



bootUR: An R Package for Bootstrap Unit Root Tests

Stephan Smeekes 
Maastricht University

Ines Wilms 
Maastricht University

Abstract

Unit root tests form an essential part of any time series analysis. We provide practitioners with a single, unified framework for comprehensive and reliable unit root testing in the R package **bootUR**. The package's backbone is the popular augmented Dickey-Fuller test paired with a union of rejections principle, which can be performed directly on single time series or multiple (including panel) time series. Accurate inference is ensured through the use of bootstrap methods. The package addresses the needs of both novice users, by providing user-friendly and easy-to-implement functions with sensible default options, as well as expert users, by giving full user-control to adjust the tests to one's desired settings. Our parallelized C++ implementation ensures that all unit root tests are scalable to datasets containing many time series.

Keywords: bootstrap, R, time series, unit roots.

1. Introduction

In this paper, we introduce the **bootUR** package (Smeekes and Wilms 2023) for R (R Core Team 2023), which implements several bootstrap tests for unit roots. Unit root testing is an essential part of any statistical analysis of time series. Given the crucial role of unit root testing in time series analysis, surprisingly few R packages exist that allow for easy and comprehensive unit root testing. The **bootUR** package aims to fill this gap by offering three major contributions to existing R packages. First, it offers a comprehensive, easy-to-use and reliable set of unit root tests not found as generally in other packages. Second, it offers accurate p values based on bootstrap methods. Third, its functions are not only directly applicable to single time series, but also to datasets consisting of a potentially large set of time series. With these contributions the **bootUR** package provides practitioners with a single source to fill their unit root testing needs.

Proper handling of unit roots is of paramount importance before commencing any form of analysis on the time series of interest. The by far most important use of unit root tests is therefore as a pre-test to determine whether differencing of the series is needed to eliminate the trend and render the time series stationary. Ignoring unit roots, or stochastic trends, essentially invalidates any subsequent statistical analysis: the stochastic trend, and associated non-decaying dependence of the present on the far past of a series, yields standard inference inapplicable. Probably the most famous consequence of ignoring unit roots is the “spurious regression phenomenon”, where one finds seemingly important relations (high R^2 s and highly significant t statistics) between unrelated time series with stochastic trends. These results have a long history and are well-established and extensively documented in the time series literature. A reader new to unit roots may, for instance, consult [Enders \(2008\)](#) for a classical textbook treatment of this spurious regression phenomenon as well as the more general problems associated with unit roots.

Currently, unit root tests are scattered across several packages in the R environment for statistical computing and graphics, making it difficult for a practitioner to find and apply an appropriate and reliable test. The most popular unit root test is the classical augmented-Dickey Fuller (ADF) test ([Dickey and Fuller 1979, 1981](#)). Implementations of the ADF test are incorporated in various packages, in particular **CADFtest** ([Lupi 2009](#)), **fUnitRoots** ([Wuertz, Setz, and Chalabi 2022](#)), **tseries** ([Trapletti and Hornik 2023](#)), and **urca** ([Pfaff 2008](#)). The **mleur** package ([Zhang, Yu, and McLeod 2013](#)) also implements the ADF test, but links to **urca** for this purpose. The package **uroot** ([de Lacalle and Boshnakov 2020](#)) used to have the ADF test implemented but it is no longer supported in the package’s current version, hence disregarded from the overview.

As we will argue in the next section, most “standard” unit root tests, such as the ones implemented in these packages, require practitioners to choose which test to use together with its various model specifications, that may have a major impact on the performance of the unit root tests. As its first major contribution, the **bootUR** package instead implements the user-friendly *union of rejections* principle ([Harvey, Leybourne, and Taylor 2009, 2012](#); [Smeekes and Taylor 2012](#)) that relieves the user from the burden of having to choose the right specification and performs this task automatically.

With the exception of the HEGY seasonal unit root test in the **uroot** package ([de Lacalle and Boshnakov 2020](#)), current R implementations of unit root tests rely on asymptotic inference when returning critical values or p values for the unit root test. Another exception is the repository **URT** ([Mallet 2017](#)), available on GitHub, which includes bootstrap unit root tests. In the remainder, we only focus on packages that are currently maintained on the Comprehensive R Archive Network (CRAN). Unit root tests are very sensitive to size distortions in smaller samples due to, for example, neglected serial correlation ([Schwert 1989](#)). Size distortions due to features such as time-varying volatility even persist asymptotically ([Cavaliere 2005](#)). Bootstrap unit root tests have therefore become a commonly used alternative to asymptotic inference. The bootstrap approximates the exact distribution of the unit root test statistic by repeatedly drawing new samples from the original sample, thereby capturing the features of the time series of interest that affect the distribution of the test, which ensures that bootstrap unit root tests have accurate size properties under very general conditions. The addition of a comprehensive set of bootstrap tests constitutes the second major contribution of our package. The “boot” in **bootUR** stands for bootstrap to emphasize that most unit roots tests we provide rely on various bootstrap methods for constructing p values.

Finally, most datasets contain multiple, sometimes even many, time series to be tested for unit roots, often leading practitioners to apply unit root tests to each time series separately. Such a practice does not only suffer from multiple testing issues, rejecting several tests by chance alone, but also disregards similarities between individual time series which, if exploited, could increase the often limited power of the individual tests. Although some packages provide joint unit root tests for multivariate or panel data (**pdR**, [Tsung-wu 2022](#); **plm**, [Croissant and Millo 2018, 2008](#)), such tests may increase power but do not allow one to determine the properties of individual series. For this goal one would need tests accounting for multiple testing, but proper implementations of multiple testing corrections are currently lacking for unit root tests. Therefore, the third major contribution of **bootUR** is to implement easy tools for applying unit root tests to multivariate time series with automatic multiple testing control. The packages **PANICr** ([Bronder 2016](#)) and **punitroots** ([Kleiber and Lupi 2012](#)) also provide panel unit root tests, but the former has been removed from CRAN and the latter is only available on R-Forge ([Theußl and Zeileis 2009](#)).

With these contributions, the **bootUR** package provides a unified framework for easy and comprehensive unit root testing based on the following philosophy. (1) for novice users, the tests should be easy to implement with sensible default options; (2) those default options should lead to reliable and accurate unit root tests, applicable in general situations; (3) expert users, familiar with the unit root literature, should be able to easily tweak and adjust the tests to their desired settings; (4) all tests should be easily scalable to large datasets without additional effort by the user, thereby providing “automatic” functionality.

To accomplish our philosophy, the package has a simple structure, yet it offers users a wide variety of unit root tests. In particular, unit root tests can *directly* be performed on *single* time series or *multiple* time series. To this end, we deliberately created separate functions that serve these purposes: the `adf()` function provides the traditional “textbook” ADF test ([Dickey and Fuller 1979](#)) based on asymptotic p values, while the functions `boot_adf()` and `boot_union()` can be used for bootstrap-based inference on single time series. The function `boot_ur()` is the main function for multiple time series without multiple-testing control, `boot_sqt()` and `boot_fdr()` are set up for multiple time series with multiple-testing control, and the `boot_panel()` function offers a panel unit root test. For each unit root test, the bootstrap method can be chosen by the end-user. To this end, all functions make use of the universal argument `bootstrap`. Via suitable warning and error messages, user-friendly advice is provided on the (non)-applicability of certain bootstrap methods in certain situations. *Model specifications* (such as deterministic components, lag length selection, detrending methods) are under the user’s full control, with the option to have them implemented automatically according to the union of unit root tests principle which ensures reliable tests across potentially heterogeneous series. Each function contains many options whose syntax is shared across the package, thereby facilitating usability and control by the end-user.

We have also added several functions, based around the core functions above, that aid in the practical implementation of the unit root tests. Most importantly, the function `order_integration()` provides an automatic way to determine the order of integration of each series in a dataset, based on a sequence of one of the aforementioned unit root tests. As it also directly outputs the correctly differenced time series that remove all stochastic trends, it provides the user with the option to conduct the entire unit root pre-analysis with a single command. Additionally, we provide several functions that easily allow the user to assess and visualize properties of the data and outcomes of the tests.

The package is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=bootUR>. In addition, the latest (development) version is available on GitHub at <https://github.com/smeekes/bootUR>. The core of the package is written in C++ (Stroustrup 2013), where we make use of the packages **Rcpp** (Eddelbuettel and François 2011; Eddelbuettel 2013; Eddelbuettel and Balamuta 2017) and **RcppArmadillo** (Eddelbuettel and Sanderson 2014) to facilitate seamless integration with R. In addition we use the **RcppParallel** (Allaire, François, Ushey, Vandenbrouck, Geelnard, and Intel 2023) package, which provides parallel functionality on all platforms to ensure scalability to large datasets. Version 1.0.0 of the **bootUR** package and version 4.2.3 of R were used in this paper.

Adhering to the four points of our philosophy not only requires thoughts on how to implement the tests and design the API, but it also requires a careful choice of the appropriate statistical methods. We therefore first consider the problem from a statistical point of view in Section 2, where we discuss the unit root test for single time series and multiple time series, and in Section 3, where we discuss the bootstrap methods. We then continue with the package’s implementation in Section 4. Section 5 uses three empirical applications to compare **bootUR**’s unit root functions to implementations in other R packages and illustrate its usefulness for practitioners. Section 6 concludes. Appendix A contains a simulation study of the finite sample performance of the methods included in the package.

2. Unit root tests

We first discuss unit root tests for individual time series (Section 2.2), followed by testing multiple series for unit roots (Section 2.3). In our discussion, paralleling Smeekes and Wijler (2020), we do not focus on theory, but on the issues that arise for practitioners when implementing these tests on their time series. For a more extensive and theoretical overview of unit root testing, we refer the interested reader to Choi (2015).

2.1. Unit roots

Consider the case where we have T observations from a time series y_t ($t = 1, \dots, T$) generated according to the data generating process (DGP)

$$y_t = x_t + \beta^\top d_t, \quad x_t = \rho x_{t-1} + u_t, \quad (1)$$

where d_t are deterministic functions of time. In particular, three cases are commonly considered: $d_t = 0$ (no deterministic components), $d_t = 1$ (intercept only), and $d_t = (1, t)^\top$ (intercept and linear trend). The error process u_t is allowed to be serially correlated and heteroskedastic. The presence of serial correlation in u_t has to be accounted for in inference. Typically, u_t is modelled as an invertible infinite order linear process, for instance as

$$u_t = \sum_{j=0}^{\infty} \psi_j \epsilon_{t-j} = \sum_{j=1}^{\infty} \phi_j u_{t-j} + \epsilon_t,$$

where ϵ_t is typically assumed to be a martingale difference sequence. This linearity motivates the use of adding lagged differences of the time series to account for the serial dependence, as in the classical ADF test (Dickey and Fuller 1979). However, Paparoditis and Politis (2018)

	Package Function	bootUR		CADFtest	fUnitRoots	tseries	urca	
		adf()	boot_adf()	CADFtest()	unitrootTest()	adf.test()	ur.df()	ur.ers()
Deterministic components	None	✓	✓	✓	✓		✓	
	Intercept only	✓	✓	✓	✓		✓	✓
	Intercept & linear trend	✓	✓	✓	✓	✓	✓	✓
	User control	✓	✓	✓	✓		✓	✓
Detrending	OLS	✓	✓	✓	✓	✓	✓	
	QD		✓					✓
Lag selection	AIC	✓	✓	✓			✓	
	BIC	✓	✓	✓			✓	
	MAIC	✓	✓	✓				
	MBIC	✓	✓					
	HQC			✓				
	Rescaled Heuristic	✓	✓			✓		
	User control	✓	✓	✓	✓	✓	✓	✓
Inference	Critical value table				✓	✓	✓	✓
	Asymptotic p value	✓		✓	✓	✓		
	Bootstrap p value		✓					
Other features	Covariate augmentation			✓				
	Explosive alternative					✓		

Table 1: Overview ADF-test functionalities in existing R packages.

show that ADF-type approaches are valid under much more general forms of dependence in u_t .

We focus on testing whether or not y_t contains a unit root, that is on testing

$$H_0 : \rho = 1 \text{ against } H_1 : |\rho| < 1$$

in Equation 1. Under the null hypothesis of a unit root, y_t contains a stochastic trend, and equivalently y_t is being said to be integrated of order 1 ($I(1)$), while the alternative postulates that y_t is integrated of order 0 ($I(0)$), which is generally taken as synonymous to y_t being stationary. Here “integrated of order d ” means that y_t should be differenced d times to achieve a process that does not contain a stochastic trend anymore. Although stationary is generally used as synonym for $I(0)$, an $I(0)$ process can still be non-stationary, for instance through a shift in the variance. Despite this distinction, we follow tradition and use “ $I(0)$ ” and “stationary” interchangeably.

