



exuber: Recursive Right-Tailed Unit Root Testing with R

Kostas Vasilopoulos 
Lancaster University

Efthymios Pavlidis 
Lancaster University

Enrique Martínez-García 
Federal Reserve Bank of Dallas

Abstract

This paper introduces the R package **exuber** for testing and date-stamping periods of mildly explosive dynamics (exuberance) in time series. The package computes test statistics for the supremum augmented Dickey-Fuller test (SADF) of Phillips, Wu, and Yu (2011), the generalized SADF (GSADF) of Phillips, Shi, and Yu (2015a,b), and the panel GSADF proposed by Pavlidis, Yusupova, Paya, Peel, Martínez-García, Mack, and Grossman (2016); generates finite-sample critical values based on Monte Carlo and bootstrap methods; and implements the corresponding date-stamping procedures. The recursive least-squares algorithm that we introduce in our implementation of these techniques utilizes the matrix inversion lemma and in that way achieves significant speed improvements. We illustrate the speed gains in a simulation experiment, and provide illustrations of the package using artificial series and a panel on international house prices.

Keywords: mildly explosive time series, right-tailed unit root tests, R.

1. Introduction

Over the last decades, and especially after the financial crisis of 2007–09, a large interest has developed in econometric tests of exuberance (explosive behavior) in asset markets. The objective of these tests is to detect periods in which the data generating process of a macroeconomic or a financial variable is characterized by (mildly) explosive dynamics.¹ Because such periods are typically followed by large market corrections which can trigger economic

¹As noted by Blanchard (1979), among others, asset market booms followed by crashes could result from rationale bubbles.

recessions, exuberance tests are of prime importance for academics who want to shed light on the functioning of asset markets, but also for policymakers for monitoring purposes. Applications of exuberance tests include, among others, stock prices, real estate valuations, precious metals, foreign exchange rates, crypto-currencies, public debt, and bond yield spreads.

Boom-bust episodes in asset markets may take all kinds of shapes and, therefore, empirically detecting their presence is a challenging task. Early empirical studies on exuberance adopted linear, time-invariant model specifications. The majority of these studies tested for explosive behavior by applying standard right-tailed unit root tests, such as the augmented Dickey-Fuller (ADF) test, to the entire sample of available data. Although widely employed, standard unit root tests have extremely low power in detecting episodes of explosive dynamics when these are interrupted by market crashes (Evans 1991).

In recent years, several new econometric methodologies have been proposed to deal with this shortcoming. Two methodologies that have gained increasing popularity are the supremum ADF test (SADF) of Phillips *et al.* (2011) and the generalized SADF test (GSADF) by Phillips *et al.* (2015a,b). In order to deal with the effect of a collapse in a time series on the test's performance, the SADF and GSADF methodologies involve a recursively evolving algorithm that estimates ADF regressions on subsamples of data. The SADF procedure sequentially tests for explosive behavior by using a forward expanding window, while its extension, the GSADF procedure, tests for exuberance using all possible subsamples of a time series given a user-specified minimum window size. The GSADF procedure is particularly attractive because it minimizes the impact of previous boom-bust episodes on the identification and thereby is consistent with multiple changes in regime. Simulation evidence in Homm and Breitung (2012), Phillips *et al.* (2015a), and Pavlidis, Paya, and Peel (2017) suggests that the GSADF test has accurate size and higher power compared to the SADF and alternative tests for changes in persistence from unit root to mildly explosive.² An attractive feature of both the SADF and GSADF methodologies is that, due to their recursive nature, they allow date-stamping of the exact periods, if any, during which the series under examination displays explosive dynamics. Therefore, they can be used to shed light on past episodes of exuberance but also for real-time market surveillance.

For the above reasons, the SADF and GSADF methodologies have become the most commonly applied in the literature on asset market exuberance and there are currently several software implementations available. These implementations, which are based on translations of the original MATLAB code of Phillips *et al.* (2015a)³ into other programming languages, include the R packages **psymonitor** (Caspi, Phillips, and Shi 2019) and **MultipleBubbles** (Araujo, Lacerda, Phillips, and Shi 2018), the EViews add-in **Rtadf** (Caspi 2017), and the Stata module **radf** (Otero and Baum 2020).

Both the SADF and GSADF are univariate testing procedures and, thus, only allow drawing conclusions one series at a time. In a number of relevant cases, exuberance is found to be (nearly) coincident across a group of assets. Prime examples include regional and international house prices in the 2000s (see Pavlidis *et al.* 2016; Greenaway-McGrevy, Grimes, and Holmes 2019; Martínez-García and Grossman 2020) and stock prices (Narayan, Mishra, Sharma, and Liu 2013). To accommodate concurrent episodes of exuberance, Pavlidis *et al.* (2016) propose an extension of the GSADF procedure to a panel setting. The panel GSADF draws

²These include the modified tests of Bhargava (1986), Buseti and Taylor (2004), Kim (2000), and the Chow-type unit root test examined by Homm and Breitung (2012).

³The code is available online at <https://sites.google.com/site/shupingshi/home/codes>.

inference on the presence of overall exuberance by exploiting the cross-sectional dimension of a dataset through a sieve bootstrap procedure. This extension can perform substantially better than univariate tests applied to aggregated series in the presence of synchronized episodes of exuberance (Pavlidis, Martínez-García, and Grossman 2019) and, like univariate tests, provides a date-stamping strategy. Despite its advantages, a software implementation of the panel recursive unit root testing procedure does not currently exist.

In this paper, we introduce the R package **exuber** (Vasilopoulos, Pavlidis, Spavound, and Martínez-García 2022) which provides a unified approach for univariate and panel testing and date-stamping periods of exuberance. Aside from extending the framework to panel data, two features that distinguish **exuber** from existing packages are its computational speed and functionality. With regard to the first feature, the core function of the package recursively estimates ADF regressions by utilizing the Sherman-Morrison-Woodbury identity, also known as the matrix inversion lemma, to efficiently update least-squares estimates as the sub-sample size changes. To improve further computational efficiency, this function is written in C++ using the **Rcpp** extension package and employs **RcppArmadillo** for high-performance linear algebra (Eddelbuettel and François 2011; Eddelbuettel 2013; Eddelbuettel and Sanderson 2014). As we show in Section 4, this implementation leads to computational speed gains of one or more orders of magnitude compared to existing implementations. The computational gains achieved are particularly important for statistical inference. This is so because the distributions of the various test statistics of exuberance are non-standard and thus their estimation requires the use of Monte Carlo simulation or bootstrap methods. In addition to the recursive least squares algorithm, **exuber** exploits the facilities of R for parallel computing, via the packages **parallel** (R Core Team 2022) and **foreach** (Kane, Emerson, and Weston 2013; Microsoft and Weston 2022), to compute critical values in a timely manner, even when the dimension of the dataset under examination is large.

With respect to its functionality, in comparison to other implementations, **exuber** has a simple and intuitive application programming interface (API) that enables easy manipulation of the estimated objects through S3 method dispatch, it can handle different data objects whose structure is wide, it allows for single index parsing, and it produces publication-ready tables and graphs. Furthermore, the two-stage econometric methodology that requires testing for explosive dynamics prior to date-stamping is hardcoded in the design which prevents users from erroneous applications. The **exuber** package is currently employed by the Federal Reserve Bank of Dallas, the International Housing Observatory, and the United Kingdom (UK) Housing Observatory for the production of exuberance indicators for international and UK regional housing markets.⁴

To demonstrate the functionality of the package and its ease of use, we provide two examples: one with artificial time series data and another with a panel of international house prices provided by the Federal Reserve Bank of Dallas (Mack, Martínez-García, and Grossman 2011). For the former, the package also includes a set of functions that allow the simulation of the periodically-collapsing bubble processes of Blanchard (1979) and Evans (1991) as well as the one- and two-bubble processes proposed by Phillips *et al.* (2015a). For the latter exercise, the evidence from the univariate methods hints at the possibility that a common factor may have contributed to house price exuberance spreading across countries. We explore

⁴Federal Reserve Bank of Dallas' International House Price Database (<https://www.dallasfed.org/institute/houseprice#tab2>), International Housing Observatory (<https://int.housing-observatory.com/>), and UK Housing Observatory (<https://uk.housing-observatory.com/>).

this issue further and illustrate the usefulness of the panel GSADF procedure to ascertain strong empirical evidence in favor of a period of global exuberance in international housing markets preceding the 2007–09 recession.

The rest of the paper is organized as follows. Section 2 gives an outline of the SADF, GSADF, and panel GSADF econometric methodologies. Section 3 provides an overview of the main functions and methods of the package, and Section 4 deals with the computational approach to estimation, software implementation, and performance. Section 5 presents the two hands-on examples. Finally, the last section discusses future extensions of the package.

