol{\omega}' \Delta \mathbf{x}_t + \sum_{i=1}^{q-1} \psi'_{xi} \Delta \mathbf{x}_{t-i} + \boldsymbol{\gamma}' \mathbf{z}_t + u_t \quad (4)$$

From the above model, we can easily recover the so-called speed-of-adjustment coefficient $\alpha = -\pi_y$ and the long-run coefficients $\boldsymbol{\theta} = \boldsymbol{\pi}_x / \alpha$. The corresponding standard errors can be computed with the delta method (Pesaran and Shin 1998). Note that (4) collapses to the well-known augmented Dickey and Fuller (1979) regression for unit-root testing when no explanatory variables \mathbf{x}_t and \mathbf{z}_t are present ($K = L = 0$).

The speed-of-adjustment coefficient α tells us how fast the process for y_t reverts to its long-run relationship when this equilibrium is distorted. $\alpha = 1$ would imply that—in the absence of any other short-run fluctuations—any deviation from the equilibrium is fully corrected immediately in the period after the distortion occurs. In contrast, $\alpha = 0$ would imply that the process never returns to its equilibrium path. Values of α between these two boundaries reflect a partial-adjustment process, where the gap to the equilibrium is gradually closed over time.⁹

Clearly, $\boldsymbol{\theta} \neq \mathbf{0}$ is not a sufficient condition for the existence of a conditional long-run relationship between the levels of y_t and \mathbf{x}_t . When $\alpha = 0$, then y_t is $I(1)$ and no such relationship exists. In the opposite scenario, when $\boldsymbol{\theta} = \mathbf{0}$ and $\alpha \in (0, 2)$, then y_t is (trend) stationary, irrespective of the integration order of the components in \mathbf{x}_t . For a long-run level relationship to exist, we need both $\boldsymbol{\theta} \neq \mathbf{0}$ and $\alpha > 0$. In this case—as long as the elements of \mathbf{x}_t are not cointegrated among themselves—the integration properties

8. When called with the option `ec1`, the `ardl` command estimates (4) but reports the coefficients for (3).

9. While we allow α to fall into the interval $[0, 2)$ in the following, we do not pay particular attention to the oscillating or overshooting case $\alpha > 1$ in this article. We also rule out explosive processes, which result under $\alpha < 0$. An estimate of α outside the reasonable region $[0, 1]$ should be seen as a warning signal for potential model misspecification.

of \mathbf{x}_t determine the integration order of y_t . If the variables in \mathbf{x}_t with nonzero long-run coefficient are $I(1)$, then y_t is $I(1)$ as well, and the conditional long-run relationship corresponds to a cointegrating relationship.

In this context, note that the assumption of \mathbf{x}_t being long-run forcing for y_t implies that there can exist at most one cointegrating relationship that involves y_t . Consider the VEC model

$$\begin{pmatrix} \Delta y_t \\ \Delta \mathbf{x}_t \end{pmatrix} = \mathbf{a}_0 + \mathbf{a}_1 t + \begin{pmatrix} \pi_{yy} & \pi'_{yx} \\ \pi_{xy} & \Pi_{xx} \end{pmatrix} \begin{pmatrix} y_{t-1} \\ \mathbf{x}_{t-1} \end{pmatrix} + \sum_{i=1}^{p^*-1} \Psi_i \begin{pmatrix} \Delta y_{t-i} \\ \Delta \mathbf{x}_{t-i} \end{pmatrix} + \Gamma \mathbf{z}_t + \begin{pmatrix} \varepsilon_{y,t} \\ \varepsilon_{x,t} \end{pmatrix} \quad (5)$$

\mathbf{x}_t is long-run forcing for y_t if it obeys the restriction $\pi_{xy} = \mathbf{0}$; that is, there is no level effect of y_{t-1} on $\Delta \mathbf{x}_t$.¹⁰ This does not rule out further cointegrating relationships among the elements of \mathbf{x}_t , $\Pi_{xx} \neq \mathbf{0}$. Thus, without further inspection, a cointegration rank larger than one for the entire system $(y_t, \mathbf{x}'_t)'$ does not necessarily imply a violation of this assumption. However, if there is reason to suspect multiple cointegrating relationships involving y_t , $\pi_{xy} \neq \mathbf{0}$, then a single-equation ARDL or EC model is inappropriate.¹¹ Instead, this would call for a multivariate cointegration analysis within the Johansen (1995) framework by fitting a VAR or VEC model.¹² In contrast, if $\pi_{xy} = \mathbf{0}$ is indeed satisfied and the interest is primarily on the long-run relationship between y_t and \mathbf{x}_t , then fitting a single-equation model is more efficient and computationally straightforward.

The remaining coefficients ψ_{yi} , ω , ψ_{xi} , and γ in (3) capture the short-run dynamics that are not prescribed by the equilibrium-reverting forces.¹³ They not only are relevant for making dynamic forecasts but also play a role for choosing appropriate CVs when testing for the existence of a long-run relationship, which we explore in section 2.3.

A complication arises if $q = 0$ for some of or all the long-run forcing variables. In that situation, $\pi_x = \omega$, which implies that the corresponding variance-covariance matrix of the coefficient estimates in (4) is rank deficient. To avoid this complication, we can equivalently formulate the EC representation with the levels of the long-run forcing variables expressed in period t instead of $t - 1$:

$$\Delta y_t = c_0 + c_1 t - \alpha(y_{t-1} - \theta \mathbf{x}_t) + \sum_{i=1}^{p-1} \psi_{yi} \Delta y_{t-i} + \sum_{i=0}^{q-1} \psi'_{xi} \Delta \mathbf{x}_{t-i} + \gamma' \mathbf{z}_t + u_t \quad (6)$$

10. See Pesaran, Shin, and Smith (2001) for a detailed discussion of the relationship between the single-equation EC model (3) and the VEC model (5).

11. Consequently, it is not permissible to run several ARDL regressions involving the variables $(y_t, \mathbf{x}'_t)'$, in which the dependent variable of one regression becomes an independent variable in other regressions.

12. See `var`, `vec`, and related Stata commands.

13. Strictly speaking, the error correction governed by the coefficient α is a short-run adjustment as well.

It has the same parameter restrictions as defined above. Note that $\omega' \Delta \mathbf{x}_t$ is replaced by $\psi'_{x0} \Delta \mathbf{x}_t$. The interpretation of the long-run coefficients θ does not change because the time subscript does not matter when the process is in equilibrium. The equation to be estimated in this case becomes

$$\Delta y_t = c_0 + c_1 t + \pi_y y_{t-1} + \pi_x \mathbf{x}_t + \sum_{i=1}^{p-1} \psi_{yi} \Delta y_{t-i} + \sum_{i=0}^{q-1} \psi'_{xi} \Delta \mathbf{x}_{t-i} + \gamma' \mathbf{z}_t + u_t \quad (7)$$

where the coefficients π_x are identical to the corresponding coefficients in (4), despite the change in the time subscript.¹⁴

2.3 Bounds test

Although we can consistently estimate all coefficients in the ARDL model (2) or its EC representations, testing for the existence of a long-run relationship involves a bit more effort. This is because the process for y_t contains a unit root under the null hypothesis of no long-run relationship; therefore, the test statistics have nonstandard distributions. Moreover, the tests depend on the choice of deterministic model components. In the ARDL model (2)—and its EC representations (3) and (6)—we have allowed for an intercept c_0 and a linear time trend $c_1 t$. We can distinguish the following five cases:

1. No deterministic model components are included ($c_0 = c_1 = 0$).
2. A restricted intercept is included ($c_0 = \alpha b_0$) but no time trend ($c_1 = 0$).
3. An unrestricted intercept is included ($c_0 \neq 0$) but no time trend ($c_1 = 0$).
4. An unrestricted intercept is included ($c_0 \neq 0$) and a restricted time trend ($c_1 = \alpha b_1$).
5. Both deterministic model components are unrestricted ($c_0 \neq 0$ and $c_1 \neq 0$).

A decision about the relevant case can often be guided by a visual inspection of the time series. Cases 1 and 2 are in line with a process y_t , which could reasonably be an $I(1)$ process without drift under the null hypothesis of no long-run level relationship. Under the alternative hypothesis, y_t would either be $I(0)$ or cointegrated with \mathbf{x}_t . Case 1 is most appropriate if y_t and \mathbf{x}_t fluctuate around a zero mean or if any nonzero means cancel out in the long-run level relationship; that is, $b_0 = b_1 = 0$ in (1). The latter condition is hard to verify ex ante, such that case 2 is often the safer option whenever some variables have a nonzero mean.

14. When called with the option `ec`, the `ardl` command estimates (7) but reports the coefficients for (6). As an aside, when option `ec1` is specified—which normally commands time subscripts $t-1$ for the long-run forcing variables—a subscript t is used in the estimation equation for those variables whose lag order is $q = 0$. However, the reported results are still reparameterized as in (3), incorporating the constraint on ω .

If y_t appears to be trending, it could be an $I(1)$ process with drift under the null hypothesis. This calls for case 3 or 4. Under the alternative hypothesis, y_t would either be trend stationary or cointegrated with \mathbf{x}_t . Case 3 is most appropriate if the trend in y_t is entirely attributable to a trend in \mathbf{x}_t ; that is, $b_1 = 0$ in (1). Again, this may be difficult to justify ex ante. Despite the fact that case 3 is most commonly applied in the empirical practice, case 4 is generally the safer option when there is insufficient knowledge about the source of the observed time trend.

Especially when the sample size is relatively small, it might be difficult to distinguish visually between a mildly drifting unit-root process under the null hypothesis and a stationary process which is fluctuating around a constant mean under the alternative hypothesis. This can be another relevant situation for case 3. Similarly, case 5 could be used to statistically discriminate between a unit-root process with faster—although hardly noticeable—than linear growth (or decline) and a trend-stationary process. For most practical applications, this might be a rather irrelevant scenario.