2.2. Individual unit root tests

To test the null hypothesis of a unit root, the classical ADF test (Dickey and Fuller 1979, 1981) remains the pre-dominant choice in practice. For this reason it also forms the backbone of the **bootUR** package. However, even in its most basic form, practitioners are required to make several non-trivial choices that affect its performance. Table 1 summarizes these choices and indicates how the various R packages address each of them. In this section, we first discuss the ADF test and the choices that need to be made, before discussing the union of unit root tests principle proposed by Harvey *et al.* (2009, 2012) which alleviates many of the concerns.

ADF test

Let Δ be the difference operator defined as $\Delta y_t := y_t - y_{t-1}$. If no deterministic components are present, the ADF regression is given by

$$\Delta y_t = \gamma y_{t-1} + \sum_{j=1}^p \phi_j \Delta y_{t-j} + \varepsilon_t, \quad t = p+1, \dots, T, \quad (2)$$

where the lagged differences of y_t are added to the regression to capture the serial correlation present in u_t . Testing the null of a unit root then boils down to testing the significance of the parameter γ in Equation 2.

If the time series y_t is suspected to have deterministic components as well, testing becomes more complicated. The traditional one-step procedure adds the relevant deterministic components directly in Equation 2. However, the one-step detrending does not lend itself to be extended to other forms of detrending such as the popular generalized least squares (GLS) or quasi-differencing (QD) framework (Elliott, Rothenberg, and Stock 1996) discussed later. Most contemporaneous literature (see e.g., Chang and Park 2002; Harvey *et al.* 2009) therefore considers a two-step approach. Here, a first stage regression is run of y_t on the deterministic components d_t , and in the second stage the ADF regression

$$\Delta y_t^d = \gamma y_{t-1}^d + \sum_{j=1}^p \phi_j \Delta y_{t-j}^d + \varepsilon_t^d, \quad t = p+1, \dots, T, \quad (3)$$

is run on the residuals of the first stage regression, $y_t^d = y_t - \hat{\beta}^\top d_t$, commonly referred to as the detrended time series. When estimated by ordinary least squares (OLS), the one-step and two-step detrending are generally used interchangeably and will give very similar results, although they are not equal, and may even have different limit behavior (cf. Harvey *et al.* 2009, Remark 5). For the traditional `adf()` test, **bootUR** offers the user the choice to implement one-step as an alternative to the default two-step detrending. This option is mainly provided to allow replicability of results obtained through other software or textbook examples. The two-step procedure is implemented for all bootstrap unit root tests in the **bootUR** package, including the `boot_adf()` function for single time series.

Elliott *et al.* (1996) proposed the DF-GLS test, in which $\hat{\beta}$ is obtained by a GLS regression where the (near-)unit root in y_t is first removed by quasi-differencing; the regression is then performed by OLS for $y_t - (1 - \frac{c}{T}) y_{t-1}$ on $d_t - (1 - \frac{c}{T}) d_{t-1}$, where c is a parameter that determines how close to differencing the GLS step is; Elliott *et al.* (1996) recommend that $c = 7$ for the case $d_t = 1$ and $c = 13.5$ for the case $d_t = (1, t)^\top$ to yield tests with good power properties. The DF-GLS, or alternatively, DF-QD test, is often considered to be more powerful than the ADF test, although as shown by Müller and Elliott (2003) and Harvey *et al.* (2009) *inter alia*, the QD test is only more powerful if the initial condition, that is the deviation of the start of the time series from equilibrium, is small. When the initial condition is large, the standard OLS-detrended ADF test is considerably more powerful. To avoid confusion with a “proper” GLS estimation that also takes into account higher-order serial dependence and heteroskedasticity, we follow Harvey *et al.* (2009) and refer to this test as the quasi-differenced (QD) test rather than GLS. Among other R packages, only **urca** provides a version of the DF-QD test, see Table 1.

While both options are implemented in the function `boot_adf()` for varying choice of d_t , the issue of including deterministic components presents the practitioner with two difficult

choices: which deterministic components to include, and how to perform the detrending. These choices can have a major impact on the performance. If too few deterministic components are included, deterministic trends are detected as stochastic trends, and the test becomes inconsistent. On the other hand, adding too many deterministic components reduces the power of the test considerably, and should also be avoided (Harvey *et al.* 2009). Yet, all unit root tests in current R packages ask the user to make these choices. Similarly, the initial condition is unobservable, such that the user has to choose which detrending method to use.

The **bootUR** package is, to the best of our knowledge, the first R package which does not force the user to make these choices, but instead offers the function `boot_union()` for a data-driven alternative via the union of rejections principle introduced by Harvey *et al.* (2009, 2012). Before discussing this in detail, we first turn to the third choice a user has to make: selecting the lag length p in Equation 3.

The lag length choice concerns a trade off between size distortions incurred from including too few lags to capture all serial correlation, and power loss incurred from including too many lags. Although theory (and some R packages such as **tseries**) generally assume p to be a deterministic function of the sample size, in practice data-driven selection allows for a more precise choice of p that can achieve a better trade off between size and power. **bootUR** offers a set of information criteria for automatic data-driven lag length selection. Information criteria trade off model fit (through the residual sum of squares) and overfitting (through a penalty on the number of parameters). The lag length is estimated as

$$\hat{p} := \underset{p_{\min} \leq k \leq p_{\max}}{\operatorname{argmin}} IC(k), \quad IC(k) = \ln \hat{\sigma}_k^2 + k \frac{C_T}{T}, \quad (4)$$

where $\hat{\sigma}_k^2 := (T - p_{\max})^{-1} \sum_{t=p_{\max}+1}^T (\hat{\varepsilon}_{k,t}^d)^2$ with $\hat{\varepsilon}_{k,t}^d$ the OLS residuals from the ADF regression with lag length k in Equation 3, and C_T is a penalty function that differs according to the information criterion used. We consider two penalties: one corresponding to the Akaike information criterion (AIC; $C_T = 2$) and the other to the Bayesian information criterion (BIC; $C_T = \ln T$).

Next to the original criteria, **bootUR** also implements their modified variants proposed by Ng and Perron (2001). These modifications are specifically motivated for lag length selection in the ADF regression. They are given by

$$MIC(k) := \ln \hat{\sigma}_k^2 + k \frac{C_T + \xi_T(k)}{T},$$

where $\xi_T(k) := (\hat{\sigma}_k^2)^{-1} \hat{\gamma}^2 \sum_{t=p_{\max}+1}^T (y_{t-1}^d)^2$. The lag length is then estimated as in Equation 4, with $IC(k)$ replaced by $MIC(k)$. The modified AIC (MAIC) is obtained by taking $C_T = 2$, the modified BIC (MBIC), by taking $C_T = \ln T$. Ng and Perron (2001) show that the *MICs* yield large size improvements over the *ICs* for the purpose of unit root testing. Perron and Qu (2007) recommend to always use the *MICs* with the OLS rather than QD-detrended data (even if the unit root test itself makes use of QD detrending) since this improves the test's power properties; **bootUR** follows this recommendation. In addition, there are various seemingly minor aspects of how the lag selection is implemented that influence its performance, such as how many observations are used to calculate the residual sum of squares. Ng and Perron (2005) provide a detailed study and guidelines for these choices; **bootUR** implements the scheme they recommend as optimal.

Cavaliere, Phillips, Smeekes, and Taylor (2015) find that heteroskedasticity affects the performance of information criteria, leading to less accurate choices of p and consequent power loss of the unit root tests. They propose rescaled information criteria, where the time series y_t is rescaled with a nonparametric estimate of its (time-varying) standard deviation, thereby eliminating the heteroskedasticity. The information criterion is then applied to this rescaled series. These rescaled ICs are generally more powerful in the presence of heteroskedasticity, yet very similar to the original ones without. **bootUR** therefore performs the rescaling by default (with the option not to consider it) since it is a safe choice and relieves the user of the burden to check whether heteroskedasticity is present.

Union of rejections test

As mentioned above, choosing the right deterministic components to include and the right detrending method to use, is crucial to obtain tests with good power properties. Harvey *et al.* (2009, 2012) take a different approach based on a very simple principle. Roughly speaking, for both specification issues, we have one powerful test and one not powerful test. A logical step would therefore be to perform both tests and reject whenever one of them rejects the null hypothesis - the logic being that the one rejecting is then the powerful one. With two tests performed simultaneously, one must control for multiple testing and adjust the tests with a Bonferroni-type adjustment to control size at the desired level. Harvey *et al.* (2009) introduced this union of rejections idea for the two specification issues separately, while Harvey *et al.* (2012) combined the two approaches to consider a union of four tests – intercept only or intercept with trend in combination with OLS or QD detrending – that guards against both uncertainty over the trend and the initial condition.

This makes the union test a safe option for quick or automatic unit root testing where careful manual specification is not viable, and makes it therefore very suitable for **bootUR**'s philosophy that the default option provides a reliable and accurate test, for which no in depth knowledge is needed about either the data or the applicability of various unit root tests. Moreover, it scales easily to large datasets with many series, where careful manual considerations about these specifications are not possible regardless of the expertise of the user.

The **bootUR** package implements the bootstrap version of the union test developed by Smeekes and Taylor (2012), which uses the bootstrap both for determining the appropriate size correction and for obtaining the test's p values. The test statistic takes the form

$$UR = \min \left(\frac{s}{c_{QD}^{\mu*}(\alpha)} QD^{\mu}, \frac{s}{c_{QD}^{\tau*}(\alpha)} QD^{\tau}, \frac{s}{c_{ADF}^{\mu*}(\alpha)} ADF^{\mu}, \frac{s}{c_{ADF}^{\tau*}(\alpha)} ADF^{\tau} \right), \quad (5)$$

where ADF and QD are the ADF and QD detrended tests, and superscript μ and τ respectively indicating whether the series are demeaned or detrended. The critical values $c_i(\alpha)$ are determined in a preliminary bootstrap step as the individual level α critical values of the four tests; weighting with their inverse is needed to bring the four tests on the same scale. The variable s is a scaling factor to which the statistics are scaled. Any $s < 0$ suffices to preserve the left-tail rejection region; in **bootUR** we scale to -1 . This bootstrap union test is made available through the function `boot_union()`. Finally, note that this union-based approach still requires one to select the lag lengths in each of the four ADF regressions. To this end, any of the four information criteria, AIC, BIC, MAIC and MBIC can be used.

Of course, various other unit root tests such as the Phillips and Perron (1988) (PP) test, the KPSS (Kwiatkowski, Phillips, Schmidt, and Shin 1992) stationarity test, the HEGY test

(Hylleberg, Engle, Granger, and Yoo 1990) for seasonal unit roots, or the covariate-augmented Dickey-Fuller test of Hansen (1995) exist. While these tests undoubtedly have a lot of value, we intentionally do not implement them in **bootUR** to avoid overloading the user with choices in tests and options. Instead, we focus on the ADF test which is by far the most popular in practice. Practitioners interested in applying these other unit root tests are recommended to consult the packages **fUnitRoots**, **stats**, **tseries** or **urca** for the PP test, and **fUnitRoots**, **tseries** or **urca** for the KPSS test. We also want to highlight the **uroot** package which focuses on the HEGY seasonal unit root test, and the **CADFtest** package which offers the covariate-augmented ADF test, which exploits correlation with known stationary covariates to improve power. These two packages are entirely devoted to these specific unit root tests and already provide excellent functionality for the practitioner in need of these specific tests.

2.3. Multiple unit root tests

Practitioners often make use of several time series in their analysis, and typically need to test all for unit roots. While performing a unit root test for each series separately is normal practice for a small number of time series, this becomes more complicated if the number of series is large. First, performing many unit root tests simultaneously suffers from multiple testing issues as the probability of incorrect classifications increases with the number of performed tests. Second, we would like to exploit the similarity between different time series to improve the power of the unit root tests, in particular if the time dimension is relatively small.

In the **bootUR** package, we consider three different ways to approach the testing problem with multiple time series. First, the simplest option of ignoring the test multiplicity issue by just performing unit root tests separately for each series. To this end, the function `boot_ur()` from **bootUR** can be used. Second, we consider the traditional approach of panel unit root tests, where one pools the information in all series to obtain a more powerful test. The function `boot_panel()` offers such a test. Third, we can consider individual tests but then with appropriate control of multiple testing error rates. **bootUR** considers two such tests, namely `boot_sqrt()` and `boot_fdr()`.

Despite the large literature on this topic, software implementations for multiple unit roots are mostly lacking. While there is some support for panel unit root testing as discussed hereafter, methods to control multiple testing in the context of unit root testing are, to the best of our knowledge, not available. While several general purpose multiple testing packages exist, using these in a proper way with unit root tests requires considerable effort and expertise from the user. For instance, some standard corrections may be overly conservative, such as the Bonferroni correction, or only applicable under specific conditions on the dependence, such as the method by Benjamini and Hochberg (1995) to control the false discovery rate. As argued by for instance Romano, Shaikh, and Wolf (2008b), bootstrap methods for controlling multiple testing allow for general forms of dependence and avoid being too conservative. However, such bootstrap methods need to be integrated with the unit root testing, which is the approach taken in **bootUR**.