2. Econometric methodologies

In this section, we provide a brief description of the SADF, GSADF, and panel GSADF econometric methodologies, outline the associated date-stamping strategies, and discuss technical details about their implementation.

2.1. Univariate right-tailed tests

At the heart of all the methodologies considered lies the following ADF regression equation:

$$\Delta y_t = a_{r_1, r_2} + \gamma_{r_1, r_2} y_{t-1} + \sum_{j=1}^k \psi_{r_1, r_2}^j \Delta y_{t-j} + \epsilon_t, \quad (1)$$

where y_t denotes a generic time series, Δy_{t-j} with $j = 1, \dots, k$ are lagged first differences of the series included to accommodate serial correlation, $\epsilon_t \sim \mathcal{N}(0, \sigma_{r_1, r_2}^2)$ are the Gaussian residuals, and a_{r_1, r_2} , γ_{r_1, r_2} and ψ_{r_1, r_2}^j with $j = 1, \dots, k$ are regression coefficients. The subscripts r_1 and r_2 denote fractions of the total sample size T that specify the starting and ending points of a subsample period.

The null hypothesis of interest is that of a unit root, $H_0 : \gamma_{r_1, r_2} = 0$, against the alternative of explosive behavior in y_t , $H_1 : \gamma_{r_1, r_2} > 0$. The ADF test statistic corresponding to this null is given by,

$$\text{ADF}_{r_1}^{r_2} = \hat{\gamma}_{r_1, r_2} / \text{SE}(\hat{\gamma}_{r_1, r_2}). \quad (2)$$

The standard ADF test

For the ADF test, the statistic in Equation 2 is obtained by estimating the regression given in Equation 1 on the full sample of observations, i.e., by setting $r_1 = 0$ and $r_2 = 1$. Under the null of a unit root, the limit distribution of ADF_0^1 is given by,

$$\frac{\int_0^1 W dW}{\left(\int_0^1 W^2\right)^{1/2}},$$

where W is a Wiener process. Testing for exuberance entails comparing the ADF_0^1 statistic with the right-tailed critical value from its limit distribution. In this setting with r_1 and r_2 fixed at 0 and 1, respectively, the alternative hypothesis is that of exuberance over the entire sample. When the ADF_0^1 test statistic exceeds the corresponding critical value, the unit root hypothesis is rejected in favor of the alternative hypothesis. Because the standard ADF test is not consistent with changes in regime, it exhibits extremely low power in the presence of

boom-bust episodes. In fact, nonlinear dynamics, such as those displayed by periodically-collapsing speculative bubbles frequently lead to finding spurious stationarity even when the process under examination is inherently explosive (see, e.g., [Evans 1991](#)).

The supremum ADF (SADF) test

[Phillips et al. \(2011\)](#) proposed a methodology that is consistent with a single boom-bust episode. Their methodology involves recursively estimating Equation 1 using a forward expanding sample. In this setting, the beginning of the subsample is held constant at $r_1 = 0$, while the end of the subsample, r_2 , increases from r_0 (the minimum window size) to one (the entire sample period).⁵ Recursive estimation of Equation 1 yields a sequence of $ADF_0^{r_2}$ statistics. The supremum of this sequence, called the SADF statistic, is defined by,

$$SADF(r_0) = \sup_{r_2 \in [r_0, 1]} ADF_0^{r_2},$$

and has a limit distribution given by,

$$\sup_{r_2 \in [r_0, 1]} \frac{\int_0^{r_2} W dW}{\left(\int_0^{r_2} W^2\right)^{1/2}}.$$

Similarly to the standard ADF test, rejection of the null of a unit root requires that the SADF statistic exceeds the right-tailed critical value from its limit distribution. However, contrary to the standard ADF test which examines the presence of explosive dynamics during the entire period, the alternative hypothesis of the SADF test is that of explosive dynamics in some part(s) of the sample.

The generalized SADF (GSADF) test

More recently, [Phillips et al. \(2015a,b\)](#) proposed an extension of the SADF, the GSADF, which has the same alternative hypothesis as the SADF but which covers a larger number of subsamples. Specifically, given a minimum window size r_0 , the methodology involves estimating the regression given in Equation 1 for all possible subsamples by allowing both the ending point, r_2 , and the starting point, r_1 , to change. This extra flexibility on the estimation window results in substantial power gains and makes the GSADF test better suited to detect the presence of multiple changes in regime.

Formally, the GSADF statistic is defined as,

$$GSADF(r_0) = \sup_{r_2 \in [r_0, 1], r_1 \in [0, r_2 - r_0]} ADF_{r_1}^{r_2},$$

and its limit distribution under the null is,

$$\sup_{r_2 \in [r_0, 1], r_1 \in [0, r_2 - r_0]} \left\{ \frac{\frac{1}{2} r_w [W(r_2)^2 - W(r_1)^2 - r_w] - \int_{r_1}^{r_2} W(r) dr [W(r_2) - W(r_1)]}{r_w^{1/2} \left\{ r_w \int_{r_1}^{r_2} W(r)^2 dr - \left[\int_{r_1}^{r_2} W(r) dr \right]^2 \right\}^{1/2}} \right\},$$

where $r_w = r_2 - r_1$ is the size of the expanding window. Again, rejection of the unit root hypothesis in favor of the alternative hypothesis of explosive behavior in some part(s) of the

⁵Given r_0 and r_1 , the expanding window size of the regression is given by $r_w = r_2 - r_1$.

sample requires that the GSADF statistic exceeds the right-tailed critical value from its limit distribution.

Date-stamping strategies

If the null of a unit root in y_t is rejected, then the SADF and GSADF methodologies can provide a chronology of episodes of exuberance. For the SADF methodology, estimates of the origin and conclusion of exuberance can be obtained by comparing the time series of the recursive backward ADF (BADF) test statistics against the right-tailed critical value of the distribution of the standard ADF statistic. Letting r_e and r_f correspond to the origination and termination dates, respectively, these estimates can be constructed from,

$$\hat{r}_e = \inf_{r_2 \in [r_0, 1]} \{r_2 : \text{ADF}_0^{r_2} > cu_{r_2}^\alpha\} \text{ and } \hat{r}_f = \inf_{r_2 \in [\hat{r}_e, 1]} \{r_2 : \text{ADF}_0^{r_2} < cu_{r_2}^\alpha\},$$

where $cu_{r_2}^\alpha$ is the $(1-\alpha)100\%$ critical value of the ADF statistic and α is the chosen significance level. A similar strategy is proposed by Phillips *et al.* (2015a,b) as part of the GSADF methodology. The strategy in this case is based on the sequence of backward SADF (BSADF) statistics,

$$\text{BSADF}_{r_2}(r_0) = \sup_{r_1 \in [0, r_2 - r_0]} \text{SADF}_{r_1}^{r_2},$$

and estimates the origination and termination dates according to,

$$\hat{r}_e = \inf_{r_2 \in [r_0, 1]} \{r_2 : \text{BSADF}_{r_2}(r_0) > scu_{r_2}^\alpha\} \text{ and } \hat{r}_f = \inf_{r_2 \in [\hat{r}_e, 1]} \{r_2 : \text{BSADF}_{r_2}(r_0) < scu_{r_2}^\alpha\},$$

where $scu_{r_2}^\alpha$ is the $(1-\alpha)100\%$ critical value of the SADF for the $[r_2T]$ observations. The consistency of the SADF dating strategy in the presence of a single period of explosive dynamics in y_t is established in Phillips *et al.* (2011), and the consistency of the GSADF strategy with one or two explosive periods is established in Phillips *et al.* (2015a,b).

Technical details

The computation of the recursive unit root test statistics necessitates the selection of the minimum window size, r_0 , and the autoregressive lag length, k . Regarding the minimum window size, this has to be large enough to allow initial estimation, but it should not be too large to avoid missing short episodes of exuberance. Phillips *et al.* (2015a,b) recommend setting the minimum window size according to the rule of thumb: $r_0 = 0.01 + 1.8/\sqrt{T}$. With respect to the selection of k , simulation evidence indicates the proposed right-tailed unit root methodologies work well when the number of lags is fixed at a small value, i.e., 0 or 1. On the contrary, lag selection based on information criteria can result in severe size distortions. The implementation of the right-tailed unit root tests also requires the limit distributions of the SADF, GSADF and BSADF test statistics. These distributions are non-standard and depend on the minimum window size. Consequently, critical values have to be obtained through either Monte Carlo simulations or bootstrapping.

2.2. The panel GSADF procedure

The SADF and GSADF tests can only be applied to individual time series and, hence, do not exploit the panel nature of many macroeconomic and financial datasets. Based on the work

of Im, Pesaran, and Shin (2003), Pavlidis *et al.* (2016) propose an extension of the GSADF procedure to heterogeneous panels.