Note that the restrictions on the intercept or linear trend under cases 2 and 4 do not affect the estimation of the ARDL model because it is irrelevant whether we treat c_0 (c_1) or b_0 (b_1) as a free parameter to be estimated. Under case 1, (2) is estimated without intercept and trend. Under cases 2 and 3, an intercept is included in the regression. Under cases 4 and 5, an intercept and linear time trend are included. However, the restrictions are incorporated into step 1 of the bounds testing procedure, which we describe in the following:

1. First, we test the joint null hypothesis

$$H_0 : \begin{cases} (\pi_y = 0) \cap (\boldsymbol{\pi}_x = \mathbf{0}), & \text{case 1, 3, or 5} \\ (\pi_y = 0) \cap (\boldsymbol{\pi}_x = \mathbf{0}) \cap c_0 = 0, & \text{case 2} \\ (\pi_y = 0) \cap (\boldsymbol{\pi}_x = \mathbf{0}) \cap c_1 = 0, & \text{case 4} \end{cases}$$

versus the alternative hypothesis

$$H_1 : \begin{cases} (\pi_y \neq 0) \cup (\boldsymbol{\pi}_x \neq \mathbf{0}), & \text{case 1, 3, or 5} \\ (\pi_y \neq 0) \cup (\boldsymbol{\pi}_x \neq \mathbf{0}) \cup c_0 \neq 0, & \text{case 2} \\ (\pi_y \neq 0) \cup (\boldsymbol{\pi}_x \neq \mathbf{0}) \cup c_1 \neq 0, & \text{case 4} \end{cases}$$

The hypotheses are not directly formulated in terms of the long-run coefficients $\boldsymbol{\theta}$, because they are not well defined when $\pi_y = 0$. Instead, the test is formulated as a test for valid exclusion of the level terms y_{t-1} and \mathbf{x}_{t-1} (or \mathbf{x}_t) in (4) or (7). The test statistic is a conventional F statistic for joint validity of the $K + 1$ (or $K + 2$) restrictions imposed under the null hypothesis. However, the nonstandard distribution requires the use of different CVs, which we discuss further below. If the null hypothesis is not rejected, we conclude that there is no statistical evidence in favor of a long-run level relationship between y_t and \mathbf{x}_t . Otherwise, we should proceed with the following steps because of the possibility of degenerate cases, which are not ruled out by the alternative hypothesis of this first step.

2. If the null hypothesis from step 1 is rejected, we need to rule out the special case that y_t is $I(1)$ but not cointegrated with any variable in \mathbf{x}_t . This is done by testing

$$H_0 : \pi_y = 0 \quad \text{versus} \quad H_1 : \pi_y < 0$$

The test statistic is a conventional t statistic for statistical insignificance of the negative speed-of-adjustment estimate with a one-sided rejection region. As in step 1, the distribution is nonstandard, and the usual CVs do not apply. If the null hypothesis is not rejected, we conclude again that there is no statistical evidence of a long-run level relationship. Otherwise, we proceed with step 3.

3. If the null hypotheses in steps 1 and 2 are both rejected, we eventually consider the degenerate case that y_t is (trend) stationary but not part of a long-run relationship with \mathbf{x}_t . For this purpose, we can use conventional Wald tests for the joint (or individual) statistical insignificance of the long-run coefficients:

$$H_0 : \boldsymbol{\theta} = \mathbf{0} \quad \text{versus} \quad H_1 : \boldsymbol{\theta} \neq \mathbf{0}$$

We base this test on the long-run coefficients $\boldsymbol{\theta}$ rather than $\boldsymbol{\pi}_x$ because the OLS estimator of $\boldsymbol{\theta}$ is asymptotically normally distributed (Pesaran and Shin 1998), irrespective of the integration orders of \mathbf{x}_t , assuming that $\alpha > 0$ as indicated by the test result from step 2. Thus, conventional CVs can be used.

The rejection of the null hypotheses from all three steps is necessary to conclude that there is statistical evidence in favor of a long-run relationship; that is, $(\alpha > 0) \cap (\boldsymbol{\theta} \neq \mathbf{0})$. It is clear that the alternative hypothesis in step 1 does not rule out the two degenerate cases, which are the subject of steps 2 and 3. Yet we should still start with step 1 because it is carried out under less restrictive assumptions on the DGP than step 2.¹⁵

For the test statistics in steps 1 and 2, Pesaran, Shin, and Smith (2001) derive the asymptotic distributions under two scenarios. In the first scenario, all long-run forcing variables \mathbf{x}_t are individually $I(0)$. In the second scenario, all of them are $I(1)$ and not mutually cointegrated. When the (co)integration properties of \mathbf{x}_t are unknown, the corresponding CVs form lower and upper bounds. Conclusive evidence is possible when the value of the test statistic falls outside these bounds. The region for not rejecting the null hypothesis is below the lower bound (closer to zero), and the rejection region is above the upper bound. The test is inconclusive if the test statistic falls between the two bounds. Because the distributions have nonstandard forms, CVs have to be obtained by simulations. This is complicated by the fact that the distributions depend on the number of variables in \mathbf{x}_t . For $K \leq 10$, Pesaran, Shin, and Smith (2001) tabulated near-asymptotic CVs for the F statistic in step 1 and the t statistic in step 2. However, the asymptotic distributions might be poor approximations when the sample size is relatively small.¹⁶

15. For technical details and a full set of assumptions, see Pesaran, Shin, and Smith (2001).

16. Because the restrictions on the deterministic components do not alter the underlying DGP, the CVs for the single-hypothesis test in step 2 are the same for cases 2 and 3 and similarly for cases 4 and 5; see again Pesaran, Shin, and Smith (2001) for further discussion.

Note that the distributions and CVs are obtained under the assumption of independent and identically normally distributed errors u_t . As mentioned earlier, a standard procedure for dealing with suspected serial correlation is to increase the lag orders p , q , or both in the ARDL model. While the $p + Kq$ short-run terms in the EC representation do not affect the asymptotic distributions of the test statistics, they are relevant for the finite-sample distributions. Consequently, different CVs are needed for each combination of T^* , K , and $p + Kq$, separately for the lower and upper bounds. Instead of tabulating vast amounts of CVs, Kripfganz and Schneider (2020) estimated response-surface regressions, which can predict CVs for any desired sample size, number of long-run forcing variables, and lag order. This includes asymptotic CVs. Another important advantage of this approach is the ability to compute approximate p -values, which facilitate statistical inference.

2.4 Practical guidelines

The following stages characterize a stylized ARDL approach to testing for the existence of a conditional long-run level relationship:

1. Decide about the candidate variables \mathbf{x}_t that are assumed to be long-run forcing for y_t . These variables can be either $I(0)$ or $I(1)$. No pretesting is necessary unless we suspect that a variable might be $I(2)$. Stationary variables \mathbf{z}_t that are suspected to affect the short-run dynamics—but not the long-run equilibrium—can be added to the ARDL model as well. If there is doubt about the (trend) stationarity of \mathbf{z}_t , unit-root tests can be carried out.
2. Decide about the deterministic model components to be included in the model and whether the constant or linear trend coefficient should be restricted; that is, choose one of the five cases above. If in doubt, choose a more flexible model.¹⁷
3. Choose a maximum lag order p^* , ensuring that sufficiently many degrees of freedom are available.¹⁸ Keeping the estimation sample fixed, use the AIC or BIC to obtain the optimal lag orders p and q . To assert that the model is dynamically complete, a serial-correlation test could be of assistance. If there is concern about remaining serial correlation, the AIC might be preferred over the BIC because it tends to select less parsimonious models. Additional specification tests—for example, tests for heteroskedasticity and normality of the errors—could be used to check whether the assumptions underlying the bounds test are met.
4. Check the plausibility of the coefficient estimates in the EC representation. For example, an implausible estimate of α , which is clearly outside of the interval $[0, 2)$, might give rise to concern about the correct model specification or a potential overparameterization of the model.

17. The higher the case number, the less restrictive the model specification.

18. The Kripfganz and Schneider (2020) CVs are available only if there are at least twice as many observations T^* than coefficients K^* . For reliable inference, a much higher ratio is usually recommended.

5. Follow the three steps of the bounds test procedure. For steps 1 and 2, do not reject the null hypothesis if the value of the test statistic is below—that is, closer to zero—the lower bound of the Kripfganz and Schneider (2020) CVs. Reject the null hypothesis (and proceed with the next testing step) if the test statistic exceeds the upper-bound CV.
6. If there is conclusive statistical evidence in favor of a long-run relationship, consider refitting a more parsimonious model with lag orders selected by the BIC. If there is evidence against a long-run level relationship, consider refitting an ARDL model in first differences to obtain more efficient estimates,

$$\Delta y_t = c_0 + c_1 t + \sum_{i=1}^{p-1} \psi_{yi} \Delta y_{t-i} + \sum_{i=0}^{q-1} \psi'_{xi} \Delta \mathbf{x}_{t-i} + \gamma' \mathbf{z}_t + u_t$$

which is a restricted version of (7) with $\pi_y = 0$ and $\boldsymbol{\pi}_x = \mathbf{0}$. In both cases, it might be worth removing variables that do not help to improve the model fit. This reestimation stage can be skipped if there is no interest in further statistical analysis—for example, forecasting—beyond the exploration of a levels relationship.

To avoid pretesting problems, keep model simplifications—like those at stage 6—to a minimum before the bounds test is performed. Also note that there is no need to separately fit a static model in levels if the bounds test provides evidence in favor of a long-run relationship. As discussed earlier, the respective long-run coefficients can be inferred directly from the EC representation (3) or (6).