Throughout this section, we use the following notation. Consider N time series for which one would like to test the presence of a unit root. We denote their respective individual unit root test statistics by UR_i , $1 \leq i \leq N$. Typically these would correspond to one of the tests discussed in Section 2.2. Without loss of generality, we assume that rejections occur for small values of the test statistic.

Panel unit root tests

Panel unit root tests view the multiple time series as a coherent panel dataset, and exploit the similarity between such time series to pool the information in them and achieve more powerful tests. They have a long tradition in econometrics, see e.g., [Breitung and Pesaran \(2008\)](#) or [Choi \(2015\)](#) for reviews. A typical panel unit root test has the null hypothesis that all series have a unit root. Rejection of this null hypothesis is then typically interpreted as evidence that a “significant proportion” of the series is stationary. However, how large that proportion is, or which series are stationary is not revealed by the test. [Pesaran \(2012\)](#) therefore suggests to use panel unit root tests as an initial screening tool for analyzing multiple series; if the panel unit root test rejects the null, this indicates that the individual series need to be examined further; if not, treating the full dataset as $I(1)$ may be a reasonable choice.

We implement the bootstrap group-mean (GM) test of [Palm, Smeekes, and Urbain \(2011\)](#)

$$GM = \frac{1}{N} \sum_{i=1}^N UR_i,$$

in the function `boot_panel()` which is based on averaging the unit root test statistics UR_i ($1 \leq i \leq N$) of the N individual time series. This test is valid under very general forms of dependence within the dataset, yet does not require modelling it. Panel unit root tests are scarcely available for R users. Currently, only two packages with panel unit root tests, namely **plm** and **pdR**, are being maintained. The package **plm** was the first to offer panel unit root tests and provides the tests introduced in [Maddala and Wu \(1999\)](#); [Choi \(2001\)](#); [Levin, Lin, and Chu \(2002\)](#); [Im, Pesaran, and Shin \(2003\)](#). However, none of them allow for cross-sectional dependence (see [Kleiber and Lupi 2011](#) for a discussion). The package **pdR** offers the panel unit root test of [Chang \(2002\)](#) and a panel version of the seasonal unit root test in [Hylleberg et al. \(1990\)](#).

Multiple testing

Multiple testing methods allow practitioners to determine the order of integration for each series in their dataset while controlling the overall probability of making false rejections. First note that in order to properly rank and compare tests for different series, the individual test statistics should have the same marginal distributions. Then, the ranking

$$UR_{(1)} \leq \dots \leq UR_{(R)} \leq UR_{(R+1)} \leq \dots \leq UR_{(N)}, \quad (6)$$

corresponds to a ranking from “most significant” to “least significant”, when the i -th order statistic of UR_1, \dots, UR_N is denoted by $UR_{(i)}$. To ensure the comparability of these statistics, nuisance parameters need to be eliminated from the distribution of the test statistics. **bootUR** does this automatically for the union test by scaling all test statistics in Equation 5 towards $s = -1$; if the user chooses to set specifications manually, it is up to the user to choose them such that any nuisance parameters are eliminated.

The goal is to find an appropriate cut-off point R such that the null of a unit root is rejected for all statistics less than or equal to $UR_{(R)}$, while it is not rejected for all statistics larger. How this threshold is determined, depends on how one controls for multiple testing. **bootUR** implements two ways to do this: the sequential testing procedure of [Smeekes \(2015\)](#), which also encompasses the Step-M method of [Romano and Wolf \(2005\)](#) to control the familywise

error rate, and the false discovery rate (FDR) controlling approach of Romano *et al.* (2008b); Moon and Perron (2012).

Sequential quantile test

Smeekes (2015) proposes the Bootstrap Sequential Quantile Test (BSQT) for multiple unit root testing, that acts as an intermediate between panel unit root testing and full multiple testing control. The method proceeds by sequentially testing groups of time series for unit roots, where the user decides the group sizes. At step 1, we test whether the first p_1 series are stationary. Here “first” does not refer to the order in the dataset (which is arbitrary), but to the most significant tests as found via Equation 6. If the null hypothesis that all p_1 units have a unit root cannot be rejected, the test stops. If we do observe a rejection, we move on to the second group where we test if the first p_2 are stationary. However, as we already concluded that the first p_1 units are stationary, in this second step the actual test is whether the next $p_2 - p_1$ units are stationary as well. We continue this testing procedure until no rejection is observed anymore or we tested all series in the dataset. The BSQT can be performed by using the function `boot_sqt()`.

More formally, let p_1, \dots, p_K be the number of series to be tested as stationary in each of the steps $k = 1, \dots, K$. In the sequential step k we then test

$$H_0 : p_{k-1} \text{ series are } I(0); \quad \text{against} \quad H_1 : p_k \text{ series are } I(0).$$

As the first test should have as H_0 that all units are $I(0)$, $p_0 = 0$ by default. Furthermore, $p_K = N$ to complete the testing procedure. The number of steps K and the intermediate numbers p_1, \dots, p_{K-1} can be chosen by the practitioner. Instead of thinking in terms of p_k series, it may be easier to think in terms of quantiles q_k , and set $p_k = [q_k N]$. A practitioner may for instance think “I want to split my series in 10 equally-sized groups.” In that case the practitioner simply sets $q_k = 0.1k$.

The choice of $\{p_k\}$ does require input and consideration from the user. It can be chosen based on the nature of the dataset and the desired level of precision of the practitioner. Smeekes (2015) shows that if p_k units are found to be $I(0)$, the probability that the true number of stationary series lies outside the interval $[p_{k-1}, p_{k+1}]$ is at most the chosen significance level of the test. Finding that p_k series are $I(0)$ should therefore be interpreted as finding that the number of $I(0)$ series is in the interval $[p_{k-1}, p_{k+1}]$. In the end, if p_2, \dots, p_{K-1} are chosen sensibly and not spaced too far apart, the series that lie in the “uncertain interval” are likely those series which are “just about” significant, and correspond to time series with a ρ parameter very close to 1. The practical consequences of incorrect classification of these series are typically small, as their behavior makes them fit reasonably well in both classes of $I(1)$ and $I(0)$ series.

One special case worth mentioning – set as the default in `boot_sqt()` – is when we set $p_k = k$, such that each series gets tested sequentially. Smeekes (2015) shows that in this case the BSQT method coincides with the popular Step-M method of Romano and Wolf (2005) to control the familywise error rate (FWE). The FWE is defined as the probability of making at least one false rejection, and is typically controlled via the Bonferroni or Holm (1979) approach.

FDR-controlled test

The false discovery rate (FDR), originally proposed by [Benjamini and Hochberg \(1995\)](#), is defined as $FDR = \mathbb{E} \left[\frac{F}{R} \mathbb{1}(R > 0) \right]$, where R denote the total number of rejections, and F the number of false rejections. It is more appropriate for larger N than the FWE, as it aims to control the proportion of false rejections to the total, rather than the probability of a single false rejection. [Romano, Shaikh, and Wolf \(2008a\)](#) develop a bootstrap method to control the FDR, and show that unlike the classical way to control FDR, the bootstrap is appropriate under very general forms of dependence between series. [Moon and Perron \(2012\)](#) applied this method to unit root testing, and it is their method that is implemented in the `boot_fdr()` function of the **bootUR** package. FDR control can be combined with any unit root test specification considered in [Section 2.2](#).

To decide on whether to use BSQT or FDR control, relative sample sizes can be considered. The Monte Carlo comparison of [Smeekes \(2015\)](#) reveals that the FDR-controlling test is somewhat more accurate when the sample size T is at least of equal magnitude as the number of time series N , whereas the BSQT method is clearly preferable when T is much smaller than N , since the FDR-controlling test then suffers from a lack of power. While [Smeekes \(2015\)](#) investigates the performance of (various) bootstrap unit root tests based on the standard ADF test with intercept only, we conduct an additional Monte Carlo simulation experiment in [Appendix A](#) that compares the asymptotic `adf()` test to the bootstrap tests `boot_ur()`, `boot_sqt()` and `boot_fdr()` based on the union of rejections test, which is the default in the package **bootUR**.

3. Bootstrap-based inference

We mainly rely on bootstrap methods to obtain critical values and/or p values for all of the unit root tests discussed in [Section 2](#). In the **bootUR** package, six bootstrap methods are implemented: the sieve bootstrap (SB), moving block bootstrap (MBB), sieve wild bootstrap (SWB), dependent wild bootstrap (DWB), block wild bootstrap (BWB) and autoregressive wild bootstrap (AWB). Their properties are summarized in [Table 2](#), and discussed more extensively below. As immediately apparent from [Table 2](#), any “off-the-shelf” time series bootstrap method may be used to counteract size distortions arising from neglected serial correlation ([Schwert 1989](#)); whereas a wild bootstrap method is needed to deal with general forms of

Bootstrap method	Serial correlation	Heteroskedasticity	Cross-sectional dependence	Unbalancedness
SB	✓			
MBB	✓		✓	
SWB	✓	✓		✓
DWB	✓	✓	✓	✓
BWB	✓	✓	✓	✓
AWB	✓	✓	✓	✓

Table 2: Bootstrap methods and their ability to deal with serial correlation, general forms of heteroskedasticity, cross-sectional dependence and unbalancedness of the data.

heteroskedasticity (Cavaliere and Taylor 2008, 2009a). General forms of cross-sectional dependence can be captured by any bootstrap method apart from the sieve ones.

Next to correcting the size of unit root tests, bootstrap methods have other advantages. First, the bootstrap offers an automatic p value. Second, the bootstrap directly allows for implementation of multiple testing techniques such as those discussed above. Third, it guards against misspecification and uncertainty regarding the lag length selection in the ADF. As **bootUR** re-selects the lag lengths within the bootstrap replications, it automatically takes effects of lag selection into account. This, coupled with the fact that the bootstrap captures any dependence missed by the lagged differences in the ADF regression, adds another layer of protection to the tests.

3.1. Sieve bootstrap

The sieve bootstrap (SB) has been extensively considered in the context of unit root testing; see among others Psaradakis (2001), Chang and Park (2003), Paparoditis and Politis (2005), Palm, Smeekes, and Urbain (2008) and Smeekes (2013). It estimates the dependence as an autoregressive (AR) process, resamples the residuals of the AR fit, and then re-applies the AR model recursively to place the dependence back into the bootstrap sample. **bootUR** determines the required order of the AR model by the order of the ADF model, combining these in a single step as they should conceptually coincide.

While it is able to capture general forms of serial dependence (Kreiss, Paparoditis, and Politis 2011), it is mostly suited for tests on single time series. Smeekes and Urbain (2014b) show that it is not suited to capture general forms of cross-sectional dependence, making it invalid for joint or multiple testing. The **bootUR** package therefore advises to only use it for unit root testing of a single series or on multivariate series without multiple testing control, throwing a warning to alert the user otherwise. When still applied multivariately, users should also realize that each time series is required to be observed over the same periods, which we refer to as balanced datasets. This often forces practitioners to delete observations for series that have been observed for a longer period. The reason for this limitation is that resampling step of the sieve bootstrap would reshuffle the missing values, creating bootstrap sample with “holes” in it.

3.2. Moving block bootstrap

The moving block bootstrap (MBB) is another traditional bootstrap method that has not only been used for univariate unit root testing in Paparoditis and Politis (2003), but also for multivariate unit root testing in Moon and Perron (2012) and Smeekes (2015), as well as for panel unit root testing in Palm *et al.* (2011). It works by dividing the data in overlapping blocks of data and resampling those blocks to create bootstrap series by laying them end-to-end. The blocks are taken in the time dimension and encompass all series. The block length ℓ is set automatically by **bootUR** as a function of the sample size, following a rule proposed by Palm *et al.* (2011) that they showed to perform well in many different circumstances. However, it is easily adjusted by the user to experiment with different lengths and assess the sensitivity of the results for varying block lengths.

The MBB can accommodate any form of serial dependence as long as it “fits” into an adequately sized block, which is a wide class. Unlike the SB, the MBB can also handle general forms of dependence between series, including but not limited to common factor structures.

From a practical point of view an attractive feature is that it can be applied without requiring one to model the serial and/or cross-sectional dependence. [Palm *et al.* \(2011\)](#) show its validity for mixed $I(1)/I(0)$ panel datasets under such general forms of dependence. The MBB still has, however, two disadvantages: it cannot handle unbalanced datasets and is sensitive to unconditional heteroskedasticity. To handle both issues, users should switch to one of the wild bootstrap methods available in **bootUR**.

3.3. Sieve wild bootstrap

The wild, or multiplier, bootstrap ([Mammen 1993](#); [Davidson and Flachaire 2008](#)) is known to be robust against general forms of heteroskedasticity, however it cannot handle serial dependence. Nonetheless, if combined with a sieve bootstrap, we get the best of both worlds. That is, by replacing the resampling step applied to the residuals of the AR model with a multiplication by independent and identically distributed (iid) random variables with mean zero and variance one, we obtain the sieve wild bootstrap (SWB). [Cavaliere and Taylor \(2009a,b\)](#) and [Smeekes and Taylor \(2012\)](#) among others apply this sieve wild bootstrap for bootstrap unit root testing. The method is perfectly suited to individual unit root testing, but due the AR estimation, suffers from the same inability to capture complex dependence across series as explained by [Smeekes and Urbain \(2014b\)](#) for the SB. For the generation of the iid random variables, **bootUR** follows the unit root papers cited above and uses the normal distribution.