Consider the panel version of the regression given in Equation 1,

$$\Delta y_{i,t} = a_{i,r_1,r_2} + \gamma_{i,r_1,r_2} y_{i,t-1} + \sum_{j=1}^k \psi_{i,r_1,r_2}^j \Delta y_{i,t-j} + \epsilon_{i,t},$$

where $i = 1, \dots, N$ denotes the panel index and the remaining variables are defined as in the previous sub-section. The panel GSADF test examines the null hypothesis of a unit root, $H_0 : \gamma_{i,r_1,r_2} = 0$, in all N series against the alternative of explosive behavior in a subset of series, $H_1 : \gamma_{i,r_1,r_2} > 0$ for some i . This alternative allows for γ_{i,r_1,r_2} to differ across panels and, in that sense, is more general (and less restrictive) than approaches that impose a homogeneous alternative hypothesis.

The testing procedure involves constructing a measure of *overall* exuberance by averaging the individual BSADF statistics at each time period,

$$\text{panel BSADF}_{r_2}(r_0) = \frac{1}{N} \sum_{i=1}^N \text{BSADF}_{i,r_2}(r_0).$$

Given Equation 2.2, the definition of the panel GSADF statistic follows naturally. It is simply the supremum of the panel BSADF,

$$\text{panel GSADF}(r_0) = \sup_{r_2 \in [r_0, 1]} \text{panel BSADF}_{r_2}(r_0).$$

Like for the univariate methods, testing for and dating episodes of overall exuberance involves comparing the panel GSADF and panel BSADF statistics with the right-tailed critical values from the corresponding limit distributions. However, as shown by Maddala and Wu (1999) and Chang (2004), the limit distribution of panel unit root tests based on mean unit root statistics is not invariant to the cross-sectional dependence of the error terms, ϵ_i . To this end, the authors adopt a sieve bootstrap that is designed specifically to allow for cross-sectional error dependence. Details of the sieve bootstrap algorithm can be found in the Appendix of Pavlidis *et al.* (2016).

3. Installation and package overview

The package **exuber** (Vasilopoulos *et al.* 2022) is available on the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=exuber>, and can be installed and accessed from the R console. The package has a number of dependencies on other R packages which are automatically loaded with **exuber**. The main dependencies include **Rcpp** (Eddelbuettel 2013), **parallel** (R Core Team 2022), **foreach** (Microsoft and Weston 2022), **dplyr** (Wickham, François, Henry, and Müller 2022), and **purrr** (Henry and Wickham 2020).

```
R> install.packages("exuber")
R> library("exuber")
```

The package **exuber** centers around the function `radf()`. This function takes as input three arguments, (`data`, `minw`, `lag`) and returns a list containing ADF, SADF, GSADF, BSADF and, whenever appropriate, panel GSADF and panel BSADF test statistics. With regard to the first argument, **exuber** is capable of working with regular `'data.frame'`, `'ts'`, numeric

‘vector’ or ‘matrix’ classes without the need for any further transformation. Estimation always includes a time dimension and, in the case of multiple time series, a cross-sectional dimension. In line with the convention in time-series econometrics, the required format is wide, where each row corresponds to a unique time period and each column represents a different time series. The `radf()` function is vectorized, i.e., it can handle multiple series at once, to improve efficiency. This property also enables the computation of panel statistics internally as a by-product of the univariate estimations with minimal additional cost incurred. When the data input is a ‘data.frame’ with an explicit `class-Date` column, **exuber** uses a similar principle as **tsibble** (Wang, Cook, and Hyndman 2020) for index search and parsing. In all other cases, an `integer` index is automatically created. The index is used by **exuber**’s methods for date-stamping and plotting.

With regards to the arguments `minw` and `lag`, these are non-negative integers that specify the minimum window size and the lag length k used in estimation. The default values for these arguments are set equal to the value suggested by the rule of Phillips *et al.* (2015a) $\lfloor T(0.01 + 1.8/\sqrt{T}) \rfloor$ and zero, respectively. Users can obtain the default minimum window size by using the helper function `psy_minw()`. Given `(data, minw, lag)`, the function `radf()` returns an object of class ‘`radf_obj`’.

Drawing inference on the estimated unit root test statistics requires knowledge of their underlying distributions. The functions `radf_mc_cv()`, `radf_wb_cv()`, and `radf_sb_cv()` generate 90%, 95%, and 99% critical values for different sample and minimum window sizes. The three functions for generating critical values return an object of class ‘`radf_cv`’. The function `radf_mc_cv()` generates Monte Carlo finite-sample critical values for univariate unit root tests approximating the Brownian motion processes in the limiting functionals specified in Section 2 by using independent $\mathcal{N}(0, 1)$ random variates. While, functions `radf_wb_cv()` and `radf_sb_cv()` implement the univariate wild bootstrap and the panel sieve bootstrap methodologies outlined in Harvey, Leybourne, Sollis, and Taylor (2016) and Pavlidis *et al.* (2016).⁶ The default values for the number of Monte Carlo repetitions (`nrep`) and bootstrap repetitions (`nboot`) in **exuber** are set to 1000 and 500, respectively. Users can also obtain the entire distribution of the various test statistics by using the corresponding functions `radf_mc_distr()`, `radf_wb_distr()` and `radf_sb_distr()`, which return an object of class ‘`radf_distr`’.⁷

Because the majority of users employ Monte Carlo critical values and adopt the default minimum window size, **exuber** comes with a stored dataset of critical values for sample sizes of up to 600 observations (`radf_crit`). These critical values can be replicated with the `radf_mc_cv()` function by setting the number of repetitions `nrep` to 2000 and the `seed` to 123. For applications which involve series of length between 600 and 2000 observations, critical values can be obtained from the **exuberdata** package (Vasilopoulos 2020a). Due to CRAN restrictions on package size, **exuberdata** is distributed through the ‘Drat’ R Archive Template

⁶The arguments of the `radf_mc_cv()` function are `(n, minw, nrep, seed)`. The argument `n` inputs the number of observations in the sample, `minw` is the minimum window size as defined in the `radf()` function, while `nrep` and `seed` set the number of Monte Carlo simulations and specify if and how the random number generator should be initialized, respectively. The arguments for `radf_wb_cv()` and `radf_sb_cv()` are `(data, minw, lag, nboot, seed)` with the same interpretation as before except for the added argument `nboot` which sets the number of bootstrap repetitions.

⁷`radf` stands for recursive ADF. The naming of these functions attaches to that one of three terms, `mc` for the Monte Carlo procedure, `wb` for the univariate wild bootstrap and `sb` for the panel sieve bootstrap, plus `cv` at the end to indicate that they are critical values or `distr` for the entire distribution.

(Eddelbuettel 2022), which works seamlessly with **exuber**, does not depend on **Rtools**, and makes it easy for package installation and upgrades.⁸ The stored dataset of critical values and **exuberdata** are convenience features, unique of the package **exuber**, added to facilitate its use in practice.⁹

Several S3 methods are available to process the output of `radf()` together with the output of one of the functions for generating critical values. The use of S3 methods reflects our preference for software scalability, as we intend to include new econometric methodologies in future versions of the package. All methods take as inputs two objects, `obj` and `cv`, where `obj` is of class ‘`radf_obj`’ and `cv` is the output of the function `radf_mc_cv()`, `radf_wb_cv()`, or `radf_sb_cv()`, i.e., an object of class ‘`radf_cv`’.

The method `summary()` returns a table of the estimated ADF, SADF, GSADF test statistics together with the corresponding critical values for each individual series in `data`. When the cross-sectional dimension of the dataset is large, the output of `summary()` can become lengthy and its interpretation tedious. In such cases, the user can work directly with objects `radf_obj` and `radf_cv` (as we will demonstrate in Section 5) or alternatively use the method `diagnostics()`. For each series in a panel, `diagnostics()` compares the estimated SADF or GSADF test statistic to the corresponding set of critical values and reports whether the null of a unit root is rejected and, if so, at which level of significance. When running `diagnostics()` the user has the ability to select, via the argument `option`, whether to report results for the univariate SADF or GSADF test, with the latter being the default.¹⁰

Date-stamping is implemented via the method `datestamp()`.¹¹ This method produces an object of class ‘`ds_radf`’ which contains the estimates of the origination, termination, and duration of episodes of exuberance. To prevent erroneous applications of the testing procedure, `datestamp()` only reports results for series for which the first-stage test statistic is significant and, when all series are found to be non-explosive, it prints a message to notify the user. In some applications, researchers may want to exclude very short episodes of exuberance by setting a minimum duration period. The package **exuber** does not set a minimum duration for periods of exuberance by default. However, minimum duration requirements can be introduced as a refinement at the discretion of the user by altering the value of the argument `min_duration`. Phillips *et al.* (2015a,b) recommend two rules for setting the minimum duration: $\delta \log(T)$ and $\delta \log(T)/T$, where δ is a frequency-dependent scaling parameter. These rules can be computed by using the helper function `psy_ds()`. Finally, similar to `diagnostics()`, `datestamp` gives the option of displaying results for either the SADF or the GSADF testing procedure.