3 The ardl command

3.1 Syntax

```
ardl depvar [indepvars] [if] [in] [, lags(numlist) exog(exogvars) ec ec1
    noconstant trendvar[(trendvarname)] restricted regstore(storename)
    perfect maxlags(numlist) aic bic maxcombs(combnum)
    matcrit(lagcombmats) nofast dots noheader display_options]
```

3.2 Options

`lags(numlist)` specifies the number of lags for some or all regressors. The first number specifies the lag length p for `depvar` (y_t), which must be larger than zero. The following numbers specify the lag lengths q for the independent variables in the order they appear in `indepvars` (\mathbf{x}_t), which can be zero or higher. Missing values are allowed; they indicate that the respective lag order is not prespecified but instead determined with information criteria. If `numlist` contains only one element, the same

lag order is applied to all variables. Otherwise, the number of elements in *numlist* must equal the number of variables in *depvar* and *indepvars*.

exog(*exogvars*) specifies additional variables (\mathbf{z}_t) to be added to regression. An automatic lag order selection is not performed for these variables.

ec displays the results in error-correction form. *indepvars* enter the long-run relationship with time subscript t , as in (6).

ec1 displays the results in error-correction form. *indepvars* enter the long-run relationship with time subscript $t - 1$, as in (3).

noconstant suppresses the constant term. Specifying this option implies that the bounds test uses CVs for case 1.

trendvar[(*trendvarname*)] specifies a linear time trend to be added to the regression. *trendvarname* must be a variable that is collinear with *timevar*, the variable that is used with **tsset** to declare the data to be time-series data. Specifying **trendvar** is equivalent to **trendvar**(*timevar*). Specifying this option implies that the bounds test uses CVs for case 4 or 5.

restricted specifies that the constant term or the time trend, if specified, will be restricted for the purpose of the bounds test. The restricted deterministic component will be displayed in the long-run section of the error-correction output. Specifying this option implies that the bounds test uses CVs for case 2 or 4.

regstore(*storename*) stores the estimation results from the underlying **regress** command. These are the OLS estimates of (4) or (7) when option **ec1** or **ec0** is specified, respectively, and (2) otherwise.

perfect omits the collinearity check among the regressors.

maxlags(*numlist*) specifies the maximum lag order p^* for the optimal lag order selection. The first number specifies the maximum lag length for *depvar* (y_t), which must be larger than zero. The following numbers specify the maximum lag lengths for the independent variables in the order they appear in *indepvars* (\mathbf{x}_t), which can be zero or higher. Missing values are allowed; they indicate that the default maximum lag order 4 is to be used. If *numlist* contains only one element, the same maximum lag order is applied to all variables. Otherwise, the number of elements in *numlist* must equal the number of variables in *depvar* and *indepvars*.

aic requests that the optimal lag lengths be determined with the AIC.

bic, the default, requests that the optimal lag lengths be determined with the BIC.

maxcombs(*combnnum*) restricts the maximum number of lag permutations for the automatic lag selection. The default is **maxcombs**(100000), or **maxcombs**(500) if option **nofast** is specified. Higher values are possible.¹⁹

19. The purpose of this option is to prevent the optimal lag order selection from taking a lot of time without explicit user consent.

`matcrit(lagcombmata)` saves the lag permutations and the respective information criterion in a matrix named *lagcombmata*.

`nofast` uses the `regress` command instead of dedicated Mata code to run the auxiliary regressions for the optimal lag order selection. This is much slower but might be numerically more robust in rare cases.

`dots` displays a progress bar for the optimal lag order selection. This is useful when there are many permutations because of many variables and high maximum lag orders. Each dot represents a 1% progress in the evaluation of candidate models.

`noctable` suppresses the display of the coefficient table.

`noheader` suppresses the display of the coefficient table header.

display_options: `noomitted`, `vsquish`, `cformat(%fmt)`, `pformat(%fmt)`, `sformat(%fmt)`, and (Stata 12+ only) `no!stretch`; see [R] **Estimation options**.

3.3 Stored results

`ardl` stores the following results in `e()`:²⁰

Scalars

<code>e(N)</code>	number of observations	<code>e(ll_0)</code>	log likelihood, constant only
<code>e(df_m)</code>	model degrees of freedom	<code>e(rank)</code>	rank of <code>e(V)</code>
<code>e(df_r)</code>	residual degrees of freedom	<code>e(F)</code>	<i>F</i> statistic *
<code>e(mss)</code>	model sum of squares	<code>e(case)</code>	case number for deterministic components *
<code>e(rss)</code>	residual sum of squares	<code>e(F_pss)</code>	bounds-test <i>F</i> statistic *
<code>e(rmse)</code>	root mean squared error	<code>e(t_pss)</code>	bounds-test <i>t</i> statistic *
<code>e(r2)</code>	R^2	<code>e(numcombs)</code>	number of lag combinations *
<code>e(r2_a)</code>	adjusted R^2	<code>e(N_gaps)</code>	number of gaps in sample
<code>e(ll)</code>	log likelihood under assumption of independent and identically distributed normal errors	<code>e(tmin)</code>	first time period in sample
		<code>e(tmax)</code>	last time period in sample

Macros

<code>e(cmd)</code>	<code>ardl</code>	<code>e(tmaxs)</code>	formatted maximum time
<code>e(cmdline)</code>	command as typed	<code>e(regressors)</code>	full set of regressors
<code>e(cmdversion)</code>	version of the <code>ardl</code> command	<code>e(det)</code>	deterministic components *
<code>e(depvar)</code>	name of dependent variable	<code>e(exogvars)</code>	exogenous variables *
<code>e(title)</code>	title in estimation output	<code>e(srvars)</code>	short-run regressors *
<code>e(estat_cmd)</code>	<code>ardl_estat</code>	<code>e(lrdet)</code>	long-run deterministic component *
<code>e(predict)</code>	<code>ardl_p</code>	<code>e(lrxvars)</code>	long-run regressors *
<code>e(tsfmt)</code>	format for the time variable	<code>e(properties)</code>	<code>b V</code>
<code>e(tvar)</code>	time variable	<code>e(model)</code>	level or <code>ec</code>
<code>e(tmins)</code>	formatted minimum time		

Matrices

<code>e(b)</code>	coefficient vector	<code>e(maxlags)</code>	maximum lag lengths
<code>e(V)</code>	variance-covariance matrix	<code>e(lags)</code>	lag lengths in ARDL model

Functions

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

20. Starred results (*) are not always stored.

4 Postestimation commands

Many standard postestimation commands for the `regress` command can be used after the `ardl` command. Importantly, the results obtained with some of them can differ depending on whether the model is specified in the ARDL level form (2) or one of the EC forms (3) or (6). For example, the `estat ovtest` includes higher-order powers of the dependent variable—which is either y_t or Δy_t —as regressors in an auxiliary regression. This complication does not apply to postestimation commands based on residuals—such as `estat bgodfrey` and `estat imtest`—because the error term u_t is unaffected by the model’s reparameterization.

The Pesaran, Shin, and Smith (2001) bounds test for the existence of a long-run level relationship with Kripfganz and Schneider (2020) CVs and approximate p -values—as discussed in section 2.3—is implemented in the postestimation command `estat ectest`. It requires the option `ec` or `ec1` to be specified with the `ardl` command.

4.1 Syntax

```
estat ectest [ , siglevels(numlist) asymptotic nocritval norule
               nodecision ]
```

4.2 Options

`siglevels(numlist)` shows CVs for levels in the `numlist`, which must have at least one element. The default is `siglevels(10 5 1)`. Levels are specified as percentiles but do allow for two digits after the decimal point. There are 221 different levels among which you can choose, indicated by the Stata numlist 0.01 0.02 0.05 0.10(0.10)0.90 1.00(0.50)98.50 99.00(0.10)99.90 99.95 99.98 99.99.

`asymptotic` requests that the sample size returned by `ardl` in `e(N)` be ignored and show asymptotic CVs instead.

`nocritval` suppresses display of the CVs table.

`norule` suppresses display of the decision rule.

`nodecision` suppresses display of the decision table.

5 Example

We illustrate the `ardl` command with an example on cryptocurrencies.²¹ Specifically, we investigate whether supply and demand factors have a long-run impact on the price

21. As another illustrative example, we replicate the empirical analysis of Pesaran, Shin, and Smith (2001), who estimate an earnings equation with macroeconomic data for the United Kingdom. This exercise can be found in the working paper version of this article (Kripfganz and Schneider 2022).

of Bitcoin (variable `bprice`, in U.S. dollars [USD]). Few would debate that Bitcoin and many other cryptocurrencies are highly speculative financial assets in the short run. Taking a long-run perspective, an often-invoked explanation for the steep rise to almost USD 69,000 per Bitcoin in November 2021 is that ever-increasing demand meets limited supply. This cannot explain, however, that the spectacular rise of the Bitcoin price was followed recently by a stark drop to almost USD 15,000 in late 2022. This raises the question whether supply and demand forces can indeed tame the Bitcoin price in the long run or whether its speculative (and therefore unpredictable) nature is prevailing. In econometric terms, we aim to investigate whether we can find evidence for a long-run equilibrium relationship between the Bitcoin price and its supply and demand factors.

The motivation for the key variables in our dataset follows Ciaian, Rajcaniova, and Kancs (2016). If a long-run equilibrium exists, the Bitcoin price is expected to be inversely proportional to its supply, which can be approximated by the historical number of mined Bitcoins (variable `supply`). On the demand side, the equilibrium price can be expected to grow proportionally to the size of the Bitcoin economy, as measured by the number of daily Bitcoin transactions (`ntrans`). It would also be inversely related to its velocity. Here the (inverse) velocity is proxied by so-called coin days destroyed (`ddestr`). Broadly speaking, this is an aggregate measure of how much time has elapsed between two transactions with the same coin. More transactions with coins that have been dormant for a longer time indicate an increase in economic activity.

Because Bitcoin is predominantly priced in USD, a depreciation of the dollar makes it cheaper to carry out Bitcoin transactions for investors in the rest of the world, therefore increasing demand. For simplicity, following Ciaian, Rajcaniova, and Kancs (2016), we just include the USD/EUR exchange rate (`fxeu_f`) in our set of explanatory variables.²² The linearization of the equilibrium relationship requires a log transformation for all the variables, indicated by the prefix `ln_` in the variable names used below.²³ In our sample, we have 3,255 daily observations from January 1, 2014, to November 29, 2022.²⁴ We do not use pre-2014 data, because Ciaian, Rajcaniova, and Kancs (2016) find evidence of a structural break in 2013.²⁵

We start with a visual inspection of the key variables.²⁶ In the left panel of figure 1, the evolution of the log Bitcoin price and its supply are shown. The latter largely follows a deterministic path, which is prescribed by the underlying Bitcoin protocol. The Bitcoin price shows all signs of a nonstationary variable. While it shares a similar

22. The `_f` in the `fxeu_f` variable name indicates that we imputed the missing weekend values by carrying forward the Friday value.