3.4. Dependent, block and autoregressive wild bootstrap

The three remaining bootstrap methods implemented in the package are all wild bootstrap methods adjusted to deal with dependence. However unlike the SWB, here the multiplicative random variables themselves are adjusted to be dependent over time. This setup allows these bootstrap methods to capture complex serial and cross-series dependence structures as well as unconditional heteroskedasticity. In addition, no resampling takes place for the dependent wild bootstrap, such that missing values “stay in their place” which makes the method applicable to unbalanced datasets.

The three wild bootstrap methods only differ in how the multiplier variables are made time-dependent. The dependent wild bootstrap method (DWB), originally introduced by [Shao \(2010\)](#), draws random variables from a T -dimensional $N(0, \Sigma)$ distribution, where the elements in Σ decrease with the distance between them. [Shao \(2010\)](#) proposes to use a kernel function to achieve this, along with a bandwidth ℓ which ensures that variables more than ℓ time points apart are independent. This way ℓ has a similar interpretation as the block length in the MBB. [Rho and Shao \(2019\)](#) and [Smeekes and Urbain \(2014a\)](#) study the DWB for unit root testing, the latter focusing on multivariate settings.

We consider two more variations. The block wild bootstrap (BWB) ([Shao 2011](#); [Zhang and Cheng 2014](#)) is a direct alternative to the MBB, where for each block of size ℓ , we use the same multiplier variable, and the variables are independent between blocks. The autoregressive wild bootstrap (AWB) ([Smeekes and Urbain 2014a](#); [Friedrich, Smeekes, and Urbain 2020](#)) generates the multiplier variables as a first-order autoregressive process. Unlike the BWB and DWB who have a block length ℓ tuning parameter, the tuning parameter of the AWB is the first-order AR parameter. To be able to use the same tuning parameter ℓ , we use the conversion formula proposed by [Smeekes and Urbain \(2014a\)](#) and [Friedrich *et al.* \(2020\)](#)

that writes the AR parameter as a function of ℓ , though **bootUR** also allows to set the AR parameter directly. The default setting for ℓ in **bootUR** uses the same rule as for the MBB, which was also tested for the three wild bootstrap methods by [Smeekes and Urbain \(2014a\)](#). They also provide theoretical results on the validity of these methods under general forms of dependence and heteroskedasticity.

For completeness, in [Algorithm 1](#) we present the six bootstrap methods and their role in the general bootstrap algorithm. Note that the outcome of the bootstrap algorithm is a collection of bootstrap unit root test statistics UR_i^b for the series $i = 1, \dots, N$ and bootstrap replications $b = 1 \dots, B$. How these are then used depends on the multiple testing approach taken. For instance, if we ignore multiple testing, we simply calculate the bootstrap p values

$$p_i^* = \frac{1}{B} \sum_{b=1}^B I(UR_i^b < UR_i), \quad i = 1, \dots, N.$$

For the BSQT and FDR tests more involved processing is needed; for details we refer to [Smeekes \(2015\)](#) and [Romano *et al.* \(2008a\)](#) respectively.

4. An introduction to the bootUR package

The **bootUR** package has a simple structure with 13 user-accessible functions. [Section 4.1](#) presents three functions to check if the data are suitable to be bootstrapped. [Sections 4.2](#) and [4.3](#) introduce the seven core functions for unit root testing on respectively individual and multiple time series. [Section 4.4](#) presents three useful functions for determining the order of integration of each series in a particular dataset.

The package's functions will now be presented together with examples of their specific use. To this end, we make use of the dataset "MacroTS" which contains a collection of 20 macroeconomic time series taken from Eurostat and comes with the package. A complete description of the data can be obtained by typing `?MacroTS` in R. The following examples require that the **bootUR** package and the data have been loaded. As random number generation is required to draw bootstrap samples, we first set the seed of the random number generator to obtain reproducible results:

```
R> library("bootUR")
R> data("MacroTS")
R> set.seed(155776)
```

4.1. Checking data suitability

The package **bootUR** offers three functions, namely `check_missing_insample_values()`, `find_nonmissing_subsample()` and `plot_missing_values()`, to check if a particular dataset is suitable to be bootstrapped. While the bootstrap tests do not work with missing data, unbalanced datasets are generally allowed (see [Table 2](#)).

The function `check_missing_insample_values()` checks if a particular dataset contains missing values. Its usage is extremely simple, as it only requires the data as input,


```
R> check_missing_insample_values(MacroTS)
```

Algorithm 1: Multivariate bootstrap unit root tests.

```

1 Let  $y_{i,t}^d = y_{i,t} - d_t \hat{\beta}$ , where  $\hat{\beta}$  is obtained by OLS;
2 for  $i \in \{1, \dots, N\}$  do
3   Estimate Equation 3 for  $\{y_{i,t}\}_{t=1}^T$ , determining  $p$  by an appropriate criterion,
   obtaining estimates  $(\hat{\gamma}_i, \hat{\phi}_{i,1}, \dots, \hat{\phi}_{i,p})$ ;
4   Set  $\hat{u}_{i,t} = \Delta y_{i,t}^d - \hat{\rho}_i y_{i,t-1}^d$  and  $\hat{\varepsilon}_{i,t} = \Delta y_{i,t}^d - \hat{\rho}_i y_{i,t-1}^d - \sum_{j=1}^p \hat{\phi}_{i,j} \Delta y_{i,t-j}^d$ , with
    $y_{-p+1}, \dots, y_0 = 0$ ;
end
5 for  $b \in \{1, \dots, B\}$  do
6   if SB then
7     Generate  $s_1, \dots, s_T$  from a uniform distribution on  $\{1, \dots, T\}$ ;
8     Set  $u_{i,t}^b = \sum_{j=1}^p u_{i,t-j}^b + \varepsilon_{i,t}^b$  with  $\varepsilon_{i,t}^b = \hat{\varepsilon}_{i,s_t}$  and  $u_{-p+1}, \dots, u_0 = 0$  for  $i = 1, \dots, N$ 
     and  $t = 1, \dots, T$ ;
   else if MBB then
9     Generate  $s_1^b, \dots, s_{\lceil T/\ell \rceil}^b$  from a uniform distribution on  $\{1, \dots, T - \ell + 1\}$ ;
10    for  $m \in \{1, \dots, \lceil T/\ell \rceil\}$  do
11      Set  $u_{i,t}^b = \hat{\varepsilon}_{i,s_t^b}$  for  $i = 1, \dots, N$  and  $t = (m-1)\ell + 1, \dots, m\ell$ ;
    end
   else if SWB then
12    Generate  $\xi_1^b, \dots, \xi_T^b$  from a  $N(0, 1)$  distribution;
13    Set  $u_{i,t}^b = \sum_{j=1}^p u_{i,t-j}^b + \varepsilon_{i,t}^b$  with  $\varepsilon_{i,t}^b = \xi_t^b \hat{\varepsilon}_{i,t}$  and  $u_{-p+1}, \dots, u_0 = 0$  for
      $i = 1, \dots, N$  and  $t = 1, \dots, T$ ;
   else if DWB then
14    Generate  $\zeta_1^b, \dots, \zeta_T^b$  from a  $N(0, 1)$  distribution and let
      $\xi^b = (\xi_1^b, \dots, \xi_T^b)^\top = \Sigma^{1/2} (\zeta_1^b, \dots, \zeta_T^b)^\top$  with  $(\sigma_{s,t})_{s,t=1}^T = K\left(\frac{|s-t|}{\ell}\right)$  for the kernel
     function  $K(\cdot)$  defined in Shao (2010);
15    Set  $u_{i,t}^b = \xi_t^b \hat{u}_{i,t}$  for  $i = 1, \dots, N$  and  $t = 1, \dots, T$ ;
   else if BWB then
16    Generate  $\xi_1^b, \dots, \xi_{\lceil T/\ell \rceil}^b$  from a  $N(0, 1)$  distribution;
17    for  $m \in \{1, \dots, \lceil T/\ell \rceil\}$  do
18      Set  $u_{i,t}^b = \xi_m^b \hat{u}_{i,t}$  for  $i = 1, \dots, N$  and  $t = (m-1)\ell + 1, \dots, m\ell$ ;
    end
   else if AWB then
19    Generate  $\zeta_2^b, \dots, \zeta_T^b$  from a  $N(0, 1 - \gamma^2)$  distribution and let  $\xi_t^b = \gamma \xi_{t-1}^b + \zeta_t^b$  with
      $\xi_1^b \sim N(0, 1)$ ;
20    Set  $u_{i,t}^b = \xi_t^b \hat{u}_{i,t}$  for  $i = 1, \dots, N$  and  $t = 1, \dots, T$ ;
21    Set  $y_{i,t}^b = \sum_{s=1}^t u_{i,s}^b$ ;
22    Let  $UR_i^b = UR(y_{i,1}^b, \dots, y_{i,T}^b)$ , where  $UR(\cdot)$  denotes the chosen unit root test.
end

```

Figure 1: Missingness for the dataset `MacroTS`.

which can be a vector, matrix, data frame or in time series format (e.g., `ts`, `zoo` or `xts`). It returns an N -dimensional Boolean vector which indicates for each series whether missing values are present (`TRUE`) or not (`FALSE`).

If a dataset contains series with different starting and end points, the bootstrap methods SWB, DWB, BWB and AWB can still be used. The function `find_nonmissing_subsample()` lets users check the start and end points of each series as follows:

```
R> sample_check <- find_nonmissing_subsample(MacroTS)
R> sample_check
```

```
$range
      GDP_BE GDP_DE GDP_FR GDP_NL GDP_UK CONS_BE CONS_DE
first      1      1      1      5      1      1      1
last     100    100    100    100    100    100    100
      CONS_FR CONS_NL CONS_UK HICP_BE HICP_DE HICP_FR
first      1      5      1      9      9      9
last     100    100    100    100    100    100
      HICP_NL HICP_UK UR_BE UR_DE UR_FR UR_NL UR_UK
first      9      9      1      1      1      1      1
last     100    100    100    100    100    100    100

$all_equal
[1] FALSE
```

The output slot `range` returns a $(2 \times N)$ matrix displaying the first and last non-missing value for each series, the logical slot `all_equal` provides a quick check to see if all time series have the same non-missing indices (`TRUE`) or not (`FALSE`).

Finally, to display missingness in the dataset, we can use

```
R> plot_missing_values(MacroTS, show_names = TRUE)
```

which displays present cell values in green, missing values at the start or end (“Unbalanced NAs”) in purple and internal missing values in red, see Figure 1. Only the latter are problematic for the wild bootstrap methods, while the purple values also need to be avoided for the resampling-based bootstraps.

4.2. Individual unit root tests

bootUR has one function, `adf()`, that implements a traditional asymptotic ADF test and two functions to perform a bootstrap unit root test on a single series: `boot_adf()` for a standard bootstrap ADF test and `boot_union()` for a bootstrap union test.

We start by discussing the standard `adf()` function followed by its bootstrap analogue `boot_adf()` with the many options users can tweak. As **bootUR** shares its syntax across the various functions, the majority of function arguments remains identical across **bootUR**’s functions, which facilitates usability and control by the end-user. In the remainder, we therefore only highlight the differences compared to the `adf()` and/or `boot_adf()` function.

Asymptotic ADF test

To perform a standard ADF test with asymptotic p value, **bootUR** offers the `adf()` function:

```
adf(data, data_name = NULL, deterministic = "intercept", min_lag = 0,
    max_lag = NULL, criterion = "MAIC", criterion_scale = TRUE,
    two_step = TRUE)
```

The minimum required input is `adf(data)`, where the time series `data` can be a vector or a time series object. The optional argument `data_name` allows users to name their data for use in the printed output. The remaining set of arguments relates to the ADF regression. The deterministic components can be tweaked via the argument `deterministic` with options "none", "intercept", "trend", lag length selection in the ADF regression is controlled via the arguments `min_lag` and `max_lag` that respectively set the minimum and maximum lag length. For the latter, we use the sample size-based default $\lfloor 12 \cdot (T/100)^{1/4} \rfloor$. The information criterion can be selected via the argument `criterion` (with options "AIC", "BIC", "MAIC", "MBIC") and the argument `criterion_scale` lets practitioners choose to use the rescaled information criteria of [Cavaliere et al. \(2015\)](#). To overwrite data-driven lag selection with a pre-specified lag length, users can simply put both `min_lag` and `max_lag` equal to the desired lag length. Finally, the argument `two_step` allows users to specify whether OLS detrending should be performed following the one-step or two-step procedure, with the latter being the default.