To produce publication-ready graphs, **exuber** uses the `ggplot2::autoplot` methods for plotting objects of classes ‘`radf_obj`’, ‘`ds_radf`’, and ‘`radf_distr`’. An attractive feature of this implementation is that it allows users to exploit the common features from the **ggplot2** package (Wickham 2016) to manipulate plots. `augment` and `augment_join` as well as `tidy` and `tidy_join` methods are also available to manipulate the data into a convenient format for custom plotting.

⁸The **exuberdata** package is hosted at <https://github.com/kvasilopoulos/exuberdata> and can be easily installed using the wrapper function in **exuber**, `install_exuberdata()`.

⁹The datasets of critical values in **exuber** are also used by the Stata module **radf**.

¹⁰The method `summary()` has inputs (`obj`, `cv`) while the method `diagnostics()` has inputs (`obj`, `cv`, `option`) all of which have been previously introduced.

¹¹The method `datestamp()` takes as inputs four arguments (`data`, `cv`, `min_duration`, `option`).

For class ‘`radf_obj`’ objects, `autoplot()` produces a faceted plot displaying BADF or BSADF statistics together with the corresponding sequence of critical values. In generating these graphs, users can set the `autoplot()` argument `shade_opt` to the function `shade()` which will demarcate the periods of exuberance identified by `datestamp()`. For class ‘`ds_radf`’ objects, `autoplot()` produces a “chronology” plot which displays a set of segmented horizontal lines, with each line representing the identified periods of exuberance for an individual series. If the data input in `radf()` includes a `Date` index then this is automatically used in the creation of x -axes.¹² We demonstrate the above plotting options in Section 5.

Finally, **exuber** provides functions for simulating data from popular bubble data generating processes. These include the periodically-collapsing bubble processes of Blanchard (1979) and Evans (1991) and the one- and two-bubble processes proposed by Phillips *et al.* (2015a). Those are implemented using the functions `sim_blan()`, `sim_evans()`, `sim_psy1()`, and `sim_psy2()`, respectively. The package **exuber** also includes `sim_data`, an artificial dataset containing series simulated from those same data generating processes, and another auxiliary function `sim_div()` to simulate log dividend series from a random walk with drift (West 1988).

4. Software implementation

The implementation of the univariate and panel right-tailed unit root testing procedures outlined in Section 2 requires the application of iterative routines which repeatedly fit the ADF regression given in Equation 1 on subsamples of data by solving a sequence of otherwise standard least-squares problems. Unlike other software implementations, **exuber** exploits the matrix inversion lemma to update least-squares estimates, thus, significantly reducing the costs that arise when dealing with long-time series or with large panels. In this section, we describe both the computational approach and the performance gains achieved against the MATLAB code of Phillips *et al.* (2015a), the EViews add-in **Rtadf** (Caspi 2017), the R packages **psymonitor** (Caspi *et al.* 2019) and **MultipleBubbles** (Araujo *et al.* 2018), and the Stata module **radf** of Otero and Baum (2020).

4.1. Computational approach

By replacing the sample fractions r_1 and r_2 with their discrete time analogues, t_1 and t_2 , the ADF regression given in Equation 1 can be rewritten in matrix form as,

$$Y_{t_1:t_2} = X_{t_1:t_2}\beta_{t_1:t_2} + \mathcal{E}_{t_1:t_2},$$

where

$$Y_{t_1:t_2} = \begin{bmatrix} \Delta y_{t_1} \\ \Delta y_{t_1+1} \\ \vdots \\ \Delta y_{t_2} \end{bmatrix}, X_{t_1:t_2} = \begin{bmatrix} x_{t_1} \\ x_{t_1+1} \\ \vdots \\ x_{t_2} \end{bmatrix} = \begin{bmatrix} 1 & y_{t_1-1} & \Delta y_{t_1-1} & \cdots & \Delta y_{t_1-k} \\ 1 & y_{t_1} & \Delta y_{t_1} & \cdots & \Delta y_{t_1-k+1} \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ 1 & y_{t_2-1} & \Delta y_{t_2-1} & \cdots & \Delta y_{t_2-k} \end{bmatrix},$$

¹²`index()` can be used to retrieve or replace the index of an object while `series_name()` allows the user to retrieve or replace the series names.

$$\mathcal{E}_{t_1:t_2} = \begin{bmatrix} \epsilon_{t_1} \\ \epsilon_{t_1+1} \\ \vdots \\ \epsilon_{t_2} \end{bmatrix}, \text{ and } \beta_{t_1:t_2} = \begin{bmatrix} \alpha_{t_1:t_2} \\ \gamma_{t_1:t_2} \\ \psi_{t_1:t_2}^1 \\ \vdots \\ \psi_{t_1:t_2}^k \end{bmatrix}.$$

For a given subsample, $t_1 : t_2$, the solution to the least-squares problem,

$$\operatorname{argmin}_{\beta_{t_1:t_2}} \|Y_{t_1:t_2} - X_{t_1:t_2}\beta_{t_1:t_2}\|^2,$$

is given by

$$\beta_{t_1:t_2} = (X_{t_1:t_2}^\top X_{t_1:t_2})^{-1} X_{t_1:t_2}^\top Y_{t_1:t_2}. \quad (3)$$

The main computational cost in estimating Equation 3 and the coefficient standard errors arises from the computation and inversion of the square matrix $D_{t_1:t_2} = (X_{t_1:t_2}^\top X_{t_1:t_2})$. Because the construction of $D_{t_1:t_2}$ takes $O(nv^2)$ (where $v = k + 2$ is the number of regression parameters and $n = t_2 - t_1 + 1$ is the subsample size) and the inversion of $D_{t_1:t_2}$ typically requires $O(v^3)$, the overall algorithmic complexity per iteration is $O(v^3 + nv^2)$. This cost may be acceptable in settings where the number of series and iterations is relatively small, but it makes the application of recursive unit root tests to long financial series and large panels cumbersome, and it becomes prohibitive for extensive power and size simulation experiments.

To get a sense of the magnitude of this problem, consider a single series of $T = 1000$ observations and a minimum window size given by the rule of thumb in Section 2, i.e., $m = 74$. In this setting, the estimation of the SADF and GSADF tests requires $T - m + 1 = 935$ and $(T - m + 1)(T - m + 2)/2 = 437580$ recursions, respectively, and the generation of critical values involves the computation of matrix $D_{t_1:t_2}$ and its inverse hundreds of millions or billions of times, depending on the number of Monte Carlo or bootstrap repetitions.

To reduce algorithmic complexity, the workhorse function of **exuber**, `radf()`, employs a recursive least-squares method which incrementally updates $D_{t_1:t_2}$ from its previous value. Following the standard notation in the statistics literature, let $G_{t_1:t_2} = D_{t_1:t_2}^{-1}$ denote the *gain* matrix. Using the matrix inversion lemma, the gain matrix can be recursively computed from its previous value according to,

$$G_{t_1:t_2} = G_{t_1:t_2-1} - (1 + x_{t_2} G_{t_1:t_2-1} x_{t_2}^\top)^{-1} (G_{t_1:t_2-1} x_{t_2}^\top) (x_{t_2} G_{t_1:t_2-1}),$$

and the coefficient vector can be incrementally updated by,

$$\beta_{t_1:t_2} = \beta_{t_1:t_2-1} - G_{t_1:t_2} x_{t_2}^\top (x_{t_2} \beta_{t_1:t_2-1} - \Delta y_{t_2}).$$

The attractiveness of the above algorithm lies in the fact that it involves solely matrix multiplications, which reduces the computational cost per recursion from $O(v^3 + nv^2)$ to $O(v^2)$, which only increases linearly with the number of recursions.