23. This is an important deviation from Ciaian, Rajcaniova, and Kancs (2016), who refrain from this transformation because of zero values for some of the variables in the early days of Bitcoin. Here we focus on a later sample period, where Bitcoin was well established.

24. The data sources are coinmetrics.io for `bprice`, `supply`, and `ntrans`, <https://blockchair.com> for `ddestr`, and <https://fred.stlouisfed.org> (Federal Reserve Bank of St. Louis) for `fxeu` (FRED code DEXUSEU). The data were downloaded in December 2022 and January 2023.

25. Our estimation results are sensitive to the inclusion of earlier data points. It can be argued that the pre-2014 period may not be representative for the subsequent dynamics, because Bitcoin was still in a maturation stage.

26. The code for replicating figure 1 can be found in the ancillary file `ard1_example.do`.

upward trend with the supply, the quasideterministic nature of the mining process precludes that the two series could be cointegrated. To avoid distorting the bounds test, it is thus advisable to exclude the supply from the long-run relationship for testing purposes. It can still enter the regression model as an exogenous price determinant—a \mathbf{z}_t variable in terms of the notation from section 2.

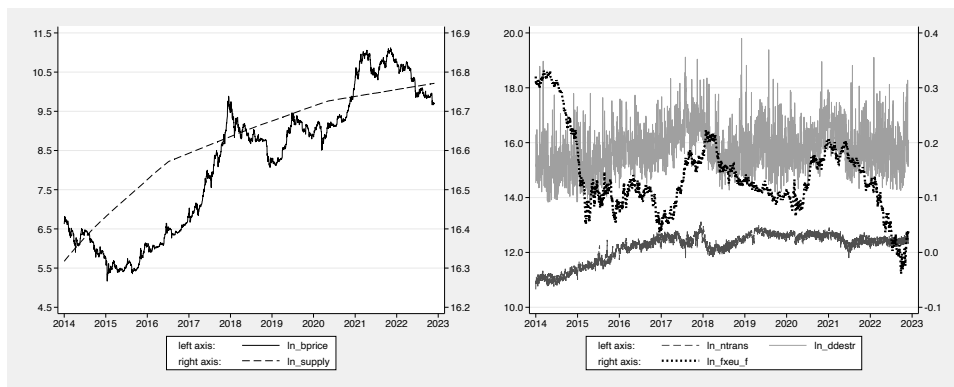


Figure 1. Time-series graphs for the main regression variables

The right panel of figure 1 depicts the demand side factors. The USD/EUR exchange rate is clearly nonstationary, while coin days destroyed look fairly stationary. The picture is less clear about the daily Bitcoin transactions series, which appears to follow different time trends at different periods in our sample and therefore is likely nonstationary. We could verify these assessments with conventional unit-root tests, but this is not necessary for ARDL estimation and bounds testing. It is one of the latter's advantages that it can deal with mixtures of $I(0)$ and $I(1)$ variables. We will confirm our initial assessment further below in the context of fitting a VEC model, where pretesting for the order of integration is required.

There is no apparent reason to believe that the observed time trend in the Bitcoin price is entirely attributable to the underlying time trend in the other variables. This calls for the inclusion of a restricted time trend—case 4—in the EC model. Under the null hypothesis of the bounds test, the log Bitcoin price would then follow a random walk with drift. Alternatively, if $\alpha \neq 0$, it can be either cointegrated or trend stationary.

Given that we have daily data, we choose a maximum lag order $p^* = 7$, such that the lags can cover up to one week. Thanks to the large number of observations, we do not have to be conservative with the degrees of freedom and therefore select the optimal lag combination with the AIC rather than the BIC. This reduces the risk of misspecifying the model dynamics, which in turn might invalidate the bounds test. We therefore add the `maxlags(7)` and `aic` options to our `ardl` command line. The quasideterministic supply variable is specified with the `exog()` option, which also constrains its lag order to zero. If we were to add it as a conventional independent variable instead, the command

would error out because of collinearity of the lags.²⁷ To complete our specification, we add a linear time trend by including the name of the time identifier in the `trend()` option. While the optimization over all 3,584 lag combinations finishes in virtually no time using our fast Mata algorithm, we illustrate how to display a progress bar with the option `dots`, which might be useful for larger models:

```
. use ardl_example
(Bitcoin price related data, -ardl-, SJ (type -notes- for sources))
. ardl ln_bprice ln_ntrans ln_ddestr ln_fxex_f, exog(ln_supply) trendvar(date)
> aic maxlags(7) dots
Optimal lag selection, % complete:
-----|-----|-----|-----|-----|
      20%      40%      60%      80%     100%
.....
AIC optimized over 3584 lag combinations
ARDL(2,1,3,2) regression
Sample: 2014-01-08 thru 2022-11-29
                                         Number of obs =      3,248
                                         F(13, 3234)   = 519948.59
                                         Prob > F      =      0.0000
                                         R-squared     =      0.9995
                                         Adj R-squared =      0.9995
                                         Root MSE    =      0.0388

Log likelihood = 5950.3079
```

ln_bprice	Coefficient	Std. err.	t	P> t	[95% conf. interval]	
ln_bprice						
L1.	.954804	.0175476	54.41	0.000	.9203984	.9892096
L2.	.038809	.0175277	2.21	0.027	.0044424	.0731756
ln_ntrans						
--.	.0234125	.0062507	3.75	0.000	.0111568	.0356682
L1.	-.0097699	.0062652	-1.56	0.119	-.022054	.0025141
ln_ddestr						
--.	-.0014938	.0010689	-1.40	0.162	-.0035896	.000602
L1.	-.0006291	.0010955	-0.57	0.566	-.0027769	.0015188
L2.	.0006551	.0010921	0.60	0.549	-.0014863	.0027964
L3.	.0027045	.0010154	2.66	0.008	.0007137	.0046953
ln_fxex_f						
--.	.4878456	.1658287	2.94	0.003	.1627056	.8129856
L1.	-.7029864	.2349611	-2.99	0.003	-1.163674	-.2422987
L2.	.2828991	.1662484	1.70	0.089	-.0430638	.608862
ln_supply	-.0283604	.0420257	-0.67	0.500	-.1107601	.0540393
date	.0000114	4.86e-06	2.34	0.020	1.82e-06	.0000209
_cons	.0863434	.5877493	0.15	0.883	-1.066055	1.238742

27. Alternatively, we could specify the supply variable as an independent variable and manually restrict its lag order to zero with the `lags()` option. While this yields identical ARDL coefficient estimates, it would later on include the supply variable in the long-run relationship when we reparameterize the model in EC form. As argued above, this is problematic with regard to the bounds test.

The optimal model chosen by the AIC is an ARDL(2,1,3,2) model.²⁸ The supply of Bitcoin has no statistically significant effect, which could justify removing this regressor at a later stage. In contrast, the linear time trend, represented by the variable `date`, is statistically significant at the 5% level, in line with our earlier observations.²⁹ Before turning our attention to the bounds test, we should inspect the residuals for potential serial correlation:

```
. estat bgodfrey, lags(1/3)
```

Breusch-Godfrey LM test for autocorrelation

lags(p)	chi2	df	Prob > chi2
1	0.161	1	0.6880
2	0.161	2	0.9224
3	0.768	3	0.8571

H0: no serial correlation

The Lagrange multiplier test does not provide reason for concern about residual serial correlation. We now refit the model in error-correction form—(6)—using the `ec` option. While the AIC would give us the same lag orders again, we can also directly specify the optimal lag orders with the `lags()` option. However, we then need to exert some caution to obtain results for the same estimation sample as above. Allowing for a maximum of seven lags, we set aside the first seven data points for the optimal lag determination. The estimation sample was held fixed by the `ardl` command even for models with lower lag orders. To base the bounds test again on the same estimation sample, we restrict it in the next step with the `e(sample)` function. To obtain the correct CVs with the bounds test for case 4, we now also need to add the option `restricted`, which includes the time trend in the long-run relationship:

28. The BIC would select an overly parsimonious ARDL(1,0,0,0) model. However, there would be evidence of remaining serial correlation in the residuals, potentially invalidating the bounds test.

29. If we were unsure about including a time trend, we could also refit the model without it and compare the specifications again with the AIC. Here the AIC is lower—and therefore preferable—for the model with trend. Note that the Bitcoin supply remains statistically insignificant (at the 5% level) even in the model without linear time trend.

```
. ardl ln_bprice ln_ntrans ln_ddestr ln_fxex_f if e(sample), exog(ln_supply)
> trendvar(date) lags(2 1 3 2) ec restricted
ARDL(2,1,3,2) regression
Sample: 2014-01-08 thru 2022-11-29
Log likelihood = 5950.3079
Number of obs = 3,248
R-squared      = 0.0175
Adj R-squared  = 0.0135
Root MSE      = 0.0388
```

D.ln_bprice	Coefficient	Std. err.	t	P> t	[95% conf. interval]	
ADJ						
ln_bprice L1.	-.006387	.0018308	-3.49	0.000	-.0099767	-.0027973
LR						
ln_ntrans	2.135988	.9257322	2.31	0.021	.3209067	3.951069
ln_ddestr	.1936262	.2219016	0.87	0.383	-.2414558	.6287081
ln_fxex_f	10.60878	2.015474	5.26	0.000	6.657047	14.56052
date	.001778	.0006272	2.83	0.005	.0005483	.0030077
SR						
ln_bprice LD.	-.038809	.0175277	-2.21	0.027	-.0731756	-.0044424
ln_ntrans D1.	.0097699	.0062652	1.56	0.119	-.0025141	.022054
ln_ddestr D1.	-.0027305	.0014205	-1.92	0.055	-.0055157	.0000547
LD.	-.0033596	.0012433	-2.70	0.007	-.0057972	-.0009219
L2D.	-.0027045	.0010154	-2.66	0.008	-.0046953	-.0007137
ln_fxex_f D1.	.4200873	.1659749	2.53	0.011	.0946607	.745514
LD.	-.2828991	.1662484	-1.70	0.089	-.608862	.0430638
ln_supply _cons	-.0283604	.0420257	-0.67	0.500	-.1107601	.0540393
	.0863434	.5877493	0.15	0.883	-1.066055	1.238742