We apply the asymptotic variant of the ADF test with both one-step and two-step OLS detrending to Dutch gross domestic product (GDP). An intercept and linear time trend are added as deterministic components:

```
R> GDP_NL <- MacroTS[, 4]
R> adf_out1 <- adf(data = GDP_NL, deterministic = "trend",
+   two_step = FALSE)
```

```
R> adf_out2 <- adf(data = GDP_NL, deterministic = "trend",
+   two_step = TRUE)
```

The function returns a ‘bootUR’ class object that inherits from the class ‘htest’. The outcome of the unit root test can be easily read from the console via the `print()` function. We slightly adjust the standard R output from the ‘htest’ class to make it better fit for our unit root tests. Practitioners interested in retrieving the standard R output from the ‘htest’ class, can use `stats:::print.htest(adf_out1)` instead.

```
R> print(adf_out1)
```

```
One-step ADF test (with trend) on a single time series
```

```
data: GDP_NL
null hypothesis: Series has a unit root
alternative hypothesis: Series is stationary

      estimate largest root statistic p-value
GDP_NL          0.9489    -2.487  0.3339
```

```
R> print(adf_out2)
```

```
Two-step ADF test (with trend) on a single time series
```

```
data: GDP_NL
null hypothesis: Series has a unit root
alternative hypothesis: Series is stationary

      estimate largest root statistic p-value
GDP_NL          0.9471    -2.515  0.3202
```

The tests differ slightly numerically but both give the same evidence regarding the presence or absence of a unit root. Return values such as sample estimate $\hat{\gamma}$, test statistic and p value can be easily retrieved via:

```
R> adf_out1$estimate
```

```
gamma
-0.05106074
```

```
R> adf_out1$statistic
```

```
tstat
-2.486711
```

```
R> adf_out1$p.value
```

```
p-value
0.333873
```

Next to the standard slots in ‘htest’, the object contains the additional slots `details` and `specifications`. The output slot `details` contains an overview of the sample estimate $\hat{\gamma}$ in the slot `individual estimates`, the test statistic in the slot `individual statistics`, the p value in `individual p-values`, and the selected number of lagged differences in the ADF regression in the slot `selected lags`. The output slot `specifications` contains the specifications (deterministics, minimum and maximum lag order, criterion for lag selection, with or without re-scaling and de-trending procedure; one-step or two-step) of the unit root test that is performed.

Bootstrap ADF test

To perform a standard ADF bootstrap unit root test on a single series, the `boot_adf()` function can be used. The function is structured as follows:

```
boot_adf(data, data_name = NULL, bootstrap = "AWB", B = 1999,
  block_length = NULL, ar_AWB = NULL, deterministic = "intercept",
  detrend = "OLS", min_lag = 0, max_lag = NULL, criterion = "MAIC",
  criterion_scale = TRUE, show_progress = TRUE, do_parallel = TRUE,
  cores = NULL)
```

Many arguments used for `adf()` re-appear in the `boot_adf()` function and play the same role. We focus the discussion on the new arguments. The arguments related to the bootstrap specifications include the bootstrap method (`bootstrap`) and number of bootstrap replications (`B`). If a user chooses the bootstrap method "MBB", "DBB", "BWB" or "AWB", the desired block length can be controlled via the argument `block_length`. By default, we use `block_length = \lfloor 1.75 \cdot T^{1/3} \rfloor`, as recommended in Palm *et al.* (2011). While for the first three, this argument concerns the genuine block length, for the latter, the block length is transformed into an autoregressive parameter `ar_AWB` via the formula $0.01^{(1/1)}$ as in Smeekes and Urbain (2014a); this can be overwritten by setting `ar_AWB` directly. The following set of arguments relates to the ADF regression and most arguments have a one-to-one correspondence with those in the `adf()` function. For `boot_adf()`, the user can additionally choose the type of detrending via `detrend`, with options "OLS" (default) and "QD".

The argument `show_progress` provides live progress updates on the bootstrap. The latter is particularly useful for large values of the argument `B`. Finally, the option `do_parallel = TRUE` allows the bootstrap to be executed in parallel via **RcppParallel**; the argument `cores` allows users to specify how many cores should be used for the parallel loops. By default, all but two cores are used.

We illustrate the bootstrap ADF test on Dutch GDP, with the sieve bootstrap (`bootstrap = "SB"`) as used by Palm *et al.* (2008) and Smeekes (2013). We use the same specification for the deterministic components as for the `adf()` test above and detrending is done via OLS.

```
R> boot_adf(data = GDP_NL, bootstrap = "SB",
+   deterministic = "trend", detrend = "OLS")
```

```
SB bootstrap OLS test (with intercept and trend)
on a single time series
```



```

data: GDP_NL
null hypothesis: Series has a unit root
alternative hypothesis: Series is stationary

      estimate largest root statistic p-value
GDP_NL          0.9471      -2.515  0.1311

```

As for `adf()`, the function returns an object of class ‘`bootUR`’ and ‘`htest`’. The bootstrap ADF test returns the same value of the sample estimate and test statistic as for the standard ADF test with two-step OLS detrending as it should, only the p value differs. Also in this case, the test indicates no rejection of the unit root.

Union of rejections test

To perform a bootstrap union unit root test on a single series, the `boot_union()` function can be used. It shares all its arguments with `boot_adf()` except for `deterministics` and `detrend` which are omitted since `boot_union()` performs tests for all the possible options, then combines the outcomes of the four unit root tests as in Equation 5 to produce a single p value. By combining these tests, the sample estimate is not available and hence set to `NA` in the output. The function introduces one new argument `union_quantile`, which allows the user to choose the α in Equation 5 used to construct the critical values that serve as weights in the union. If the user has a particular significance level in mind when performing the test, α should be set to this. If not, the default 0.05 is used.

The bootstrap union test for Dutch GDP with the sieve wild bootstrap as proposed by [Smeekes and Taylor \(2012\)](#) can be obtained via

```

R> boot_union(data = GDP_NL, bootstrap = "SWB")

      SWB bootstrap union test on a single time series

```

```

data: GDP_NL
null hypothesis: Series has a unit root
alternative hypothesis: Series is stationary

      estimate largest root statistic p-value
GDP_NL          NA      -0.6701  0.6433

```

Finally, for the union test, the output slot `details` contains an overview of the estimates, test statistics, p values and selected lag orders of the four tests, namely the ADF unit root tests with intercept only or intercept with trend in combination with OLS or GD trending.

4.3. Multiple unit root tests

Below, we discuss the various approaches `bootUR` offers to approach the testing problem with multiple series.

Separate unit root tests

To perform individual ADF tests on multiple time series simultaneously without multiple testing control, the function `boot_ur()` can be used:

```
boot_ur(data, data_name = NULL, bootstrap = "AWB", B = 1999,
        block_length = NULL, ar_AWB = NULL, level = NULL, union = TRUE,
        union_quantile = 0.05, deterministic = NULL, detrend = NULL, min_lag = 0,
        max_lag = NULL, criterion = "MAIC", criterion_scale = TRUE,
        show_progress = TRUE, do_parallel = TRUE, cores = NULL)
```

Compared to the syntax of `boot_adf()`, it has three additional arguments. The argument `union` controls whether a bootstrap union test is used (`TRUE`) or not (`FALSE`). If `union = TRUE` (default), the arguments `deterministic` and `detrend` are ignored, and a warning message is returned if the user would have provided specifications for these anyway. If set to `FALSE`, the deterministic components and detrending methods can be specified as for the `boot_adf()` function. The argument `union_quantile` is the same as for `boot_union()`. Finally, the optional argument `level` offers the user the possibility to give a significance level for the tests to be performed. Doing so adds a slot in the output object which contains a Boolean vector indicating for each series if the unit root null hypothesis is rejected or not. This may be convenient for automated post-processing if the unit root tests are performed as an initial step in an analysis. If a significance level is given, `union_quantile` is overwritten with this value.

Furthermore, since the bootstrap is performed for all series simultaneously, the bootstrap methods "SB" or "MBB", that cannot handle unbalanced datasets, should not be used. If the user were to specify these anyway, the function will revert to splitting the bootstrap up and performing it separately for each time series. A warning message is then returned to alert the user. If a vector (or univariate time series) instead of a matrix (multivariate time series) is given for `data`, a single unit root test is performed; in this case the function acts as an alternative to `boot_adf()` and `boot_union()`. In fact, internally, these functions call `boot_ur()`.

We illustrate the function's usage by performing individual ADF tests with the "MBB" bootstrap on the first five series of the unbalanced dataset `MacroTS`, which correspond to the real GDP in Belgium, Germany, France, the Netherlands and the United Kingdom respectively.

```
R> boot_ur_out <- boot_ur(data = MacroTS[, 1:5], bootstrap = "MBB")
```

Warning message:

```
In check_inputs(data = data, boot_sqt_test = boot_sqt_test,
boot_ur_test = boot_ur_test, : Missing values cause resampling bootstrap to
be executed for each time series individually.
```

The warning message alerts the user about the resampling "MBB" bootstrap method being unable to handle unbalanced datasets and the corrective action that is taken to this end.

```
R> boot_ur_out
```

```

      MBB bootstrap union test on each individual series
      (no multiple testing correction)

```

```

data: MacroTS[, 1:5]
null hypothesis: Series has a unit root
alternative hypothesis: Series is stationary

```

```

Tests performed on each series:

```

	estimate	largest root	statistic	p-value
GDP_BE	NA	-0.8135	0.36618	
GDP_DE	NA	-1.1076	0.08804	
GDP_FR	NA	-0.6301	0.76188	
GDP_NL	NA	-0.8211	0.41371	
GDP_UK	NA	-0.7207	0.53877	

An object of class ‘bootUR’ and ‘mult_hctest’, for multiple hypothesis testing, is returned. We constructed the class ‘mult_hctest’ as an extension of ‘hctest’ to make it suitable for multiple hypothesis testing. In addition to the return values as for the class ‘hctest’, it has additional output slots `series.names`, which contains the names of the individual time series, and `rejections`, which optionally contains the aforementioned vector with rejections in case this is requested.

Looking at the results, none of the time series is stationary at the 5% significance level, as can be seen from the printed p values. Information on, for instance, the test statistics and p values can also be accessed directly via:

```
R> boot_ur_out$statistic
```

GDP_BE	GDP_DE	GDP_FR	GDP_NL	GDP_UK
-0.8135022	-1.1076021	-0.6301366	-0.8210610	-0.7207147

```
R> boot_ur_out$p.value
```

GDP_BE	GDP_DE	GDP_FR	GDP_NL	GDP_UK
0.36618309	0.08804402	0.76188094	0.41370685	0.53876938

Panel unit root test

To perform a panel unit root test, the function `boot_panel()` can be used. It shares its syntax with `boot_ur()`. Unlike for the latter, usage of the "MBB" or "SB" bootstrap methods for a panel unit root test on unbalanced datasets will result in an error— not a warning —since the unbalancedness cannot be reverted. Therefore, users should switch to one of the wild bootstrap methods or balance their dataset. Besides, sieve bootstrap methods can be used, but they are not suited to capture general forms of dependence across units (see Table 2). The code therefore warns users against their usage.

We illustrate the usage of the panel unit root test on the five GDP time series with the "DWB" bootstrap of [Shao \(2010\)](#) and [Rho and Shao \(2019\)](#):

```
R> boot_panel(data = MacroTS[, 1:5], bootstrap = "DWB")
```

```
Panel DWB bootstrap group-mean union test
```

```
data: MacroTS[, 1:5]
null hypothesis: All series have a unit root
alternative hypothesis: Some series are stationary

              estimate largest root statistic p-value
MacroTS[, 1:5]          NA    -0.8372  0.2986
```

Since the null is not rejected at any reasonable level, treating all five GDP series as $I(1)$ is reasonable.

Sequential quantile test

To perform the BSQT for multiple unit root testing, the function `boot_sqt()` should be used. It has one additional argument compared to the `boot_panel()` function, namely `steps` which sets the group sizes. These can either be set in units or in quantiles. To split the series in, for instance, K equally sized groups, use `steps = 0:K/K`. By the convention of [Smeekes \(2015\)](#), the first entry of the vector should be equal to zero, while the second entry indicates the end of the first group, and so on. If the initial zero value or the final value (N or 1 for quantiles) are accidentally omitted, the function automatically adds them back. The default `steps = 0:NCOL(data)` corresponds to the Step-M method of [Romano and Wolf \(2005\)](#). Regarding the bootstrap methods, the same warning and error messaging as for the `boot_panel()` apply. To set the significance level of the individual tests performed sequentially, the argument `SQT_level` can be set.

We illustrate the BSQT on the five GDP series with the "AWB" (default) bootstrap method of [Smeekes and Urbain \(2014a\)](#) and [Friedrich et al. \(2020\)](#):

```
R> boot_sqt(data = MacroTS[, 1:5])
```

```
AWB bootstrap sequential quantile union test
```

```
data: MacroTS[, 1:5]
null hypothesis: Series has a unit root
alternative hypothesis: Series is stationary
```

```
Sequence of tests:
```

```
      H0: # I(0) H1: # I(0)  tstat p-value
Step 1          0          1 -1.046  0.3347
```

The details on the number of series p_k to be tested as stationary in step k under H_0 and H_1 (first two columns) is printed to the console, together with, the test-statistic and p value (last two columns) for each of the sequential steps until no rejection occurs. The latter information is also accessible through the output slot `details$$SQT`, information on the (non) rejection of the unit root null for each of the series separately can be accessed via the slot `rejections`.