4.2. Computational performance

For interpreted languages such as R, iterative function calls typically come with a heavy computational overhead due to high-level language properties like dynamic typing, bounds checking, and garbage collection (see, e.g., [Sridharan and Patel 2014](#)). To avoid substantial speed penalties and obtain high-level computational efficiency, `radf()` is written in C++ using

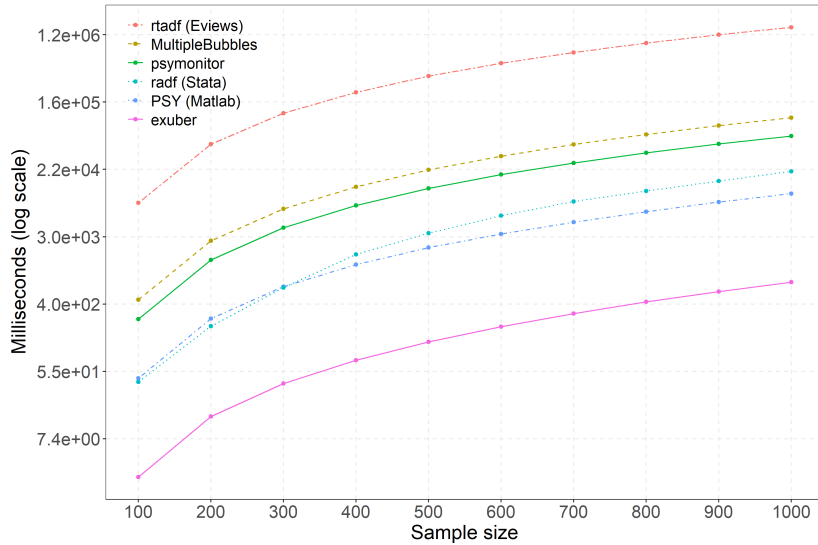


Figure 1: Package performance comparison.

the **Rcpp** extension package in conjunction with **RcppArmadillo** for high-performance linear algebra. As we show next, this implementation leads to substantial gains in computational efficiency. Figure 1 displays the run-time performance of **exuber** for samples of size $T = \{100, \dots, 1000\}$, lag length $k = 1$, and a minimum window size of 30 observations. It also shows the performance of the original MATLAB code **PSY** of Phillips *et al.* (2015a), the EViews add-in **Rtadf** (Caspi 2017), the two other R packages **psymonitor** (Caspi *et al.* 2019) and **MultipleBubbles**¹³ (Araujo *et al.* 2018), and the Stata module **radf** (Otero and Baum 2020). As previously mentioned, unlike **exuber**, these packages use a batch least-squares algorithm. Median times are generated on an Intel (R) Core(TM) i7-8700 CPU@3.20GHz with 64.0 GB RAM using 100 simulations. As is evident from the figure, **exuber** performs remarkably well both in absolute and relative terms, with the difference in computation time across packages becoming particularly apparent as the sample size increases. For a sample size of 1000 observations, the computation time for **exuber** is just 0.78 seconds. The MATLAB code **PSY** of Phillips *et al.* (2015a) ranks second, with a median time of 10.7 seconds (14 times slower), and the Stata **radf** module ranks third, with a median time of 20.8 seconds (27 times slower). The median time of the two R packages, **psymonitor** and **MultipleBubbles**, are 59.2 seconds (76 times slower) and 102.4 seconds (132 times slower), respectively. Finally, the EViews add-in **Rtadf** is by far the slowest, as it requires 1496.5 seconds (1919 times slower).

The computational speed of **exuber** on the one hand and the recursive nature of our estimation algorithm on the other render parallel estimation for a single series nontrivial and unnecessary. Parallel computing is particularly beneficial, however, for statistical inference. **exuber** uses the package **parallel** and the looping construct of **foreach** to acquire the non-standard finite-sample distributions of the various test statistics via Monte Carlo or bootstrap simulations. These distributions are generated using the functions `radf_mc_cv()`, `radf_wb_cv()`, and `radf_sb_cv()`.

¹³The function `gsadf_sadf()` of **MultipleBubbles** does not allow for a user-specified minimum window size. To provide meaningful comparisons across packages, in our simulation experiments we have used a slightly modified version of the function. The modification does not affect run-time performance.

5. Examples with artificial and real data

To demonstrate the functionality of the package, this section provides two applications. The first application involves an artificial series, while the second explores the presence of explosive dynamics in a panel of international house price data.

5.1. Simulated bubble series

For our first application, we simulate data from the one-bubble data generating process of Phillips *et al.* (2015a) using the function `sim_psy1()`. This data generating process switches from a martingale to mildly explosive and then collapses back to a martingale at predetermined sample points. We consider a series of 100 observations, and set the start and end dates of the bubble period at observations 50 and 70. The time evolution of the series can be visualized in Figure 2.

```
R> sim <- sim_psy1(100, te = 50, tf = 70, seed = 145)
R> autoplot(sim)
```

To illustrate the ease of use of **exuber** by non-expert users, we adopt the default values in the estimation of the unit root test statistics.

```
R> radf_sim <- radf(sim)
```

This produces an object of class ‘`radf_obj`’ that contains the entire set of test statistics for the univariate testing procedures (i.e., the ADF, SADF, GSADF, BADF, and BSADF). Following the econometric methodology of Section 2, the first stage of our analysis involves comparing the ADF, SADF, and GSADF statistics contained in the `radf_obj` object to their corresponding right-tailed critical values. This can be easily done using the method `summary()`. Because the estimation is based on the default window size and the sample size




Figure 2: Simulation of a single-bubble process.

is less than 600 observations, `summary()` does not require as input a `radf_cv` object of critical values. Instead, it automatically uses the Monte Carlo critical values stored in the package.

```
R> summary(radf_sim)
```

```
Using `radf_crit` for `cv`.
```

```
-- Summary (minw = 19, lag = 0) ----- Monte Carlo (nrep = 2000) --
```

```
series1 :
# A tibble: 3 x 5
  name  tstat  `90`  `95`  `99`
<fct> <dbl> <dbl> <dbl> <dbl>
1 adf   -2.38 -0.413 -0.0812 0.652
2 sadf  10.6  0.988  1.29  1.92
3 gsadf 11.6  1.71  1.97  2.57
```

A comparison of the SADF and GSADF test statistics to their critical values suggests that the null hypothesis of a unit root against the alternative of explosive behavior is rejected at all conventional levels of significance. On the contrary, the standard ADF test fails to reject the null, which is in line with the low power of the ADF test to detect bubbles that burst in-sample.

Having established the statistical significance of the SADF and GSADF test statistics at the 5% level, we can turn to the identification of the episode(s) of exuberance. To visualize these episodes, we simply pass the estimated `radf_obj` object to `autoplot()`. This returns a plot of the estimated BSADF statistics together with the sequence of 95% critical values. Shaded areas in the plot indicate periods during which the BSADF statistic exceeds its critical value and, thus, we see that there is evidence of exuberance. A similar plot can be obtained for the SADF testing procedure by altering the value of the `option` argument.

```
R> library("ggplot2")
```

```
R> autoplot_sim_gsadf <- autoplot(radf_sim) + labs(title = NULL)
```

```
Using `radf_crit` for `cv`.
```

```
R> autoplot_sim_sadf <- autoplot(radf_sim, option = "sadf") +
+   labs(title = NULL)
```

```
Using `radf_crit` for `cv`.
```

```
R> library("patchwork")
```

```
R> autoplot_sim_sadf + autoplot_sim_gsadf
```

A visual inspection of the two graphs in Figure 3 reveals that once the bubble starts, the BADF and BSADF test statistics gradually increase and eventually exceed their critical values, rejecting the null hypothesis of no explosiveness. When the bubble bursts, both statistics fall below the critical bound almost instantaneously and remain there until the end of the sample. To obtain the exact dates of the origination and collapse of the bubble and the bubble duration, we use the `datestamp()` method.




Figure 3: Datestamping with the SADF (left) and GSADF (right) tests.

```
R> datestamp(radf_sim)
```

```
Using `radf_crit` for `cv`.
```

```
-- Datestamp (min_duration = 0) ----- Monte Carlo --
```

```
series1 :
  Start End Duration
1    54  71      17
```

```
R> datestamp(radf_sim, option = "sadf")
```

```
Using `radf_crit` for `cv`.
```

```
-- Datestamp (min_duration = 0) ----- Monte Carlo --
```

```
series1 :
  Start End Duration
1    54  71      17
```

The results indicate that the bubble had started by observation 54 under the GSADF and the SADF procedures and continued until observation 70. In sum, both tests perform well in detecting and dating the period of exuberance.