The first coefficient in the ADJ section of the regression output is the negative speed-of-adjustment coefficient $\pi_y = -\alpha$. Its magnitude is very small. At this stage, we should not be fooled by the reported p -value and confidence interval into believing that it is statistically significant. The t statistic for this coefficient does not have a standard distribution under the null hypothesis. In fact, this is the test statistic that we consider under the second step of the bounds test. The long-run coefficients θ in the LR section all have the expected sign. Coin days destroyed is statistically insignificant, which is not too surprising given the suspected differences in the integration orders. The trend coefficient (**date**) is reported in the LR section because of the option **restricted**. Without that option, it would be reported in the SR section together with the other short-run coefficients ω , ψ_{xi} , and γ . This does not affect any of the other coefficients, but it will matter for the bounds test. Because of first differencing, the lag orders for the short-run terms are one less than those in the ARDL level representation. Also, note that the exogenous supply variable is not transformed into first differences and just enters the

model with the same coefficient as in the ARDL specification. Before we start interpreting the coefficient estimates, we first need to establish with the bounds test whether a long-run relationship exists. We do this with the `estat ectest` postestimation command:

```
. estat ectest
Pesaran, Shin, and Smith (2001) bounds test
H0: no level relationship                      F =      6.156
Case 4                                       t =     -3.489

Finite sample (3 variables, 3248 observations, 8 short-run coefficients)
Kripfganz and Schneider (2020) critical values and approximate p-values
```

	10%		5%		1%		p-value	
	I(0)	I(1)	I(0)	I(1)	I(0)	I(1)	I(0)	I(1)
F	2.957	3.721	3.391	4.209	4.315	5.235	0.000	0.002
t	-3.128	-3.819	-3.412	-4.128	-3.964	-4.709	0.041	0.186

```
do not reject H0 if
    either F or t are closer to zero than critical values for I(0) variables
    (if either p-value > desired level for I(0) variables)
reject H0 if
    both F and t are more extreme than critical values for I(1) variables
    (if both p-values < desired level for I(1) variables)
decision: no rejection (.a), inconclusive (.), or rejection (.r) at levels:
```

	10%	5%	1%
decision	.	.	.a

In the top-right corner, the bounds-test output displays the test statistics for the first two testing steps, as outlined in section 2.3. The command reports the Kripfganz and Schneider (2020) CVs for finite samples. However, because of the large sample size, they are virtually identical with the asymptotic CVs.³⁰ First, we consider the F statistic for the joint null hypothesis $\pi_y = 0$, $\pi_x = \mathbf{0}$, and $c_1 = 0$. The last coefficient captures the restriction on the time trend. The test statistic is larger than the upper-bound CVs—which would be the exact CVs if all long-run forcing variables were $I(1)$ —for the conventional significance levels. This is most easily seen by looking at the approximate p -values, which are computed with the response-surface methodology of Kripfganz and Schneider (2020). We therefore reject the null hypothesis in this first step. However, this is not yet sufficient evidence in favor of a long-run relationship, because we need to rule out the degenerate cases.

30. To use asymptotic CVs, we must specify the option `asymptotic` with `estat ectest`.

Second, we need to consider the individual null hypothesis $\pi_y = 0$. The test statistic is the same as in the ADJ section of the regression output, but only `estat ectest` provides the appropriate CVs. Here the conclusion depends on the chosen significance level. If we take a conservative stance with the 1% level, the t statistic is closer to zero than the lower-bound CV. We would therefore not reject the null hypothesis. The statistical significance of the long-run coefficients in the EC regression output then becomes irrelevant, because the equilibrium correction term $y_{t-1} - \theta \mathbf{x}_t$ drops out from (6) with $\pi_y = \alpha = 0$ under the null hypothesis. In the long run, the log Bitcoin price can thus be characterized by a unit-root process that is not driven by the independent variables in our model.

As we move to a more relaxed stance on the risk of committing a type-I error—rejecting the null hypothesis when it is actually true—the bounds test becomes inconclusive at the 5% significance level. Here the value of the t statistic falls inside the two bounds, although it exceeds the lower bound only narrowly. Given the presence of $I(1)$ independent variables, the evidence still points more strongly toward not rejecting the null hypothesis. When we move further to the 10% level, the test statistic remains within the two bounds. Because the long-run coefficient of the only $I(0)$ variable, coin days destroyed, is statistically insignificant, the upper-bound CV carries much more weight than the lower bound. As a way of resolving this inconclusiveness, we can redo the bounds test for a model without coin days destroyed:

```
. ardl ln_bprice ln_ntrans ln_fxex_f, exog(ln_supply) trendvar(date) aic
> maxlags(7) ec restricted
(output omitted)
. estat ectest, norule
Pesaran, Shin, and Smith (2001) bounds test
H0: no level relationship                                F =      5.794
Case 4                                                    t =     -3.219
Finite sample (2 variables, 3248 observations, 11 short-run coefficients)
Kripfganz and Schneider (2020) critical values and approximate p-values
```

	10%		5%		1%		p-value	
	I(0)	I(1)	I(0)	I(1)	I(0)	I(1)	I(0)	I(1)
F	3.362	4.016	3.882	4.588	5.002	5.805	0.003	0.010
t	-3.128	-3.647	-3.412	-3.952	-3.963	-4.531	0.081	0.222

decision: no rejection (.a), inconclusive (.), or rejection (.r) at levels:

	10%	5%	1%
decision	.	.a	.a

Assuming that all integration orders are known to be $I(1)$, the bounds test still fails to reject the null hypothesis because the t statistic is less negative than the upper-bound CV at all significance levels. At the 5% level, it now even falls short of the lower bound. Consequently, the statistical significance of the long-run coefficients is not informative. Evidence of a cointegrating relationship could not be established.³¹

We can cross-check results with the Johansen (1995) framework using the `vecrank` and `vec` commands. Here, unlike the ARDL framework, we should specify only nonstationary variables. To be on the safe side, we might want to initially run the augmented Dickey and Fuller (1979) unit-root pretest for each of the variables. This can be done with the `dfuller` command. Because this unit-root test is a special case of the bounds test when there is only one variable, we can also use our `ardl` command for this purpose. As an example, we show this for the Bitcoin price:³²

```
. ardl ln_bprice, trendvar(date) aic maxlags(7) ec restricted
(output omitted)
. estat ectest, norule nodecision
Pesaran, Shin, and Smith (2001) bounds test
H0: no level relationship                      F =      1.382
Case 4                                       t =     -1.663
Finite sample (0 variables, 3248 observations, 1 short-run coefficients)
Kripfganz and Schneider (2020) critical values and approximate p-values
```

	10%		5%		1%		p-value	
	I(0)	I(1)	I(0)	I(1)	I(0)	I(1)	I(0)	I(1)
F	5.364	5.390	6.305	6.349	8.376	8.481	0.907	0.906
t	-3.130	-3.132	-3.414	-3.423	-3.964	-4.000	0.769	0.771

```
. dfuller ln_bprice if e(sample), lags(`=e(lags)[1,1]-1`) trend
Augmented Dickey-Fuller test for unit root
Variable: ln_bprice                      Number of obs = 3,248
                                         Number of lags = 1
H0: Random walk with or without drift
```

Test statistic	Dickey-Fuller critical value		
	1%	5%	10%
Z(t)	-1.663	-3.960	-3.120

MacKinnon approximate p -value for Z(t) = 0.7669.

31. If we had modeled an unrestricted trend by removing the option `restricted`, the qualitative conclusions would have remained unchanged because the decision primarily rests on the t statistic, whose CVs are invariant to restricting the time trend.

32. Note that the `dfuller` command requests the lag order to be specified for the first-differenced model, which is one less than the optimal lag order obtained with the `ardl` command for the level model.

The t statistic from the bounds test equals the Dickey and Fuller (1979) test statistic. The CVs also virtually coincide.³³ The F test reported by `estat ectest` corresponds to the Dickey and Fuller (1981) F statistic, which is not implemented elsewhere in Stata. Both tests confirm our prior assessment that the log Bitcoin price is nonstationary. While not shown here, similar tests for the other variables also support our initial classification. We can now proceed with the cointegration rank tests. For the best comparability with the previous results, we choose a lag order of 3—which was the maximum lag order for any variable selected by the AIC in the ARDL model—and restrict the estimation sample again to coincide with the one above:

```
. vecrank ln_bprice ln_ntrans ln_fxau_f if date >= td(08jan2014), lags(3)
> trend(rtrend) max levela
```

Johansen tests for cointegration
Trend: Restricted
Sample: 2014-01-08 thru 2022-11-29

					Number of obs = 3,248		
					Number of lags = 3		

Maximum				Trace	Critical value	
rank	Params	LL	Eigenvalue	statistic	5%	1%
0	21	21528.446		109.1701	42.44	48.45
1	27	21574.8	0.02814	16.4624*1*5	25.32	30.45
2	31	21581.604	0.00418	2.8545	12.25	16.26
3	33	21583.031	0.00088			

Maximum				Eigenvalue	Critical value	
rank	Params	LL		Maximum	5%	1%
0	21	21528.446		92.7077	25.54	30.34
1	27	21574.8	0.02814	13.6079	18.96	23.65
2	31	21581.604	0.00418	2.8545	12.52	16.26
3	33	21583.031	0.00088			

* selected rank

The Johansen (1995) trace and maximum-eigenvalue tests both indicate a cointegration rank of one. However, this is only a necessary but not sufficient condition for the presence of an error-correction mechanism in the process of the log Bitcoin price.³⁴ In the next step, we fit the VEC model:

33. Theoretically, the lower-bound and upper-bound CVs reported by `estat ectest` should coincide for this special case. However, because they were obtained from separate response-surface regressions, minor numerical discrepancies occur; see Kripfganz and Schneider (2020) for details. For smaller sample sizes, these CVs can be more accurate than the original Dickey and Fuller (1979) ones, especially when the regression is augmented with multiple lags.