FDR-controlled test

To perform a multiple unit root test by controlling the FDR, the function `boot_fdr()` should be used. Its arguments are the same as for the other multivariate unit root tests, with the addition of the argument `FDR_level` which sets the FDR level. We illustrate it here with the "BWB" bootstrap method of [Shao \(2011\)](#) and [Smeekes and Urbain \(2014a\)](#):

```
R> boot_fdr(data = MacroTS[, 1:5], bootstrap = "BWB")
```

```
      BWB bootstrap union test with false discovery rate control
```

```
data: MacroTS[, 1:5]
null hypothesis: Series has a unit root
alternative hypothesis: Series is stationary
```

```
Sequence of tests:
      tstat critical value
GDP_DE -0.9814          -1.346
```

The procedure developed by [Romano *et al.* \(2008b\)](#); [Moon and Perron \(2012\)](#) does not provide p values, therefore critical values are returned instead here.

4.4. Determining series' order of integration

Finally, **bootUR** offers three useful functions for determining the order of integration of each series in dataset: `order_integration()`, `diff_mult()` and `plot_order_integration()`. The main function is `order_integration()` which applies the "Pantula principle" ([Pantula 1989](#)) to determine the order of integration of each series

```
order_integration(data, data_name = NULL, max_order = 2, method = "boot_ur",
  level = 0.05, plot_orders = FALSE, ...)
```

The argument `max_order` sets the maximum order of integration that should be considered for each series. Generally the default of two should generally suffice, with series of order three or higher only very rarely occurring in practice. The user can choose the unit root test through the argument `method` depending on whether a single ("adf", "boot_adf" or "boot_union") or a multiple time series ("boot_ur", "boot_sqt" or "boot_fdr") is considered. All arguments used in these functions can also be passed on via `order_integration()`. As setting the order of integration automatically requires "yes/no" decisions about the presence of unit roots, a significance level has to be set via the argument `level`. For "boot_sqt" or "boot_fdr" this is automatically transformed to `SQT_level` and `FDR_level`, respectively.

The Pantula principle works as follows. It starts by setting $d = \text{max_order} - 1$ and testing for a unit root on the $\Delta^d y_t$ series. The series for which the unit root null cannot be rejected are classified as $I(d+1)$ and subsequently removed from the dataset. In the next step, $d = d - 1$ and the remaining series are tested and classified accordingly. Under the default `max_order = 2`, this second round involves testing the series in levels and classifying them as either $I(1)$ (if the unit root null is not rejected) or $I(0)$ (if the null is rejected).

The function returns an object of class ‘`bootUR`’ and ‘`order_integration`’ with two accessible elements. The slot `diff_data` contains a matrix whose columns are $\Delta^{d_i} y_{i,t}$ with d_i indicating the order of integration of the i -th series ($i = 1, \dots, N$). This matrix is generated by the user-accessible function `diff_mult(data, d)`, where `data` is the original dataset and `d` is an N -dimensional vector indicating each series’ order of integration. It contains the same number of rows as the original dataset (since the default setting `keep_NAs = TRUE` in `diff_mult()` is used), thereby indicating lost observations as missing. It can be tweaked if a practitioner directly makes use of this function. Note that the object of the input series (e.g., matrix, data frame or time series object) is preserved for the differenced series. The output slot `order_int` makes the vector containing the found orders of integration available to the end-user.

Finally, if the argument `plot_orders` in the function `order_integration()` is set to `TRUE`, a plot is provided which displays each series’ order of integration. To this end, it uses the function `plot_order_integration()` which takes the object that `order_integration()` outputs as its input. This function is also made accessible if the end-user wishes to further adjust the display of the variable names, legend and colours through its optional arguments `show_names`, `show_legend`, `names_size`, `legend_size` and `cols`.

5. Applications

We illustrate the methods on three datasets: the income data used in [Smeekes \(2015\)](#), the `MacroTS` dataset which comes with the package, and the FRED-QD dataset, which is widely used for macro-economic analysis.

5.1. Replication of [Smeekes \(2015\)](#)

We replicate the empirical application of [Smeekes \(2015\)](#) based on income data from the Panel Study of Income Dynamics (PSID) used in [Pesaran \(2007\)](#). The income dataset represents a balanced micro-panel with $T = 22$ time periods and $N = 181$ units split into three subsamples consisting of college graduates (CLG, $N = 58$), high school graduates (HSG, $N = 87$), and high school dropouts (HSD, $N = 36$).

We use the functions `boot_ur()`, `boot_sqt()` and `boot_fdr()` with the MBB bootstrap to determine the proportion of stationary units in the PSID dataset with the same specifications

	<code>boot_ur()</code>	<code>boot_sqt()</code>	<code>boot_fdr()</code>
Total rejections ($N = 181$)	49	60	1
Proportion of rejections	0.27	0.33	0.01
Rejection in subsamples			
CLG ($N = 58$)	13	16	1
Proportion	0.22	0.28	0.02
HSG ($N = 87$)	25	29	0
Proportion	0.29	0.33	0.00
HSD ($N = 36$)	11	15	0
Proportion	0.31	0.42	0.00

Table 3: Rejections of the unit root null hypothesis in the PSID dataset.

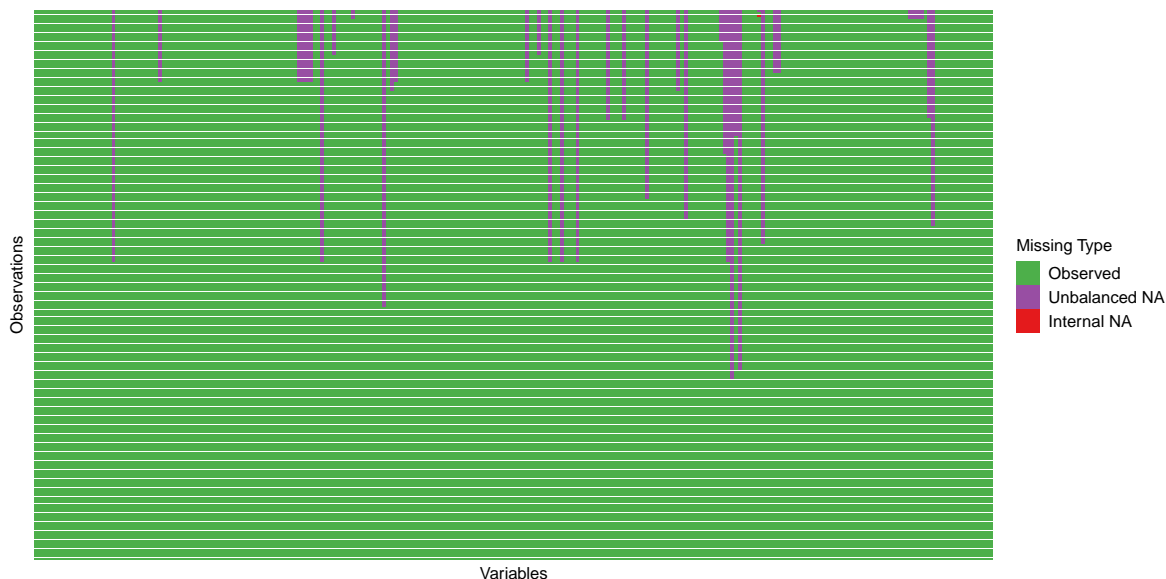


Figure 2: Missingness for the dataset FRED-QD.

as in [Smeekes \(2015\)](#): the bootstrap replications B are set to 4999, the block size to five, and the lags in the ADF regression are selected using the rescaled MAIC criterion with a maximal lag length of three.

Table 3 presents the rejection results of a unit root. `boot_ur()` and `boot_fdr()` indicate that around 30% of the units are stationary. As discussed in [Smeekes \(2015\)](#), `boot_fdr()` returns only few units as stationary and thereby likely suffers from low power due to the small sample size. Finally, in line with [Smeekes \(2015\)](#), the proportion of stationary units are fairly similar across subsamples.

5.2. MacroTS

The `MacroTS` dataset contains $N = 20$ macro-economic time series collected from Eurostat (<https://ec.europa.eu/eurostat/data/database>) and is included in the package. Quarterly observations from 1995–2019 ($T = 100$) are available on GDP, consumption, inflation and unemployment for Belgium, Germany, France, the Netherlands and the United Kingdom. The dataset is unbalanced, see Figure 1.

5.3. FRED-QD

This is a quarterly version of the monthly Federal Reserve Economic Data (FRED) database introduced in [McCracken and Ng \(2016\)](#). It contains $N = 248$ macro-economic time series and was imported into R using the commands

```
R> FRED_url <- paste0("https://files.stlouisfed.org/files/htdocs/",
+ "fred-md/quarterly/2020-06.csv")
R> FRED_QD <- read.csv(FRED_url)
```

This paper uses the data from 1959 Quarter 2 to 2019 Quarter 4 ($T = 244$) to avoid possible

Series	MacroTS	FRED-QD
In first differences	< 0.001	< 0.001
In levels	0.097	0.197

Table 4: p values of the panel unit root test on all series in differences and all series in levels, for both the MacroTS and FRED-QD dataset.

structural breaks due to the COVID-19 pandemic in 2020. If a researcher wishes to import the up-to-date version of the dataset, `2020-06.csv` should be changed to `current.csv`. As can be seen from Figure 2, the dataset contains one internal NA, since the third observation of variable 188 (UMCSENTx: Consumer Expectations) is missing while the second observation is not. **bootUR** cannot handle internal missing values but this can be easily fixed by setting the second observation to NA, which results in the first three observations of this variable being “unbalanced NAs” that can be handled by **bootUR**. The resulting dataset then contains 38 macro-economic indicators with missing values at the start of the sample. Finally, note that all FRED-QD series have been classified into $I(0), I(1), I(2)$ by the transformation codes provided in McCracken and Ng (2020). However, the authors themselves indicate several discrepancies between these codes and the outcome of unit root tests. We therefore use the transformation codes as a benchmark for the classifications obtained through the unit root tests but do not necessarily consider the classification closest to theirs to be the best.

Since some of the macro-economic series are likely to be $I(2)$, we use the `order_integration()` function (with its defaults) to implement the Pantula principle. All unit root tests in the **bootUR** are performed with their default settings, which means that union tests are performed with the AWB bootstrap method, and lag length selection is done via the re-scaled MAIC. Throughout this section, a significance level of 5% is used. For `boot_sqt()`, the default (i.e., Step-M method) is reported as well as results for evenly spaced 0.1 quantiles (`steps = 0:10/10`, for MacroTS), and 0.05 quantiles (`steps = 0:20/20`, for FRED-QD).

We compare **bootUR**’s unit root tests to the R packages reported in Table 1. We hereby use the following specifications: For the function `CADFtest()` (package **CADFtest**), we perform ADF regressions with intercept and trend (`type = "trend"`), and lag length selection with MAIC (`criterion = "MAIC"`) thereby considering a maximum of $\lfloor 12 \cdot (T/100)^{1/4} \rfloor$ lags, set via the argument `max.lag.y`. These lag length specifications correspond to the defaults used in **bootUR**. For `unitrootTest()` (package **fUnitRoots**), we perform ADF regressions with intercept and trend (`type = "ct"`). By default, one lagged difference is included. For `adf.test()` (package **tseries**), we use its default settings which implies ADF-regressions with intercept and trend and the number of lags fixed to $\lfloor (T-1)^{1/3} \rfloor$, a deterministic function of the sample size. For `ur.df()` (package **urca**), we use ADF-regressions with intercept and trend (`type="trend"`), lag length selection via AIC (`selectlags = "AIC"`), thereby considering a maximum of $\lfloor 12 \cdot (T/100)^{1/4} \rfloor$ lags, set via the argument `lags`. Finally, for `ur.ers()` (package **urca**), we use an intercept and trend for detrending (`model = "trend"`). By default, four lagged differences are included in the ADF-regression. Unlike the other packages, **urca** only comes with critical values to judge the significance of the unit root test, the p value is not reported, see Table 1. As discussed in Lupi (2009), the p value reported under `summary()` is computed using the t distribution, which is incorrect under the unit root null. Finally, only the packages **CADFtest** and **fUnitRoots** can handle missing values, for the other packages, we removed missing values prior to performing the unit root tests.

Non-bootstrap tests		Bootstrap tests	
<code>CADFtest()</code>	134.27	<code>boot_ur()</code>	343.75
<code>unitrootTest()</code>	1.70	<code>boot_sqrt()</code>	306.68
<code>adf.test()</code>	0.50	<code>boot_sqrt_20step()</code>	343.01
<code>ur.df()</code>	5.71	<code>boot_fdr()</code>	351.72
<code>ur.ers()</code>	0.90		
Benchmark: <code>adf()</code>		1.33 (absolute time in seconds)	

Table 5: Computing time of the unit root tests relative to the benchmark `adf()` test for the FRED-QD dataset on a MacBook Pro (macOS 11.6.7) with Quad-Core Intel Core i5 2,4 GHz processor under the default parallel computing settings.