5.2. International house prices

In our second application, we employ data on real house prices provided by the International House Price Database of the Federal Reserve Bank of Dallas (Mack *et al.* 2011). The dataset

has a large cross-sectional dimension (23 countries at the time of our implementation) and starts in the first quarter of 1975, thus allowing an international perspective on the evolution of housing markets over the last four decades. We use the **ihpdr** package (Vasilopoulos 2020b) to download the 2015:Q1 data release directly from the R console and, then, use **dplyr** (Wickham *et al.* 2022) to select the columns corresponding to real house prices (**rhpi**), excluding the Aggregate series. Finally, we use **tidyr** (Wickham 2014) to pivot the data into wide format.¹⁴

```
R> library("dplyr")
R> library("tidyr")
R> library("ihpdr")
R> hprices <- ihpdr::ihpd_get(version = "1501") %>%
+   filter(country != "Aggregate") %>%
+   select(Date, country, rhpi) %>%
+   pivot_wider(names_from = country, values_from = rhpi)
R> hprices
```

```
# A tibble: 161 x 24
  Date      Australia Belgium Canada Switzerland Germany Denmark Spain
  <date>    <dbl>   <dbl> <dbl>      <dbl>    <dbl>   <dbl> <dbl>
1 1975-01-01  39.2    44.3   59.3        93.8     111.    57.4  86.2
2 1975-04-01  38.5    45.4   59.0        91.8     110.    57.6  93.0
3 1975-07-01  38.6    46.5   59.7        90.4     111.    59.1  91.5
4 1975-10-01  37.8    48.1   59.4        88.9     111.    58.1  94.7
5 1976-01-01  38.0    50.6   59.1        86.8     112.    58.4  92.9
6 1976-04-01  38.2    52.0   59.7        85.9     112.    57.3  99.6
7 1976-07-01  38.4    53.4   58.7        85.0     113.    57.3  98.5
8 1976-10-01  37.9    54.8   57.5        85.2     115.    58.6  93.7
9 1977-01-01  37.9    55.9   55.8        84.5     116.    57.5  89.6
10 1977-04-01  37.8    57.5   55.2        85.1     117.    59.3  81.2
# ... with 151 more rows, and 16 more variables: Finland <dbl>,
#   France <dbl>, UK <dbl>, Ireland <dbl>, Italy <dbl>, Japan <dbl>, `S.
#   Korea` <dbl>, Luxembourg <dbl>, Netherlands <dbl>, Norway <dbl>, `New
#   Zealand` <dbl>, Sweden <dbl>, US <dbl>, `S. Africa` <dbl>,
#   Croatia <dbl>, Israel <dbl>
```

Figure 4 displays the time evolution of the cross-sectional mean of real house prices together with the lower and upper quartiles. As can be seen from the figure, mean real house prices troughed in the mid-1990s and peaked around 2006. Furthermore, by looking at the upper and lower quartiles, we observe that the run-up in real house prices during this period was not driven by the behavior of a few housing markets, but it was a widespread phenomenon. This last observation has led to a consensus among economists that this boom-bust episode in housing played a central role in the 2007–09 financial crisis around the world, and has naturally generated a vast interest in modelling the dynamics of house prices.

To further demonstrate the functionality of **exuber**, in our econometric analysis of real house

¹⁴The house price indexes are expressed in real terms using the personal consumption expenditure deflator of the corresponding country and have the same base year, 2005.

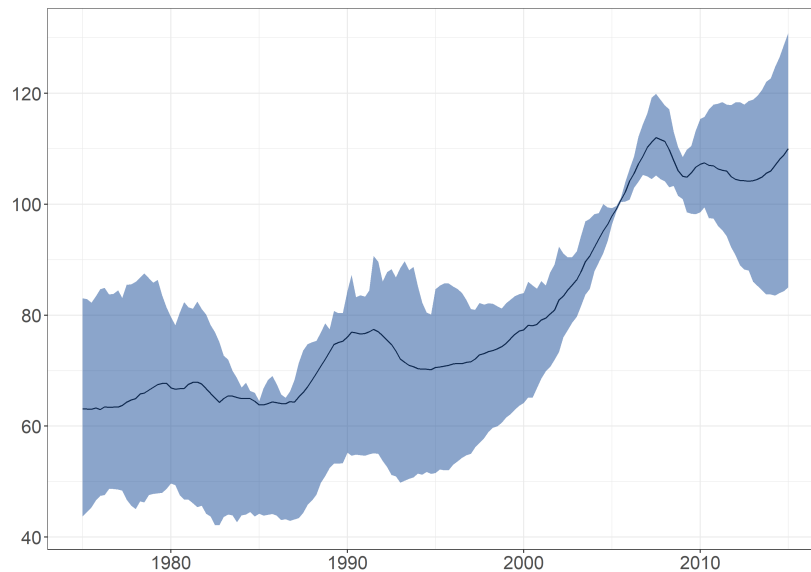


Figure 4: Real house prices.

prices we set the minimum window size equal to 36 quarters, the lag length equal to one, and the minimum duration of an episode of exuberance using the rule of Phillips *et al.* (2015a):

```
R> min_dur <- psy_ds(hprices)
R> min_dur
```

```
[1] 5
```

The fact that the minimum window size differs from its default value implies that we cannot use the stored critical values to draw inference. Instead, we need to generate a new `radf_cv` object. The estimation of the right-tailed unit root test statistics for all countries and the generation of Monte Carlo critical values can be easily achieved by using the functions `radf()` and `radf_mc_cv()`.

```
R> radf_hprices <- radf(hprices, minw = 36, lag = 1)
```

Using ``Date`` as index variable.

```
R> mc_critical_values <- radf_mc_cv(nrow(hprices), minw = 36, seed = 145)
```

Having obtained the `radf_obj` and `radf_cv` objects, there are now a number of different ways to display the estimation results. We first focus on the presentation of country-specific results and then illustrate how the user can obtain an overall picture of exuberance for the entire dataset. Following the same approach as for the simulated data, we apply the `summary()`, `datestamp()`, and `autoplot()` methods to the `radf_obj` and `radf_cv` objects. The former two methods return a list of elements, with each element containing results for a specific country. For illustration purposes, we focus on the United Kingdom (UK) and the United States (US) housing markets and report their chronology of exuberance under the GSADF

procedure.¹⁵ This choice reflects the economic size and significance of these two countries, but also the fact that their housing markets exemplify the distinct pattern of the boom-bust episode of the late 1990s and 2000s.

```
R> summary_results <- summary(radf_hprices, mc_critical_values)
R> summary_results[c("US", "UK")]
```

\$US

```
# A tibble: 3 x 5
  name  tstat  `90`  `95`  `99`
<fct> <dbl> <dbl> <dbl> <dbl>
1 adf   -1.03 -0.362 0.0530 0.787
2 sadf   4.15  0.993 1.31    1.77
3 gsadf  6.04  1.55  1.78    2.39
```

\$UK

```
# A tibble: 3 x 5
  name  tstat  `90`  `95`  `99`
<fct> <dbl> <dbl> <dbl> <dbl>
1 adf  -0.0907 -0.362 0.0530 0.787
2 sadf   2.94   0.993 1.31    1.77
3 gsadf  3.87   1.55  1.78    2.39
```

```
R> datestamp_results <- datestamp(radf_hprices, mc_critical_values,
+   min_duration = min_dur)
R> datestamp_results[c("US", "UK")]
```

\$US

```
      Start      End Duration
1 1998-10-01 2007-07-01      35
```

\$UK

```
      Start      End Duration
1 1987-07-01 1989-10-01      9
2 2000-04-01 2008-04-01     32
```

```
R> autoplot(radf_hprices, mc_critical_values, select_series = c("UK", "US"),
+   min_duration = min_dur)
```

According to the output of `summary()`, the null hypothesis of a unit root in real house prices is rejected in favor of explosive dynamics at the 5% significance level by the SADF and GSADF tests for both the US and UK, but not by the ADF test. Furthermore, the output of `datestamp()` and Figure 5 produced with `autoplot()` indicate that an episode of exuberance took place from the late 1990s/early 2000s to the late 2000s in both countries. For the UK, an episode of exuberance also took place in the first part of the sample period, from 1987 to 1989. Turning to the entire dataset, a way to concisely report first-stage results is via the function `tidy()`.

¹⁵As stated earlier, GSADF is the default option in `datestamp()` and `autoplot()`.

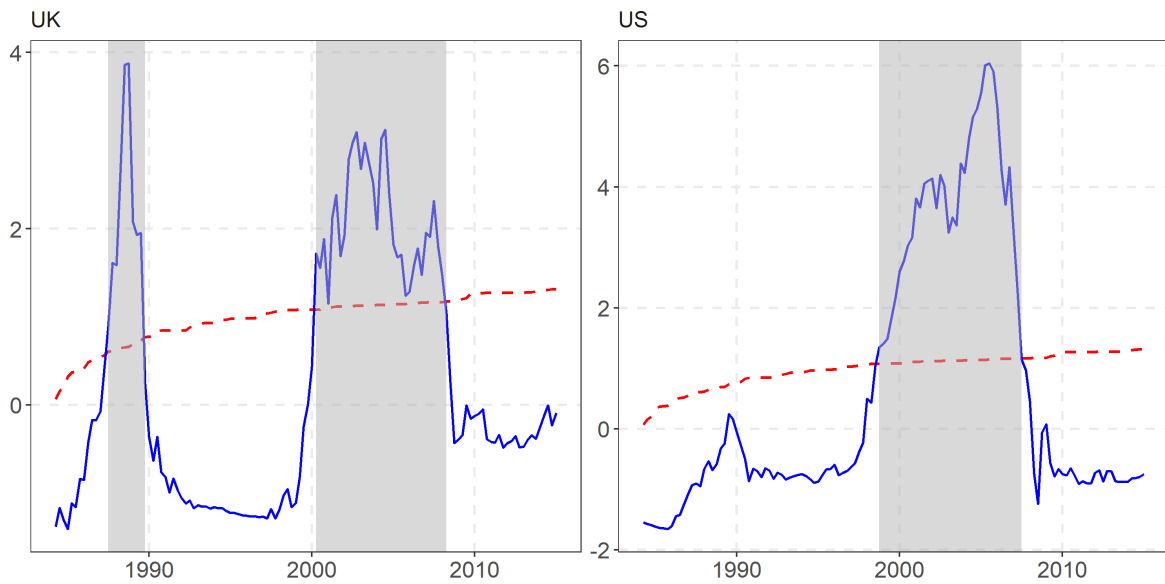


Figure 5: Date-stamping periods of exuberance in the UK and US housing markets.

```
R> tidy(radf_hprices)
```

```
# A tibble: 23 x 4
```