34. It is possible that the cointegrating relationship comprises only variables other than the Bitcoin price. Conversely, while there should be at most one cointegrating relationship involving the Bitcoin price for the ARDL approach to be applicable, which is satisfied here, it is worth reiterating that a cointegration rank larger than one for the entire system by itself would not automatically invalidate the ARDL procedure. This is again due to the possibility of cointegrating relationships among the other variables themselves; see also the discussion in section 2.2.

```
. vec ln_bprice ln_ntrans ln_fxex_f if date >= td(08jan2014), rank(1) lags(3)
> trend(rtrend) alpha noetable
```

Vector error-correction model

Sample: 2014-01-08 thru 2022-11-29

Number of obs	=	3,248
AIC	=	-13.26835
HQIC	=	-13.25023
SBIC	=	-13.21776

Cointegrating equations

Equation	Parms	chi2	P>chi2
_ce1	2	138.2818	0.0000

Identification: beta is exactly identified

Johansen normalization restriction imposed

beta	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
_ce1						
ln_bprice	1
ln_ntrans	-3.169487	.3007237	-10.54	0.000	-3.758895	-2.580079
ln_fxex_f	-15.18478	1.632316	-9.30	0.000	-18.38406	-11.9855
_trend	-.0010287	.0001443	-7.13	0.000	-.0013116	-.0007458
_cons	34.41289

Adjustment parameters

Equation	Parms	chi2	P>chi2
D_ln_bprice	1	21.04656	0.0000
D_ln_ntrans	1	68.06329	0.0000
D_ln_fxex_f	1	.0319188	0.8582

alpha	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
D_ln_bprice						
_ce1						
L1.	-.0034951	.0007618	-4.59	0.000	-.0049882	-.0020019
D_ln_ntrans						
_ce1						
L1.	.0186699	.002263	8.25	0.000	.0142345	.0231053
D_ln_fxex_f						
_ce1						
L1.	.0000144	.0000806	0.18	0.858	-.0001435	.0001723

Clearly, according to the bottom table for the speed-of-adjustment coefficients, the USD/EUR exchange rate is not loading onto the cointegrating relationship. The respective coefficient for the Bitcoin price is also very small. For practical matters, the Bitcoin price hardly reacts to deviations from the equilibrium relationship. Thus, the statistical question of whether there exists a long-run relationship should not bear too much weight in the final assessment, because it would take a very long time for the Bitcoin price to return to such an equilibrium. Somewhat problematic is the wrong sign of the statistically significant adjustment coefficient for the number of transactions.

This points toward an instability in the system. By focusing only on the equation for the Bitcoin price, the ARDL approach avoids this issue. Another disadvantage of the `vecrank` command is that it does not allow inclusion of exogenous $I(0)$ variables, unlike `ardl`. Furthermore, by estimating a system of equations, the number of coefficients to be estimated can be substantially larger in the VEC model. Notably, the long-run coefficients are broadly consistent with the ARDL results, even though their relevance should not be overstated because of the minuscule or nonexistent error adjustment.

So far, there is no convincing evidence in favor of a long-run relationship of the Bitcoin price with traditional supply and demand side characteristics. However, the demand for Bitcoin may generally depend on other or additional factors than those for well-established currencies and investment assets. For example, it may depend on how well the cryptocurrency market is understood and trusted by potential investors. Furthermore, macrofinancial developments can affect the willingness to invest in high-risk assets. Ciaian, Rajcaniova, and Kancs (2016) therefore include the number of views of Bitcoin's Wikipedia page (`wikiw`) as a measure of investment attractiveness, the Dow Jones Industrial Average stock market index (`djon_f`) as a proxy for investor sentiment, and the Brent crude oil price (`oprc_f`) as an indicator of macroeconomic risks.³⁵ Given its statistical insignificance, we remove the Bitcoin supply in the following specifications.

```
. ardl ln_bprice ln_ntrans ln_ddestr ln_fxex_f ln_wikiw ln_djon_f ln_oprc_f,
> trend(date) aic maxlags(7) ec restricted maxcombs(2000000) dots
(output omitted)
. estat ectest, norule
(output omitted)
```

To economize on space, we summarize the results for the speed-of-adjustment coefficient (ADJ) and the long-run coefficients (LR, excluding the linear time trend) in table 1 instead of showing detailed Stata output. For ease of comparability, column 1 repeats the results from our initial regression further above. Let us first look at column 2. The stock market index and the oil price index do not appear to be relevant long-run forcing variables for the Bitcoin price, irrespective of whether any long-run relationship exists in the first place. In contrast, the long-run coefficient of Wikipedia views is statistically significant. With its inclusion, the bounds test now also conclusively rejects the null hypothesis, although only at the 10% level. However, the economic significance of this result remains limited because of the very slow speed of adjustment. Compared with the benchmark specification in column 1, the main statistical reason for the reduction in the bounds-test p -values is the near doubling of the speed-of-adjustment coefficient, which is immediately reflected in a larger t statistic. However, this cannot mask the fact that the economic effect size is still negligible.

35. Data sources are <https://wikishark.com> (Vardi et al. 2021) for `wikiw` and <https://fred.stlouisfed.org> (Federal Reserve Bank of St. Louis) for `djon` (S&P Dow Jones Indices LLC, FRED code DJIA) and `oprc` (FRED code DCOILBRENTU).

Table 1. ARDL long-run estimation results in EC representation

		Column				
	D.ln_bprice	1	2	3	4	5
date					>10jul2016	>10jul2016 <12may2020
ADJ	L.ln_bprice	-0.006	-0.011	-0.011	-0.010	-0.011
LR	ln_ntrans	2.136**	1.285***	1.256***	1.556**	1.504*
	ln_ddestr	0.194	0.027	0.025	0.005	0.190
	ln_fxex_f	10.609***	7.862***	7.838***	8.538***	9.054**
	ln_wikiw		0.315***	0.320***	0.422**	0.389*
	ln_djon_f		1.396	1.514	0.453	0.431
	ln_oprc_f		0.065	.	.	
<i>F</i>		6.156***	5.198***	5.934***	3.985**	2.872**
<i>t</i>		-3.489**	-4.521***	-4.522***	-3.355*	-2.706

NOTE: Stars indicate the significance level (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.) Conventional p -values for the ADJ coefficient are invalid and thus not reported. For the bounds test, both upper-bound and lower-bound significance levels are indicated. A conclusive decision requires either significance or insignificance with respect to both the upper and lower bounds. Significance only with respect to the lower bound indicates inconclusiveness.

In column 3, we exclude the irrelevant oil prices from the model. Despite their insignificant long-run coefficients, we keep the stock market index and coin days destroyed because they still have significant short-run effects. The estimates hardly differ from the previous specification.

```
. ardl ln_bprice ln_ntrans ln_ddestr ln_fxex_f ln_wikiw ln_djon_f, trend(date) aic
> maxlags(7) ec restricted maxcombs(300000) dots
(output omitted)
. estat ectest, norule
(output omitted)
```

Over time, Bitcoin (and cryptocurrencies in general) became more and more accessible to a wider audience and also attracted the interest of professional investors. This may have lead to a gradual change in the fundamental relationship between the Bitcoin price and its determinants. In econometric terms, we may have to worry about parameter instability. Stata offers several diagnostics for structural breaks, which we can use here because the `ardl` command supports all standard postestimation commands for `regress`. A routinely applied tool is the cumulative sum (CUSUM) test:

```
. estat sbcusum, ylabel(, angle(horizontal)) tttitle("") name(sb1)
Cumulative sum test for parameter stability
Sample: 08jan2014 thru 29nov2022      Number of obs = 3,248
H0: No structural break
```

Type	Test statistic	1%	5%	10%
Recursive	1.0067	1.1430	0.9479	0.8499

```
. estat sbcusum, ols ylabel(, angle(horizontal)) tttitle("") name(sb2)
Cumulative sum test for parameter stability
Sample: 08jan2014 thru 29nov2022      Number of obs = 3,248
H0: No structural break
```

Type	Test statistic	Critical value		
		1%	5%	10%
OLS	0.4834	1.6276	1.3581	1.2238

```
. graph combine sb1 sb2, ysize(2) xsize(5)
```

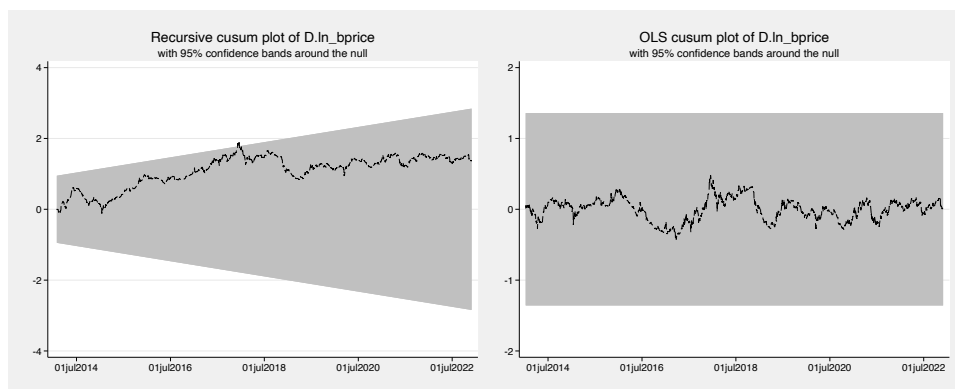


Figure 2. CUSUM plots

The CUSUM test based on OLS residuals does not trigger a warning sign. In contrast, the test based on recursive residuals rejects the null hypothesis of parameter stability at the 5% significance level. However, because the recursive CUSUM process travels beyond the 95% confidence bounds only very briefly, we may not have to worry too much. Figure 2 shows that the drift away from zero occurs rather gradually over time. This does not suggest a specific date for a structural break, other than that it may have occurred relatively early during our sample period. However, potential break points can be spotted in figure 1. While the Bitcoin supply did not turn out to be a relevant predictor in the earlier regressions, the discrete slowdowns in the mining of new Bitcoin at July 10, 2016, and May 12, 2020—so-called “halving dates”³⁶—could possibly have wider repercussions.