Figure 3: Classification of the MacroTS dataset into $I(0)$, $I(1)$, $I(2)$.

Before applying the various unit root tests to the two datasets, we perform the `boot_panel()` (with default settings) to all series taken in first differences, and to all series in levels. Table 4 reports the p values of the panel unit root tests. For both datasets, the panel unit root tests on the series in first differences indicates that the unit root null is rejected, thereby indicating that a “significant proportion” of the series is stationary in first differences (hence not $I(2)$). The panel unit root test on the series in levels indicates non-rejection of the unit root null.

To shed further light on the order of integration for each of the individual series, the bootstrap unit root tests are applied and compared to the implementations from other R packages. Figure 3 presents the obtained orders of integration on the MacroTS dataset, Figures 4 and 5 on the FRED-QD dataset. Note that on such a large dataset, performing the bootstrap tests – even twice in `order_integration()` if $d = 2$ – can become computationally intensive. Table 5 gives the computation times relative to our `adf()` function, as recorded on a MacBook Pro (macOS 11.6.7) with Quad-Core Intel Core i5 2,4 GHz processor. While the `adf()` function took just over a second to run, it took roughly between 7 and 8 minutes to perform each of the bootstrap tests. There is therefore a price to pay for bootstrap accuracy. A simple way to save computation time, while still preserving reasonable accuracy, is to reduce the number of bootstrap replications, by altering the argument `B`.

Globally speaking, most unit root tests agree upon a series’ classification into $I(0)$, $I(1)$, $I(2)$, which is comforting. Still, several interesting remarks can be made. First, the results of




Figure 4: Classification of the first half of time series in the FRED-QD dataset into $I(0)$, $I(1)$, $I(2)$.

`boot_ur()` are fairly similar to `adf()`, `boot_sqrt()` and `boot_fdr()` but it classifies a considerable amount of series as $I(0)$ instead of $I(1)$ on the FRED-QD dataset. This illustrates that ignoring multiple testing can quickly lead to a considerable number of misclassifications on such large datasets. Second, among the `boot_sqrt()` procedures, the default Step-M method tends to classify more series as $I(2)$ than the other procedure. On the smaller MacroTS dataset the two versions of the `boot_sqrt()` show more agreement than on the larger FRED-

6. Summary

This paper presents the R package **bootUR** that provides a unified framework for bootstrap unit root testing on single and multiple time series. To this end, the package builds upon the popular ADF test with a union of rejections principle. Unlike existing packages on unit root tests, **bootUR** (i) provides a large collection of easy-to-use, fully-controllable and reliable unit root tests, including the union of rejections test which is set as default to enable quick, automatic unit root testing, (ii) ensures accurate inference through bootstrap methods with easy-to-read output (including p values), (iii) allows for testing the presence of unit roots in datasets containing many time series by relying on fast C++ implementations.

Acknowledgments

We thank the editors and referees for their thorough review and highly appreciate their constructive comments which substantially improved the quality of the package and paper. The first author was financially supported by the Netherlands Organization for Scientific Research (NWO) under grant number 452-17-010, the second author by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 832671. We gratefully acknowledge the comments and checks provided by Robert Adamek, Rui Jorge Almeida, Nalan Baştürk, Caterina Schiavoni and Étienne Wijler on earlier versions of the package. All remaining errors are our own.

References

- Allaire JJ, François R, Ushey K, Vandenbrouck G, Geelnard M, Intel (2023). **RcppParallel**: *Parallel Programming Tools for Rcpp*. R package version 5.1.7, URL <https://CRAN.R-project.org/package=RcppParallel>.
- Benjamini Y, Hochberg Y (1995). “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing.” *Journal of the Royal Statistical Society B*, **57**(1), 289–300. doi:10.1111/j.2517-6161.1995.tb02031.x.
- Breitung J, Pesaran MH (2008). “Unit Roots and Cointegration in Panels.” In *The Econometrics of Panel Data*, pp. 279–322. Springer-Verlag.
- Bronder S (2016). **PANICr**: *PANIC Tests of Nonstationarity*. R package version 1.0.0, URL <https://CRAN.R-project.org/src/contrib/Archive/PANICr/>.
- Cavaliere G (2005). “Unit Root Tests under Time-Varying Variances.” *Econometric Reviews*, **23**(3), 259–292. doi:10.1081/etc-200028215.
- Cavaliere G, Phillips PCB, Smeekes S, Taylor AMR (2015). “Lag Length Selection for Unit Root Tests in the Presence of Nonstationary Volatility.” *Econometric Reviews*, **34**(4), 512–536. doi:10.1080/07474938.2013.808065.
- Cavaliere G, Taylor AMR (2008). “Bootstrap Unit Root Tests for Time Series with Nonstationary Volatility.” *Econometric Theory*, **24**(1), 43–71. doi:10.1017/s0266466608080043.

- Cavaliere G, Taylor AMR (2009a). “Bootstrap M Unit Root Tests.” *Econometric Reviews*, **28**(5), 393–421. doi:10.1080/07474930802467167.
- Cavaliere G, Taylor AMR (2009b). “Heteroskedastic Time Series with a Unit Root.” *Econometric Theory*, pp. 1228–1276. doi:10.1017/s026646660809049x.
- Chang Y (2002). “Nonlinear IV Unit Root Tests in Panels with Cross-Sectional Dependency.” *Journal of Econometrics*, **110**(2), 261–292. doi:10.1016/s0304-4076(02)00095-7.
- Chang Y, Park JY (2002). “On the Asymptotics of ADF Tests for Unit Roots.” *Econometric Reviews*, **21**, 431–447. doi:10.1081/etc-120015385.
- Chang Y, Park JY (2003). “A Sieve Bootstrap for the Test of a Unit Root.” *Journal of Time Series Analysis*, **24**(4), 379–400. doi:10.1111/1467-9892.00312.
- Choi I (2001). “Unit Root Tests for Panel Data.” *Journal of International Money and Finance*, **20**(2), 249–272. doi:10.1016/s0261-5606(00)00048-6.
- Choi I (2015). *Almost All about Unit Roots: Foundations, Developments, and Applications*. Cambridge University Press.
- Croissant Y, Millo G (2008). “Panel Data Econometrics in R: The **plm** Package.” *Journal of Statistical Software*, **27**(2), 1–43. doi:10.18637/jss.v027.i02.
- Croissant Y, Millo G (2018). *Panel Data Econometrics with R*. John Wiley & Sons. doi:10.1002/9781119504641.
- Davidson R, Flachaire E (2008). “The Wild Bootstrap, Tamed At Last.” *Journal of Econometrics*, **146**, 162–169. doi:10.1016/j.jeconom.2008.08.003.
- de Lacalle JL, Boshnakov GN (2020). **uroot**: Unit Root Tests for Seasonal Time Series. R package version 2.1-2, URL <https://CRAN.R-project.org/package=uroot>.
- Dickey DA, Fuller WA (1979). “Distribution of Estimators for Autoregressive Time Series with a Unit Root.” *Journal of the American Statistical Association*, **74**(366a), 427–431. doi:10.1080/01621459.1979.10482531.
- Dickey DA, Fuller WA (1981). “Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root.” *Econometrica*, pp. 1057–1072. doi:10.2307/1912517.
- Eddelbuettel D (2013). *Seamless R and C++ Integration with Rcpp*. Springer-Verlag, New York. doi:10.1007/978-1-4614-6868-4.
- Eddelbuettel D, Balamuta JJ (2017). “Extending R with C++: A Brief Introduction to Rcpp.” *PeerJ Preprints*, **5**, e3188v1. ISSN 2167-9843. doi:10.7287/peerj.preprints.3188v1.
- Eddelbuettel D, François R (2011). “Rcpp: Seamless R and C++ Integration.” *Journal of Statistical Software*, **40**(8), 1–18. doi:10.18637/jss.v040.i08.
- Eddelbuettel D, Sanderson C (2014). “RcppArmadillo: Accelerating R with High-Performance C++ Linear Algebra.” *Computational Statistics & Data Analysis*, **71**, 1054–1063. doi:10.1016/j.csda.2013.02.005.

- Elliott G, Rothenberg TJ, Stock JH (1996). “Efficient Tests for an Autoregressive Unit Root.” *Econometrica*, **64**(4), 813–836. doi:10.2307/2171846.
- Enders W (2008). *Applied Econometric Time Series*. 4th edition. John Wiley & Sons.
- Friedrich M, Smeekes S, Urbain JP (2020). “Autoregressive Wild Bootstrap Inference for Nonparametric Trends.” *Journal of Econometrics*, **214**(1), 81–109. doi:10.26481/umagsb.2017010.
- Hansen BE (1995). “Rethinking the Univariate Approach to Unit Root Testing: Using Covariates to Increase Power.” *Econometric Theory*, pp. 1148–1171. doi:10.1017/s0266466600009993.
- Harvey DI, Leybourne SJ, Taylor AMR (2009). “Unit Root Testing in Practice: Dealing with Uncertainty over the Trend and Initial Condition.” *Econometric Theory*, **25**(3), 587–636. doi:10.1017/s026646660809018x.
- Harvey DI, Leybourne SJ, Taylor AMR (2012). “Testing for Unit Roots in the Presence of Uncertainty over Both the Trend and Initial Condition.” *Journal of Econometrics*, **169**(2), 188–195. doi:10.1016/j.jeconom.2012.01.018.
- Holm S (1979). “A Simple Sequentially Rejective Multiple Test Procedure.” *Scandinavian Journal of Statistics*, **6**, 65–70. doi:10.2307/2532027.
- Hylleberg S, Engle RF, Granger CWJ, Yoo BS (1990). “Seasonal Integration and Cointegration.” *Journal of Econometrics*, **44**(1-2), 215–238. doi:10.1016/0304-4076(90)90080-d.
- Im KS, Pesaran MH, Shin Y (2003). “Testing for Unit Roots in Heterogeneous Panels.” *Journal of Econometrics*, **115**(1), 53–74. doi:10.1016/s0304-4076(03)00092-7.
- Kleiber C, Lupi C (2011). “Panel Unit Root Testing with R.” Vignette, version 2011-05-18, R package **punitroots**, URL <https://rdrr.io/rforge/punitroots/f/inst/doc/panelUnitRootWithR.pdf>.
- Kleiber C, Lupi C (2012). **punitroots: Tests for Unit Roots in Panels of (Economic) Time Series, with and without Cross-Sectional Dependence**. R package version 0.0-2/r42, URL <https://R-Forge.R-project.org/projects/punitroots/>.
- Kreiss JP, Paparoditis E, Politis DN (2011). “On the Range of Validity of the Autoregressive Sieve Bootstrap.” *The Annals of Statistics*, **39**, 2103–2130. doi:10.1214/11-aos900.
- Kwiatkowski D, Phillips PCB, Schmidt P, Shin Y (1992). “Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root.” *Journal of Econometrics*, **54**(1-3), 159–178. doi:10.1016/0304-4076(92)90104-y.
- Levin A, Lin CF, Chu CSJ (2002). “Unit Root Tests in Panel Data: Asymptotic and Finite-Sample Properties.” *Journal of Econometrics*, **108**(1), 1–24. doi:10.1016/s0304-4076(01)00098-7.
- Lupi C (2009). “Unit Root CADF Testing with R.” *Journal of Statistical Software*, **32**(2), 1–19. doi:10.18637/jss.v032.i02.