	id	adf	sadf	gsadf
	<fct>	<dbl>	<dbl>	<dbl>
1	Australia	0.502	2.82	5.47
2	Belgium	-0.0410	0.996	3.16
3	Canada	0.648	0.751	4.40
4	Switzerland	-2.19	2.39	3.27
5	Germany	-1.22	-0.140	2.50
6	Denmark	-1.02	2.00	3.01
7	Spain	-1.49	0.0704	4.01
8	Finland	-1.01	2.84	2.87
9	France	-0.618	3.34	5.32
10	UK	-0.0907	2.94	3.87

```
# ... with 13 more rows
```

```
R> tidy(mc_critical_values)
```

```
# A tibble: 3 x 4
```

	sig	adf	sadf	gsadf
	<fct>	<dbl>	<dbl>	<dbl>
1	90	-0.362	0.993	1.55
2	95	0.0530	1.31	1.78
3	99	0.787	1.77	2.39

Alternatively, one can use `diagnostics()` to quickly check the significance of the GSADF and SADF statistics.

```
R> diagnostics(radf_hprices, mc_critical_values)
```

```
-- Diagnostics (option = gsadf) ----- Monte Carlo --
```

```
Australia:  Rejects H0 at the 1% significance level
Belgium:    Rejects H0 at the 1% significance level
Canada:     Rejects H0 at the 1% significance level
Switzerland: Rejects H0 at the 1% significance level
Germany:    Rejects H0 at the 1% significance level
Denmark:    Rejects H0 at the 1% significance level
Spain:      Rejects H0 at the 1% significance level
Finland:    Rejects H0 at the 1% significance level
France:     Rejects H0 at the 1% significance level
UK:         Rejects H0 at the 1% significance level
Ireland:    Rejects H0 at the 1% significance level
Italy:      Cannot reject H0
Japan:      Rejects H0 at the 1% significance level
S. Korea:   Cannot reject H0
Luxembourg: Rejects H0 at the 1% significance level
Netherlands: Rejects H0 at the 1% significance level
Norway:     Rejects H0 at the 1% significance level
New Zealand: Rejects H0 at the 1% significance level
Sweden:     Rejects H0 at the 1% significance level
US:         Rejects H0 at the 1% significance level
S. Africa:  Rejects H0 at the 1% significance level
Croatia:    Rejects H0 at the 5% significance level
Israel:     Rejects H0 at the 10% significance level
```

```
R> diagnostics(radf_hprices, mc_critical_values, option = "sadf")
```

```
-- Diagnostics (option = sadf) ----- Monte Carlo --
```

```
Australia:  Rejects H0 at the 1% significance level
Belgium:    Rejects H0 at the 10% significance level
Canada:     Cannot reject H0
Switzerland: Rejects H0 at the 1% significance level
Germany:    Cannot reject H0
Denmark:    Rejects H0 at the 1% significance level
Spain:      Cannot reject H0
Finland:    Rejects H0 at the 1% significance level
France:     Rejects H0 at the 1% significance level
UK:         Rejects H0 at the 1% significance level
Ireland:    Rejects H0 at the 1% significance level
Italy:      Cannot reject H0
Japan:      Rejects H0 at the 5% significance level
S. Korea:   Cannot reject H0
Luxembourg: Rejects H0 at the 1% significance level
```





Figure 6: Date-stamping periods of exuberance in international housing markets.

Netherlands: Cannot reject H_0
 Norway: Rejects H_0 at the 1% significance level
 New Zealand: Rejects H_0 at the 1% significance level
 Sweden: Cannot reject H_0
 US: Rejects H_0 at the 1% significance level
 S. Africa: Cannot reject H_0
 Croatia: Cannot reject H_0
 Israel: Rejects H_0 at the 5% significance level

A comparison of the results of the two econometric tests reveals large differences. For the GSADF, there is strong evidence of exuberance in real house prices with the null hypothesis of a unit root being rejected for all but three countries at the 5% significance level. On the contrary, the number of rejections of the null hypothesis is substantially smaller for the SADF test. In particular, the test cannot reject the null for almost half of the countries in our sample (13 out of the 23) at the 5% significance level. Given the superior power properties of the GSADF procedure, the conclusion that emerges is that episodes of exuberance were widespread across housing markets.

Figure 6 provides a chronology of the identified periods of exuberance for each country. By far, the most interesting observation is the near-simultaneous exuberance in real house prices from the early to the mid-2000s. As evident from the figure, this phenomenon has no precedent at least in our sample period and hints to the possibility that a common factor may have contributed to house price exuberance spreading across countries.

```
R> autoplot(datestamp_results)
```

We explore this issue further by employing the panel recursive unit root testing procedure. This procedure requires the estimation of the panel GSADF and BSADF statistics and the




Figure 7: Date-stamping periods of overall exuberance in international housing markets.

generation of the corresponding critical values. Panel test statistics are already computed as a by-product of the univariate estimation and are included in the `radf` object `radf_hprices`. We next compute sieve bootstrap critical values with `radf_sb_cv()` and display the results.

```
R> panel_critical_values <- radf_sb_cv(hprices, minw = 36, lag = 1, seed = 145)
```

Using `Date` as index variable.

```
R> summary(radf_hprices, panel_critical_values)
```

```
-- Summary (minw = 36, lag = 1) ----- Sieve Bootstrap (nboot = 500) --
```

```
panel :
# A tibble: 1 x 5
  name      tstat `90` `95` `99`
<fct>    <dbl> <dbl> <dbl> <dbl>
1 gsadf_panel 2.06 0.105 0.122 0.155
```

```
R> autoplot(radf_hprices, cv = panel_critical_values, min_duration = min_dur)
R> datestamp(radf_hprices, panel_critical_values, min_duration = min_dur)
```

```
-- Datestamp (min_duration = 5) ----- Sieve Bootstrap --
```

```
panel :
      Start      End Duration
1 2001-01-01 2008-10-01      31
```

The panel results indicate that the null hypothesis of a unit root can be rejected at all conventional levels, providing strong evidence in favor of global exuberance. Furthermore, the date-stamping results illustrated in Figure 7 demonstrate in a clear manner the three phases of the international boom-bust episode in housing markets. The panel BSADF statistic sequence starts below its critical value at the beginning of the sample period; it increases rapidly after the mid-1990s and becomes significant during the early 2000s, with the period of global exuberance in housing markets continuing until 2006-07; and, finally, it collapses below its critical value.

6. Conclusion

Econometric tests of exuberance in asset prices have been attracting increasing interest over the last decade. This paper presents the R package **exuber** for recursive right-tailed unit root testing using the popular SADF and GSADF methodologies proposed by Phillips *et al.* (2011) and Phillips *et al.* (2015a,b), respectively, and the panel GSADF of Pavlidis *et al.* (2016). What differentiates **exuber** from existing implementations is its computational speed, the simple and intuitive API, and its ease of use. We demonstrated these features using simulation experiments, as well as with applications to artificial and real world data.

While several methodological advances have occurred over the last years, testing for and dating periods of exuberance remains a topic of ongoing econometric research. Our intention is to extend the statistical toolkit of **exuber** in future versions of the package to provide users with the state-of-the-art methodologies in this area. For the near future, we plan to include the reverse-regression algorithm for detecting bubble implosion of Phillips and Shi (2018), the end-of-sample bubble detection test of Astill, Harvey, Leybourne, and Taylor (2017), and the transmission of exuberance (contagion) test of Greenaway-McGrevy and Phillips (2016). We also plan to incorporate the Chow-type test and the modified versions of the Buseti and Taylor (2004) and Kim (2000) tests proposed by Homm and Breitung (2012).

References

- Araujo P, Lacerda G, Phillips PCB, Shi SP (2018). *MultipleBubbles: Test and Detection of Explosive Behaviors for Time Series*. R package version 0.2.0, URL <https://CRAN.R-project.org/package=MultipleBubbles>.
- Astill S, Harvey DI, Leybourne SJ, Taylor AMR (2017). “Tests for an End-of-Sample Bubble in Financial Time Series.” *Econometric Reviews*, **36**(6-9), 651–666. doi:10.1080/07474938.2017.1307490.
- Bhargava A (1986). “On the Theory of Testing for Unit Roots in Observed Time Series.” *The Review of Economic Studies*, **53**(3), 369–384. doi:10.2307/2297634.
- Blanchard OJ (1979). “Speculative Bubbles, Crashes and Rational Expectations.” *Economics Letters*, **3**(4), 387–389. doi:10.1016/0165-1765(79)90017-x.
- Buseti F, Taylor AMR (2004). “Tests of Stationarity against a Change in Persistence.” *Journal of Econometrics*, **123**(1), 33–66. doi:10.1016/j.jeconom.2003.10.028.