36. The Bitcoin mining algorithm ensures that the supply of new Bitcoins will ultimately come to an end. After a certain number of new Bitcoin blocks have been created, the rewards for the mining of a new Bitcoin are halved, driving inefficient miners out of the market and therefore reducing the supply growth.

```
. estat sbknown, break(td(10jul2016) td(12may2020))
Wald test for a structural break: Known break dates
Sample: 08jan2014 thru 29nov2022
Break dates: 10jul2016 12may2020
H0: No structural break
Number of obs = 3,248
      chi2(42) = 52.2187
      Prob > chi2 = 0.0072
(output omitted)

. estat sbknown, break(td(10jul2016) td(12may2020))
> breakvars(L.ln_bprice ln_ntrans ln_ddestr ln_fxex_f ln_wikivw ln_djon_f)
Wald test for a structural break: Known break dates
Sample: 08jan2014 thru 29nov2022
Break dates: 10jul2016 12may2020
H0: No structural break
Number of obs = 3,248
      chi2(12) = 4.9395
      Prob > chi2 = 0.9642
(output omitted)
```

Indeed, a parameter stability test with these known structural-break dates rejects the null hypothesis. However, if we restrict the test to the speed-of-adjustment and long-run coefficients, no instability is found. The latter is reassuring regarding our earlier results. Accounting for structural breaks in the short-run coefficients would become a potential issue if we were interested in a more detailed analysis of the short-run dynamics. Nevertheless, as a robustness check, we refit the model by considering only the observations after the first halving date. In another specification, we further curtail the sample with the second halving date.

```
. ardl ln_bprice ln_ntrans ln_ddestr ln_fxex_f ln_wikivw ln_djon_f
> if date >= td(10jul2016), trend(date) aic maxlags(7) ec restricted
> maxcombs(300000) dots
(output omitted)

. estat ectest, norule
(output omitted)

. ardl ln_bprice ln_ntrans ln_ddestr ln_fxex_f ln_wikivw ln_djon_f
> if date > td(10jul2016) & date < td(12may2020), trend(date) aic maxlags(7)
> ec restricted maxcombs(300000) dots
(output omitted)

. estat ectest, norule
(output omitted)
```

The main results are shown in columns 4 and 5 of table 1. The effect sizes hardly changed, especially for the speed of adjustment, which is instrumental for existence of an equilibrium correction mechanism. Interestingly though, the bounds test now does conclusively not reject the null hypothesis of no level relationship at least at the 5% significance level. Compared with the specifications in columns 2 and 3, this is mainly driven by the larger standard error of the speed-of-adjustment coefficient, partly because of the smaller sample size. Turning the argument around, the large size of the unrestricted sample—daily observations for almost nine years—previously enabled us to statistically detect (at the 10% significance level) an economically insignificant effect.

As a word of caution, the reliability of the bounds test could be hampered by the nonnormality of the regression errors. Heteroskedasticity and normality tests with the postestimation commands `estat hettest` and `estat imtest` tend to point in that direction (not shown here).³⁷ This is not unexpected when working with financial data. Ideas for further exploration include the incorporation of potential asymmetric effects and other nonlinearities, a quest for alternative explanatory variables, or a generalized autoregressive conditional heteroskedasticity modeling approach. We leave these avenues to the interested reader.

Overall, based on the results presented here, there do not seem to be strong forces in place that keep the log Bitcoin price in an equilibrium relationship with the candidate long-run forcing variables. Even if we accept column 2 or 3 as our preferred specification and take a liberal stand on the type-I error probability, the economic relevance of the rejected bounds test remains negligible because of the slow speed of adjustment. It appears that the price of Bitcoin is hardly driven by the underlying fundamentals but might be following the path of a predominantly speculative asset. If we accept the statistical conclusion from one of the other specifications that there is no significant long-run relationship present, we could proceed by refitting a more parsimonious version of the model purely in first differences, potentially also using the BIC instead of the AIC as a lag order selection criterion. This could then be used for forecasting purposes or further analyses of the dynamic adjustment processes. For the purpose of this article, however, our curiosity shall end here.³⁸

6 Conclusion

In this article, we have described the `ardl` command for the estimation of ARDL models with time-series data. The lag orders can be prespecified or chosen optimally with the AIC or BIC. For this purpose, the command is able to fit tens of thousands of candidate models in virtually no time. Two useful reparameterizations of the model in error-correction form allow for an interpretation of the coefficients as short-run and long-run effects. The command further enables testing for the existence of a long-run level relationship using the popular bounds test, which is implemented as a postestimation

37. For our application, increasing the ARDL lag orders does not provide an improvement.

38. In the working paper version of this article (Kripfganz and Schneider 2022), we present a forecasting example with a different dataset.

feature. For nonstationary variables, this amounts to cointegration testing. Yet the ARDL approach is flexible to allow for both stationary and nonstationary variables. The package provides the recently improved Kripfganz and Schneider (2020) CVs for the bounds test, which allow accurate inference for almost all practically relevant combinations of sample size, number of long-run forcing variables, lag orders, and deterministic model components.

7 Acknowledgments

We thank Michael Binder for his support and guidance during early stages of this project. Moreover, we are grateful for numerous comments and suggestions from the Stata community that helped to improve our `ardl` package. This includes countless email communications, discussions on the Statalist forum, and exchanges of ideas at the 2016 Stata Conference in Chicago, the 2017 and 2018 German Stata Users Group meetings in Berlin and Konstanz, respectively, and the 2018 U.K. Stata Conference in London.

8 Programs and supplemental material

To install the software files as they exist at the time of publication of this article, type

```
. net sj 23-4
. net install st0734      (to install program files, if available)
. net get st0734          (to install ancillary files, if available)
```

9 References

- Adeleye, N., E. Osabuohien, E. Bowale, O. Matthew, and E. Oduntan. 2018. Financial reforms and credit growth in Nigeria: Empirical insights from ARDL and ECM techniques. *International Review of Applied Economics* 32: 807–820. <https://doi.org/10.1080/02692171.2017.1375466>.
- Ang, J. B. 2007. CO₂ emissions, energy consumption, and output in France. *Energy Policy* 35: 4772–4778. <https://doi.org/10.1016/j.enpol.2007.03.032>.
- . 2010. Finance and inequality: The case of India. *Southern Economic Journal* 76: 738–761. <https://doi.org/10.4284/sej.2010.76.3.738>.
- Aruga, K., M. M. Islam, and A. Jannat. 2020. Effects of COVID-19 on Indian energy consumption. *Southern Economic Journal* 12: 5616. <https://doi.org/10.3390/su12145616>.
- Bahmani-Oskooee, M., and T. J. Brooks. 1999. Bilateral *J*-curve between U.S. and her trading partners. *Weltwirtschaftliches Archiv* 135: 156–165. <https://doi.org/10.1007/BF02708163>.

- Bekun, F. V., F. Emir, and S. A. Sarkodie. 2019. Another look at the relationship between energy consumption, carbon dioxide emissions, and economic growth in South Africa. *Science of the Total Environment* 655: 759–765. <https://doi.org/10.1016/j.scitotenv.2018.11.271>.
- Belloumi, M. 2014. The relationship between trade, FDI and economic growth in Tunisia: An application of the autoregressive distributed lag model. *Economic Systems* 38: 269–287. <https://doi.org/10.1016/j.ecosys.2013.09.002>.
- Blackburne, E. F., III, and M. W. Frank. 2007. Estimation of nonstationary heterogeneous panels. *Stata Journal* 7: 197–208. <https://doi.org/10.1177/1536867X0700700204>.
- Büyüksahin, B., and M. A. Robe. 2014. Speculators, commodities and cross-market linkages. *Journal of International Money and Finance* 42: 38–70. <https://doi.org/10.1016/j.jimonfin.2013.08.004>.
- Ciaian, P., M. Rajcaniova, and d. Kancs. 2016. The economics of BitCoin price formation. *Applied Economics* 48: 1799–1815. <https://doi.org/10.1080/00036846.2015.1109038>.
- . 2018. Virtual relationships: Short- and long-run evidence from BitCoin and altcoin markets. *Journal of International Financial Markets, Institutions and Money* 52: 173–195. <https://doi.org/10.1016/j.intfin.2017.11.001>.
- Davis, E. P., and H. Zhu. 2011. Bank lending and commercial property cycles: Some cross-country evidence. *Journal of International Money and Finance* 30: 1–21. <https://doi.org/10.1016/j.jimonfin.2010.06.005>.
- De Vita, G., and A. Abbott. 2004. Real exchange rate volatility and U.S. exports: An ARDL bounds testing approach. *Economic Issues* 9: 69–78.
- Dickey, D. A., and W. A. Fuller. 1979. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74: 427–431. <https://doi.org/10.2307/2286348>.
- . 1981. Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica* 49: 1057–1072. <https://doi.org/10.2307/1912517>.
- Ditzen, J. 2018. Estimating dynamic common-correlated effects in Stata. *Stata Journal* 18: 585–617. <https://doi.org/10.1177/1536867X1801800306>.
- . 2021. Estimating long-run effects and the exponent of cross-sectional dependence: An update to xtdcce2. *Stata Journal* 21: 687–707. <https://doi.org/10.1177/1536867X211045560>.
- Doğan, B., M. B. Jebli, K. Shahzad, T. H. Farooq, and U. Shahzad. 2020. Investigating the effects of meteorological parameters on COVID-19: Case study of New Jersey, United States. *Environmental Research* 191: 110148. <https://doi.org/10.1016/j.envres.2020.110148>.