- Maddala GS, Wu S (1999). “A Comparative Study of Unit Root Tests with Panel Data and a New Simple Test.” *Oxford Bulletin of Economics and Statistics*, **61**(S1), 631–652. doi:10.1111/1468-0084.0610s1631.
- Mallet O (2017). *URT: Fast Unit Root Tests and OLS Regression in C++ with Wrappers for R and Python*. URL <https://github.com/olmallet81/URT>.
- Mammen E (1993). “Bootstrap and Wild Bootstrap for High Dimensional Linear Models.” *The Annals of Statistics*, **21**, 255–285. doi:10.1214/aos/1176349025.
- McCracken M, Ng S (2020). “FRED-QD: A Quarterly Database for Macroeconomic Research.” *Working Paper 26872*, National Bureau of Economic Research.
- McCracken MW, Ng S (2016). “FRED-MD: A Monthly Database for Macroeconomic Research.” *Journal of Business & Economic Statistics*, **34**(4), 574–589. doi:10.1080/07350015.2015.1086655.
- Moon HR, Perron B (2012). “Beyond Panel Unit Root Tests: Using Multiple Testing to Determine the Nonstationarity Properties of Individual Series in a Panel.” *Journal of Econometrics*, **169**(1), 29–33. doi:10.1016/j.jeconom.2012.01.008.
- Müller UK, Elliott G (2003). “Tests for Unit Roots and the Initial Condition.” *Econometrica*, **71**(4), 1269–1286. doi:10.1111/1468-0262.00447.
- Ng S, Perron P (2001). “Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power.” *Econometrica*, **69**(6), 1519–1554. doi:10.1111/1468-0262.00256.
- Ng S, Perron P (2005). “A Note on the Selection of Time Series Models.” *Oxford Bulletin of Economics and Statistics*, **67**, 115–134. doi:10.1111/j.1468-0084.2005.00113.x.
- Palm FC, Smeekes S, Urbain JP (2008). “Bootstrap Unit Root Tests: Comparison and Extensions.” *Journal of Time Series Analysis*, **29**(1), 371–401. doi:10.1111/j.1467-9892.2007.00565.x.
- Palm FC, Smeekes S, Urbain JP (2011). “Cross-Sectional Dependence Robust Block Bootstrap Panel Unit Root Tests.” *Journal of Econometrics*, **163**(1), 85–104. doi:10.1016/j.jeconom.2010.11.010.
- Pantula SG (1989). “Testing for Unit Roots in Time Series Data.” *Econometric Theory*, **5**(2), 256–271. doi:10.1017/s0266466600012421.
- Paparoditis E, Politis DN (2003). “Residual-Based Block Bootstrap for Unit Root Testing.” *Econometrica*, **71**(3), 813–855. doi:10.1111/1468-0262.00427.
- Paparoditis E, Politis DN (2005). “Bootstrapping Unit Root Tests for Autoregressive Time Series.” *Journal of the American Statistical Association*, **100**, 545–553. doi:10.1198/016214504000001998.
- Paparoditis E, Politis DN (2018). “The Asymptotic Size and Power of the Augmented Dickey-Fuller Test for a Unit Root.” *Econometric Reviews*, **37**(9), 955–973. doi:10.1080/00927872.2016.1178887.

- Perron P, Qu Z (2007). “A Simple Modification to Improve the Finite Sample Properties of Ng and Perron’s Unit Root Tests.” *Economics Letters*, **94**(1), 12–19. doi:10.1016/j.econlet.2006.06.009.
- Pesaran MH (2007). “A Simple Panel Unit Root Test in the Presence of Cross-Section Dependence.” *Journal of Applied Econometrics*, **22**(2), 265–312. doi:10.1002/jae.951.
- Pesaran MH (2012). “On the Interpretation of Panel Unit Root Tests.” *Economics Letters*, **116**(3), 545–546. doi:10.1016/j.econlet.2012.04.049.
- Pfaff B (2008). *Analysis of Integrated and Cointegrated Time Series with R*. 2nd edition. Springer-Verlag, New York.
- Phillips PCB, Perron P (1988). “Testing for a Unit Root in Time Series Regression.” *Biometrika*, **75**(2), 335–346. doi:10.1093/biomet/75.2.335.
- Psaradakis Z (2001). “Bootstrap Tests for an Autoregressive Unit Root in the Presence of Weakly Dependent Errors.” *Journal of Time Series Analysis*, **22**, 577–594. doi:10.1111/1467-9892.00242.
- R Core Team (2023). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Rho Y, Shao X (2019). “Bootstrap-Assisted Unit Root Testing with Piecewise Locally Stationary Errors.” *Econometric Theory*, **35**(1), 142–166. doi:10.1017/s0266466618000038.
- Romano JP, Shaikh AM, Wolf M (2008a). “Control of the False Discovery Rate under Dependence Using the Bootstrap and Subsampling.” *Test*, **17**(3), 417–442. doi:10.1007/s11749-008-0126-6.
- Romano JP, Shaikh AM, Wolf M (2008b). “Formalized Data Snooping Based on Generalized Error Rates.” *Econometric Theory*, **24**(2), 404–447. doi:10.1017/s0266466608080171.
- Romano JP, Wolf M (2005). “Stepwise Multiple Testing as Formalized Data Snooping.” *Econometrica*, **73**(4), 1237–1282. doi:10.1111/j.1468-0262.2005.00615.x.
- Schwert GW (1989). “Tests for Unit Roots: A Monte Carlo Investigation.” *Journal of Business and Economic Statistics*, **7**(1), 147–159. doi:10.1080/07350015.1989.10509723.
- Shao X (2010). “The Dependent Wild Bootstrap.” *Journal of the American Statistical Association*, **105**(489), 218–235. doi:10.1198/jasa.2009.tm08744.
- Shao X (2011). “A Bootstrap-Assisted Spectral Test of White Noise under Unknown Dependence.” *Journal of Econometrics*, **162**(2), 213–224. doi:10.1016/j.jeconom.2011.01.001.
- Smeekes S (2013). “Detrending Bootstrap Unit Root Tests.” *Econometric Reviews*, **32**(8), 869–891. doi:10.1080/07474938.2012.690693.
- Smeekes S (2015). “Bootstrap Sequential Tests to Determine the Order of Integration of Individual Units in a Time Series Panel.” *Journal of Time Series Analysis*, **36**(3), 398–415. doi:10.1111/jtsa.12110.

- Smeekes S, Taylor AMR (2012). “Bootstrap Union Tests for Unit Roots in the Presence of Nonstationary Volatility.” *Econometric Theory*, **28**(2), 422–456. doi:10.1017/S0266466611000387.
- Smeekes S, Urbain JP (2014a). “A Multivariate Invariance Principle for Modified Wild Bootstrap Methods with an Application to Unit Root Testing.” *GSBE Research Memorandum RM/14/008*, Maastricht University.
- Smeekes S, Urbain JP (2014b). “On the Applicability of the Sieve Bootstrap in Time Series Panels.” *Oxford Bulletin of Economics and Statistics*, **76**(1), 139–151. doi:10.1111/obes.12005.
- Smeekes S, Wijler E (2020). “Unit Roots and Cointegration.” In P Fuleky (ed.), *Macroeconomic Forecasting in the Era of Big Data*, volume 52 of *Advanced Studies in Theoretical and Applied Econometrics*, chapter 17, pp. 541–584. Springer-Verlag.
- Smeekes S, Wilms I (2023). **bootUR**: *Bootstrap Unit Root Tests*. R package version 1.0.0, URL <https://CRAN.R-project.org/package=bootUR>.
- Stroustrup B (2013). *The C++ Programming Language*. 4th edition. Addison-Wesley.
- Theußl S, Zeileis A (2009). “Collaborative Software Development Using R-Forge.” *The R Journal*, **1**(1), 9–14. doi:10.32614/rj-2009-007.
- Trapletti A, Hornik K (2023). **tseries**: *Time Series Analysis and Computational Finance*. R package version 0.10-53, URL <https://CRAN.R-project.org/package=tseries>.
- Tsung-wu H (2022). **pdR**: *Threshold Model and Unit Root Tests in Cross-Section and Time Series Data*. R package version 1.8, URL <https://CRAN.R-project.org/package=pdR>.
- Wuertz D, Setz T, Chalabi Y (2022). **fUnitRoots**: *Rmetrics – Modelling Trends and Unit Roots*. R package version 4021.80, URL <https://CRAN.R-project.org/package=fUnitRoots>.
- Zhang X, Cheng G (2014). “Bootstrapping High Dimensional Time Series.” *arXiv 1406.1037*, arXiv.org E-Print Archive. doi:10.48550/arxiv.1406.1037.
- Zhang Y, Yu H, McLeod AI (2013). “Developments in Maximum Likelihood Unit Root Test.” *Communications in Statistics – Simulation and Computation*, **42**(5), 1088–1103. doi:10.1080/03610918.2012.655828.

A. Simulation study

We perform a simulation study to investigate the small sample performance of the different unit root tests, thereby largely following the design choices in [Smeekes \(2015\)](#). We consider three designs: a (i) micro panel with $T = 25$ and $N = 200$ (ii) macro panel with $T = 100$ and $N = 50$, and a (iii) high-dimensional panel with $T = 100$ and $N = 100$. Data are generated from model (1) with $\beta = 0$ (hence no deterministic components), and the error terms drawn from a non-central t distribution with non-centrality parameter 10 and 4 degrees of freedom. For the stationary units $i \leq k_0$ where $k_0 = \lfloor q_0 N \rfloor$ in the panel, we set the AR parameter $\rho = 0.6$ in setting (i) and $\rho = 0.9$ in settings (ii) and (iii). For the non-stationary units $i > k_0$, we take $\rho = 1$. As the true proportions of stationary units, we consider $q_0 = 0, 0.2, 0.5, 0.9$.

Results are reported for the asymptotic `adf()` test, and the bootstrap `boot_ur()`, `boot_sqt()` and `boot_panel()` tests. For all tests, we apply the `order_integration()` function with `max_order = 1` to test for a unit root in the series in levels. For the bootstrap tests, we use `B = 499` bootstrap replications. For the `boot_sqt()` test, we take `steps = 0:8/8` in the micro and high-dimensional panel, `steps = 0:4/4` for the macro panel. All other arguments are kept at their defaults. Results are based on 200 Monte Carlo simulations.

We assess the performance of the unit root tests on three metrics. We report the average proportion of units (1) correctly classified (“correct”), (2) incorrectly found to be stationary (“size”), (3) correctly found to be stationary (“power”). Results are reported in Table 6.

For the micro panel, all tests produce a roughly similar, high proportion of correct classifications and a fairly low probability of classifying units incorrectly as $I(0)$ when there are few stationary units. All tests have rather low power, which is expected given the small number of time series observations, but this is most severe for the `boot_fdr()` test. As the proportion of stationary units increases, `boot_sqt()` starts to perform better than the other tests; the opposite occurs for `boot_fdr()`.

For the macro and high-dimensional panel, results for the tests `adf()`, `boot_ur()` and

Design	Unit root test	$q_0 = 0$			$q_0 = 0.2$			$q_0 = 0.5$			$q_0 = 0.9$		
		correct	size	power	correct	size	power	correct	size	power	correct	size	power
micro	<code>adf()</code>	0.96	0.04	0.26	0.82	0.04	0.26	0.61	0.04	0.26	0.33	0.04	0.26
	<code>boot_ur()</code>	0.94	0.06	0.44	0.84	0.06	0.44	0.69	0.06	0.45	0.49	0.06	0.44
	<code>boot_sqt()</code>	1.00	0.00	0.31	0.83	0.04	0.31	0.68	0.06	0.42	0.57	0.09	0.53
	<code>boot_fdr()</code>	1.00	0.00	0.00	0.80	0.00	0.00	0.50	0.00	0.00	0.11	0.00	0.01
macro	<code>adf()</code>	0.96	0.04	0.23	0.81	0.04	0.23	0.60	0.04	0.23	0.30	0.04	0.23
	<code>boot_ur()</code>	0.94	0.06	0.36	0.82	0.06	0.36	0.65	0.06	0.35	0.42	0.07	0.36
	<code>boot_sqt()</code>	0.99	0.01	0.31	0.80	0.08	0.31	0.68	0.12	0.49	0.70	0.28	0.69
	<code>boot_fdr()</code>	1.00	0.00	0.03	0.80	0.00	0.03	0.52	0.00	0.03	0.15	0.01	0.06
high-dim	<code>adf()</code>	0.96	0.04	0.23	0.81	0.04	0.23	0.60	0.04	0.24	0.30	0.04	0.23
	<code>boot_ur()</code>	0.94	0.06	0.36	0.82	0.06	0.36	0.65	0.06	0.36	0.42	0.06	0.36
	<code>boot_sqt()</code>	0.99	0.01	0.27	0.81	0.05	0.27	0.65	0.06	0.37	0.51	0.10	0.46
	<code>boot_fdr()</code>	1.00	0.00	0.01	0.80	0.00	0.01	0.51	0.00	0.02	0.13	0.00	0.03

Table 6: Simulation results for the four unit root tests across the different simulation designs (micro, macro and high-dim panel), as well as varying proportion of stationary units q_0 .

`boot_fdr()` are largely comparable to those of the micro panel. The `boot_sqt()` test selects too many stationary units for the macro panel (especially when their true proportion is large), which occurs due to the choice of the argument `steps` set to four equally spaced quantiles. Note that the `boot_fdr()` test still displays very low power in the macro panel, whereas its power reported in [Smeekes \(2015\)](#) is considerably better. Note that we consider a different value under the alternative here; we use a fixed AR parameter of 0.9 whereas [Smeekes \(2015\)](#) drew the AR parameter from a uniform distribution between zero and 0.9. In addition, the usage of the union of rejection test instead of the ADF test with intercept in [Smeekes \(2015\)](#) as well as the different choice of the maximum number of lagged differences in the ADF regression (we use the sample size dependent default in `bootUR`, 12 in this case, whereas the maximum was set to four in [Smeekes 2015](#)) may also affect the performance of the `boot_fdr()` test.

Affiliation:

Stephan Smeekes, Ines Wilms
Quantitative Economics, School of Business and Economics
Maastricht University
Tongersestraat 53
6211 LM Maastricht, The Netherlands
E-mail: s.smeekes@maastrichtuniversity.nl, i.wilms@maastrichtuniversity.nl
URL: <https://www.stephansmeekes.nl/>, <https://sites.google.com/view/iwilms>