- Caspi I (2017). “**Rtadf**: Testing for Bubbles with EViews.” *Journal of Statistical Software*, **81**(1), 1–16. doi:10.18637/jss.v081.c01.
- Caspi I, Phillips PCB, Shi S (2019). **psymonitor**: *Real Time Monitoring of Asset Markets*. R package version 0.0.2, URL <https://CRAN.R-project.org/package=psymonitor>.
- Chang Y (2004). “Bootstrap Unit Root Tests in Panels with Cross-Sectional Dependency.” *Journal of Econometrics*, **120**(2), 263–293. doi:10.1016/s0304-4076(03)00214-8.
- Eddelbuettel D (2013). *Seamless R and C++ Integration with Rcpp*. Springer-Verlag, New York. doi:10.1007/978-1-4614-6868-4.
- Eddelbuettel D (2022). **drat**: ‘Drat’ R Archive Template. R package version 0.2.3, URL <https://CRAN.R-project.org/package=drat>.
- 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.
- Evans GW (1991). “Pitfalls in Testing for Explosive Bubbles in Asset Prices.” *American Economic Review*, **81**(4), 922–930. URL <https://www.jstor.org/stable/2006651>.
- Greenaway-McGrevy R, Grimes A, Holmes M (2019). “Two Countries, Sixteen Cities, Five Thousand Kilometres: How Many Housing Markets?” *Papers in Regional Science*, **98**(1), 353–370. doi:10.1111/pirs.12353.
- Greenaway-McGrevy R, Phillips PCB (2016). “Hot Property in New Zealand: Empirical Evidence of Housing Bubbles in the Metropolitan Centres.” *New Zealand Economic Papers*, **50**(1), 88–113. doi:10.1080/00779954.2015.1065903.
- Harvey DI, Leybourne SJ, Sollis R, Taylor AMR (2016). “Tests for Explosive Financial Bubbles in the Presence of Non-Stationary Volatility.” *Journal of Empirical Finance*, **38**, 548–574. doi:10.1016/j.jempfin.2015.09.002.
- Henry L, Wickham H (2020). **purrr**: *Functional Programming Tools*. R package version 0.3.4, URL <https://CRAN.R-project.org/package=purrr>.
- Homm U, Breitung J (2012). “Testing for Speculative Bubbles in Stock Markets: A Comparison of Alternative Methods.” *Journal of Financial Econometrics*, **10**(1), 198–231. doi:10.1093/jjfinec/nbr009.
- 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.
- Kane MJ, Emerson J, Weston S (2013). “Scalable Strategies for Computing with Massive Data.” *Journal of Statistical Software*, **55**(14), 1–19. doi:10.18637/jss.v055.i14.
- Kim JY (2000). “Detection of Change in Persistence of a Linear Time Series.” *Journal of Econometrics*, **95**(1), 97–116. doi:10.1016/s0304-4076(99)00031-7.

- Mack A, Martínez-García E, Grossman V (2011). “A Cross-Country Quarterly Database of Real House Prices: A Methodological Note.” *Technical report*, Globalization and Monetary Policy Institute. doi:10.24149/gwp99. Working Paper No. 99.
- 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.
- Martínez-García E, Grossman V (2020). “Explosive Dynamics in House Prices? An Exploration of Financial Market Spillovers in Housing Markets around the World.” *Journal of International Money and Finance*, **101**(102103). doi:10.1016/j.jimonfin.2019.102103.
- Microsoft, Weston S (2022). **foreach**: Provides Foreach Looping Construct. R package version 1.5.2, URL <https://CRAN.R-project.org/package=foreach>.
- Narayan PK, Mishra S, Sharma S, Liu R (2013). “Determinants of Stock Price Bubbles.” *Economic Modelling*, **35**, 661–667. doi:10.1016/j.econmod.2013.08.010.
- Otero J, Baum CF (2020). “**radf**: Stata Module to Calculate Unit Root Tests for Explosive Behaviour.” Statistical Software Components, Boston College Department of Economics. URL <https://EconPapers.repec.org/RePEc:boc:bocode:s458836>.
- Pavlidis EG, Martínez-García E, Grossman V (2019). “Detecting Periods of Exuberance: A Look at the Role of Aggregation with an Application to House Prices.” *Economic Modelling*, **80**, 87–102. doi:10.1016/j.econmod.2018.07.021.
- Pavlidis EG, Paya I, Peel DA (2017). “Testing for Speculative Bubbles Using Spot and Forward Prices.” *International Economic Review*, **58**(4), 1191–1226. doi:10.1111/iere.12249.
- Pavlidis EG, Yusupova A, Paya I, Peel DA, Martínez-García E, Mack A, Grossman V (2016). “Episodes of Exuberance in Housing Markets: In Search of the Smoking Gun.” *The Journal of Real Estate Finance and Economics*, **53**(4), 419–449. doi:10.1007/s11146-015-9531-2.
- Phillips PCB, Shi S, Yu J (2015a). “Testing for Multiple Bubbles: Historical Episodes of Exuberance and Collapse in the S&P 500.” *International Economic Review*, **56**(4), 1043–1078. doi:10.1111/iere.12132.
- Phillips PCB, Shi S, Yu J (2015b). “Testing for Multiple Bubbles: Limit Theory of Real-Time Detectors.” *International Economic Review*, **56**(4), 1079–1134. doi:10.1111/iere.12131.
- Phillips PCB, Shi SP (2018). “Financial Bubble Implosion and Reverse Regression.” *Econometric Theory*, **34**(4), 705–753. doi:10.1017/s0266466617000202.
- Phillips PCB, Wu Y, Yu J (2011). “Explosive Behavior in the 1990s NASDAQ: When Did Exuberance Escalate Asset Values?” *International Economic Review*, **52**(1), 201–226. doi:10.1111/j.1468-2354.2010.00625.x.
- R Core Team (2022). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.

- Sridharan S, Patel JM (2014). “Profiling R on a Contemporary Processor.” *Proceedings of the VLDB Endowment*, **8**(2), 173–184. doi:10.14778/2735471.2735478.
- Vasilopoulos K (2020a). *exuberdata*: Data for Critical Values in *exuber*. R package version 0.2.0, URL <https://github.com/kvasilopoulos/exuberdata>.
- Vasilopoulos K (2020b). *ihpdr*: Download Data from the International House Price Database. R package version 1.2.1, URL <https://CRAN.R-project.org/package=ihpdr>.
- Vasilopoulos K, Pavlidis E, Spavound S, Martínez-García E (2022). *exuber*: Econometric Analysis of Explosive Time Series. R package version 1.0.0, URL <https://CRAN.R-project.org/package=exuber>.
- Wang E, Cook D, Hyndman RJ (2020). “A New Tidy Data Structure to Support Exploration and Modeling of Temporal Data.” *Journal of Computational and Graphical Statistics*, **29**(3), 466–478. doi:10.1080/10618600.2019.1695624.
- West KD (1988). “Dividend Innovations and Stock Price Volatility.” *Econometrica*, **56**(1), 37–61. URL <https://www.jstor.org/stable/1911841>.
- Wickham H (2014). “Tidy Data.” *Journal of Statistical Software*, **59**(10), 1–23. doi:10.18637/jss.v059.i10.
- Wickham H (2016). *ggplot2: Elegant Graphics for Data Analysis*. 2nd edition. Springer-Verlag, New York. doi:10.1007/978-3-319-24277-4.
- Wickham H, François R, Henry L, Müller K (2022). *dplyr: A Grammar of Data Manipulation*. R package version 1.0.9, URL <https://CRAN.R-project.org/package=dplyr>.

Affiliation:

Kostas Vasilopoulos, Efthymios Pavlidis
Lancaster University

E-mail: k.vasilopoulos@lancaster.ac.uk, e.pavlidis@lancaster.ac.uk

URL: <https://kvasilopoulos.com/>, <https://sites.google.com/site/etpavlidis/>

Enrique Martínez-García

Federal Reserve Bank of Dallas

E-mail: emg.economics@gmail.com

URL: <https://sites.google.com/view/emgeconomics>

Journal of Statistical Software

published by the Foundation for Open Access Statistics

August 2022, Volume 103, Issue 10

doi:10.18637/jss.v103.i10

<https://www.jstatsoft.org/>

<https://www.foastat.org/>

Submitted: 2021-01-09

Accepted: 2021-10-20