- Engle, R. F., and C. W. J. Granger. 1987. Co-integration and error correction: Representation, estimation, and testing. *Econometrica* 55: 251–276. <https://doi.org/10.2307/1913236>.
- Enisan, A. A., and A. O. Olufisayo. 2009. Stock market development and economic growth: Evidence from seven sub-Sahara African countries. *Journal of Economics and Business* 61: 162–171. <https://doi.org/10.1016/j.jeconbus.2008.05.001>.
- Esaku, S. 2022. Institutionalized democracy and the shadow economy in the short- and long-run: Empirical analysis from Uganda. *Humanities and Social Sciences Communications* 9: 165. <https://doi.org/10.1057/s41599-022-01128-1>.
- Fatai, K., L. Oxley, and F. G. Scrimgeour. 2004. Modelling the causal relationship between energy consumption and GDP in New Zealand, Australia, India, Indonesia, The Philippines and Thailand. *Mathematics and Computers in Simulation* 64: 431–445. [https://doi.org/10.1016/S0378-4754\(03\)00109-5](https://doi.org/10.1016/S0378-4754(03)00109-5).
- Fedderke, J. W., P. Perkins, and J. M. Luiz. 2006. Infrastructural investment in long-run economic growth: South Africa 1875–2001. *World Development* 34: 1037–1059. <https://doi.org/10.1016/j.worlddev.2005.11.004>.
- Frankel, J., S. L. Schmukler, and L. Servén. 2004. Global transmission of interest rates: Monetary independence and currency regime. *Journal of International Money and Finance* 23: 701–733. <https://doi.org/10.1016/j.jimonfin.2004.03.006>.
- Halicioglu, F. 2009. An econometric study of CO₂ emissions, energy consumption, income and foreign trade in Turkey. *Energy Policy* 37: 1156–1164. <https://doi.org/10.1016/j.enpol.2008.11.012>.
- Hassler, U., and J. Wolters. 2006. Autoregressive distributed lag models and cointegration. *Allgemeines Statistisches Archiv* 90: 59–74. <https://doi.org/10.1007/s10182-006-0221-5>.
- He, F., Z. Hu, W. Zhang, L. Cai, G. Cai, and K. Aoyagi. 2017. Construction and evaluation of two computational models for predicting the incidence of influenza in Nagasaki Prefecture, Japan. *Scientific Reports* 7: 7192. <https://doi.org/10.1038/s41598-017-07475-3>.
- Jalil, A., and S. F. Mahmud. 2009. Environment Kuznets curve for CO₂ emissions: A cointegration analysis for China. *Energy Policy* 37: 5167–5172. <https://doi.org/10.1016/j.enpol.2009.07.044>.
- Johansen, S. 1995. *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford: Oxford University Press. <https://doi.org/10.1093/0198774508.001.0001>.
- Jordan, S., and A. Q. Philips. 2018. Cointegration testing and dynamic simulations of autoregressive distributed lag models. *Stata Journal* 18: 902–923. <https://doi.org/10.1177/1536867X1801800409>.

- Katircioglu, S. T. 2009. Revisiting the tourism-led-growth hypothesis for Turkey using the bounds test and Johansen approach for cointegration. *Tourism Management* 30: 17–20. <https://doi.org/10.1016/j.tourman.2008.04.004>.
- Kirikaleli, D., H. Güngör, and T. S. Adebayo. 2022. Consumption-based carbon emissions, renewable energy consumption, financial development and economic growth in Chile. *Business Strategy and the Environment* 31: 1123–1137. <https://doi.org/10.1002/bse.2945>.
- Kripfganz, S., and V. Sarafidis. 2021. Instrumental-variable estimation of large-T panel-data models with common factors. *Stata Journal* 21: 659–686. <https://doi.org/10.1177/1536867X211045558>.
- Kripfganz, S., and D. C. Schneider. 2020. Response surface regressions for critical value bounds and approximate p -values in equilibrium correction models. *Oxford Bulletin of Economics and Statistics* 82: 1456–1481. <https://doi.org/10.1111/obes.12377>.
- . 2022. ardl: Estimating autoregressive distributed lag and equilibrium correction models. Research Center for Policy Design Discussion Paper 2TUPD-2022-006, Tohoku University. <http://hdl.handle.net/10097/00135205>.
- Lee, G. H. Y., and S. P. Lee. 2014. Childcare availability, fertility and female labor force participation in Japan. *Journal of the Japanese and International Economies* 32: 71–85. <https://doi.org/10.1016/j.jjie.2014.01.002>.
- Morley, B. 2006. Causality between economic growth and immigration: An ARDL bounds testing approach. *Economics Letters* 90: 72–76. <https://doi.org/10.1016/j.econlet.2005.07.008>.
- Murthy, V. N. R., and A. A. Okunade. 2016. Determinants of U.S. health expenditure: Evidence from autoregressive distributed lag (ARDL) approach to cointegration. *Economic Modelling* 59: 67–73. <https://doi.org/10.1016/j.econmod.2016.07.001>.
- Narayan, P. K. 2005. The saving and investment nexus for China: Evidence from cointegration tests. *Applied Economics* 37: 1979–1990. <https://doi.org/10.1080/00036840500278103>.
- Narayan, P. K., and R. Smyth. 2005. Electricity consumption, employment and real income in Australia evidence from multivariate Granger causality tests. *Energy Policy* 33: 1109–1116. <https://doi.org/10.1016/j.enpol.2003.11.010>.
- Narayan, P. K., R. Smyth, and M. Nandha. 2004. Interdependence and dynamic linkages between the emerging stock markets of South Asia. *Accounting and Finance* 44: 419–439. <https://doi.org/10.1111/j.1467-629x.2004.00113.x>.
- Ntanos, S., M. Skordoulis, G. Kyriakopoulos, G. Arabatzis, M. Chalikias, S. Galatsidas, A. Batzios, and A. Katsarou. 2018. Renewable energy and economic growth: Evidence from European countries. *Sustainability* 10: 2626. <https://doi.org/10.3390/su10082626>.

- Obstfeld, M., J. C. Shambaugh, and A. M. Taylor. 2005. The trilemma in history: Tradeoffs among exchange rates, monetary policies, and capital mobility. *Review of Economics and Statistics* 87: 423–438. <https://doi.org/10.1162/0034653054638300>.
- Oteng-Abayie, E. F., and J. M. Frimpong. 2006. Bounds testing approach to cointegration: An examination of foreign direct investment trade and growth relationships. *American Journal of Applied Sciences* 3: 2079–2085.
- Persyn, D., and J. Westerlund. 2008. Error-correction-based cointegration tests for panel data. *Stata Journal* 8: 232–241. <https://doi.org/10.1177/1536867X0800800205>.
- Pesaran, M. H., and Y. Shin. 1998. An autoregressive distributed-lag modelling approach to cointegration analysis. In *Econometrics and Economic Theory in the Twentieth Century: The Ragnar Frisch Centennial Symposium*, ed. S. Steiner, 371–413. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CCOL521633230.011>.
- Pesaran, M. H., Y. Shin, and R. J. Smith. 2001. Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics* 16: 289–326. <https://doi.org/10.1002/jae.616>.
- Phillips, P. C. B., and B. E. Hansen. 1990. Statistical inference in instrumental variables regression with $I(1)$ processes. *Review of Economic Studies* 57: 99–125. <https://doi.org/10.2307/2297545>.
- Sari, R., S. Hammoudeh, and U. Soytas. 2010. Dynamics of oil price, precious metal prices, and exchange rate. *Energy Economics* 32: 351–362. <https://doi.org/10.1016/j.eneco.2009.08.010>.
- Schaffer, M. E. 2010. egranger: Stata module to perform Engle–Granger cointegration tests and 2-step ECM estimation. Statistical Software Components S457210, Department of Economics, Boston College. <https://ideas.repec.org/c/boc/bocode/s457210.html>.
- Shambaugh, J. C. 2004. The effect of fixed exchange rates on monetary policy. *Quarterly Journal of Economics* 119: 301–352. <https://doi.org/10.1162/003355304772839605>.
- Shin, Y., B. Yu, and M. Greenwood-Nimmo. 2014. Modelling asymmetric cointegration and dynamic multipliers in a nonlinear ARDL framework. In *Festschrift in Honor of Peter Schmidt: Econometric Methods and Applications*, ed. W. C. Horrace and R. C. Sickles, 281–314. New York: Springer. https://doi.org/10.1007/978-1-4899-8008-3_9.
- Song, H., S. Lin, S. F. Witt, and X. Zhang. 2011. Impact of financial/economic crisis on demand for hotel rooms in Hong Kong. *Tourism Management* 32: 172–186. <https://doi.org/10.1016/j.tourman.2010.05.006>.
- Stoian, A., and F. Iorgulescu. 2020. Fiscal policy and stock market efficiency: An ARDL bounds testing approach. *Economic Modelling* 90: 406–416. <https://doi.org/10.1016/j.econmod.2019.12.023>.

- Vardi, E., L. Muchnik, A. Conway, and M. Breakstone. 2021. WikiShark: An online tool for analyzing Wikipedia traffic and trends. In *WWW '21: Companion Proceedings of the Web Conference 2021*, 558–571. New York: Association for Computing Machinery. <https://doi.org/10.1145/3442442.3452341>.
- Varona, L., and J. R. Gonzales. 2021. Dynamics of the impact of COVID-19 on the economic activity of Peru. *PLOS ONE* 16: e0244920. <https://doi.org/10.1371/journal.pone.0244920>.
- Wang, Q. 2012. Long-run covariance and its applications in cointegration regression. *Stata Journal* 12: 515–542. <https://doi.org/10.1177/1536867X1201200312>.
- Wang, Y.-S. 2009. The impact of crisis events and macroeconomic activity on Taiwan's international inbound tourism demand. *Tourism Management* 30: 75–82. <https://doi.org/10.1016/j.tourman.2008.04.010>.
- Wolde-Rufael, Y. 2006. Electricity consumption and economic growth: A time series experience for 17 African countries. *Energy Policy* 34: 1106–1114. <https://doi.org/10.1016/j.enpol.2004.10.008>.
- Zhang, H., H. Song, L. Wen, and C. Liu. 2021. Forecasting tourism recovery amid COVID-19. *Annals of Tourism Research* 87: 103149. <https://doi.org/10.1016/j.annals.2021.103149>.
- Zhang, X., E. A. Davidson, D. L. Mauzerall, T. D. Searchinger, P. Dumas, and Y. Shen. 2015. Managing nitrogen for sustainable development. *Nature* 528: 51–59. <https://doi.org/10.1038/nature15743>.

About the authors

Sebastian Kripfganz received his PhD in economics from Goethe University Frankfurt in 2015 and is now a senior lecturer in econometrics at the University of Exeter Business School. He has a strong research interest in methods for estimation and inference in dynamic econometric models.

Daniel Schneider received his PhD in economics from Goethe University Frankfurt in 2013 and is now a statistical programmer and research scientist at the Max Planck Institute for Demographic Research in Rostock. His research interests include time-series analysis and discrete-time multistate models.