International Journal of Statistics and Applications 2022, 12(2): 30-41
DOI: 10.5923/j.statistics.20221202.02

Investigating the Persistence of Shocks: Global Warming, Economic Growth and Migration

Erhard Reschenhofer[1,*], Barbara Katharina Reschenhofer[2]

[1]Department of Statistics and Operations Research, University of Vienna, Oskar-Morgenstern-Platz 1, Vienna, Austria
[2]Department of English and American Studies, University of Vienna, Spitalgasse 2-4, Vienna, Austria

Abstract Choosing the value of 0.5 for the fractional differencing operator can be helpful for the determination of the stationarity of a time series. A pole at frequency zero of the spectral density of the fractionally differenced time series may indicate nonstationarity of the original time series (underdifferencing) whereas a vanishing spectral density at frequency zero may indicate stationarity of the original time series (overdifferencing). In addition to this frequency-domain analysis, it is advantageous to check in the time domain whether the autocorrelation function of the fractionally differenced time series is positive and slowly decaying. Unfortunately, carrying out fractional differencing is not a simple task unless the time series is extremely long, which is rarely the case in practice. We therefore propose a simple approximation which is based on a parsimonious ARMA(1,1) model. The new method is applied to climatological and socioeconomic datasets. The hypothesis of stationarity is rejected for the global surface temperature, economic growth, and migration.

Keywords Nonstationarity, Fractional differencing, Root differencing, Global surface temperature, GDP per capita, Immigration, Emigration

1. Introduction

The problem of determining whether a time series has a deterministic trend, a stochastic trend, or no trend at all is extremely difficult. Even when we are prepared to settle for vague answers, we need large samples. Unfortunately, there are very few time series of at least annual frequency which span over a period of hundreds of years. In this paper, we investigate three interesting examples, namely the Earth's global surface temperature from 1850 to 2021, the UK's GDP per capita from 1252 to 2018, as well as Swedish migration data from 1851 to 2020. In principle, climatological and socioeconomic variables ideally lend themselves to a related analysis due to the fact that these variables can very well impact one another. For example, Drake (2017) argues that the periodic weakening of the North Atlantic Oscillation (NAO) may have negatively affected the climate in parts of Europe and caused (at least in part) waves of migration to Italy, which eventually led to the fall of the Western Roman Empire in 476. A more recent example would be the Syrian drought from 2007 to 2010, which sparked mass movements of migration from rural farming areas to urban centers. This, in turn, may have moreover contributed to the unrest in Syria which began in 2011 and ended in a war resulting in millions of refugees (see Kelley et al., 2015). Of course, there is never only one, singular trigger causing a major historical event. Reports about strong relationships between environmental shocks and conflicts or wars (see, e.g., Miguel et al., 2004; Burke et al., 2009; Hsiang et al., 2011) must therefore be interpreted with caution (as argued by Buhaug, 2010, Theisen et al., 2011; see also Solow, 2013).

Although the investigation of correlations and causations between global warming, national income, and mass migration is undoubtedly a feat worth pursuing, the present paper will treat the three datasets separately. Aside from the fact that the datasets stem from different regions, the nature of this paper does not call for a joint analysis, as it is primarily concerned with the demonstrations of statistical methodology rather than the implications of its findings. More precisely, the main goal of our paper is the development of a new procedure for assessing the stationarity of a time series, whereby the individual time series merely serve to illustrate the usefulness of our method. Separate analyses are therefore entirely justified. This new method will be introduced in the next section, before then being applied to our three datasets in Section 3. Section 4 features a concluding discussion.

* Corresponding author:
erhard.reschenhofer@univie.ac.at (Erhard Reschenhofer)
Received: Feb. 12, 2022; Accepted: Mar. 2, 2022; Published: Mar. 15, 2022
Published online at http://journal.sapub.org/statistics

2. Methods

2.1. Checking Stationarity by Root Differencing

Rewriting a discrete-time stochastic process $y = (y_t)_{t\in\mathbb{Z}}$ which satisfies the difference equation

$$y_t - \phi y_{t-1} = u_t, \quad (1)$$

where $u = (u_t)_{t\in\mathbb{Z}}$ is white noise with mean 0 and variance σ^2, as

$$y_t = \phi y_{t-1} + u_t = \phi(\phi y_{t-2} + u_{t-1}) + u_t = \cdots = \sum_{j=0}^{\infty} \phi^j u_{t-j}, \quad (2)$$

we see that the effect of a past shock u_{t-j} on the current value y_t is only temporary and vanishes as $j \to \infty$ if $|\phi| < 1$. In contrast, if $\phi = 1$, we have

$$y_t = \sum_{j=0}^{\infty} u_{t-j} \quad (3)$$

and the effect of u_{t-j} on y_t is therefore persistent. In the former case, y is a stationary autoregressive (AR) process of order 1 with variance

$$var(y_t) = \sum_{j=0}^{\infty} (\phi^j)^2 var(u_{t-j}) = \sigma^2 \sum_{j=0}^{\infty} (\phi^2)^j = \frac{\sigma^2}{1-\phi^2} < \infty \quad (4)$$

whereas in the latter case, y is a random walk with infinite variance. However, based on a finite sample $y_1, \ldots, y_n$ of size *n*, it is often extremely difficult to distinguish between the two cases. Indeed, there is hardly any difference between the two samples of size $n = 100$ shown in Figure 1.a ($\phi = 1$) and Figure 1.c ($\phi = 0.95$), respectively. At first glance, both look nonstationary. The situation improves when the sample size is increased to $n = 1000$. While the sample from the random walk still looks nonstationary (see Figure 1.b), the sample from the AR process now looks quite stationary (see Figure 1.d). In practice, the dataset is given and can usually not be increased arbitrarily. Only if we are lucky and the parameter ϕ of the data generating process is sufficiently smaller than 1 for the given sample size, the unit root hypothesis

$$H_0: \phi = 1 \quad (5)$$

can be rejected with reasonable confidence either by a formal statistical test or by mere visual inspection of the time series plot (see, e.g., Figure 1.e). However, already when we move from the simple autoregressive scheme (1) to the slightly more general equation

$$y_t - \phi y_{t-1} = u_t + \theta u_{t-1}, \quad (6)$$

we can no longer hope for luck. Regardless of the sample size n, we can always choose a suitable value for the parameter θ so that the sample will look stationary even if $\phi = 1$ (see Figure 1.h). This is due to the fact that the terms in the numerator and denominator of the lag polynomial representation

$$y_t = \frac{1+\theta L}{1-\phi L} u_t \quad (7)$$

of (6) will nearly cancel out if θ is chosen only slighty greater than -1. Thus, we can never be sure whether a rejection of the unit root hypothesis is due to a small value of ϕ or a value of θ close to -1 (for a more thorough line of reasoning see Pötscher, 2002).

So in general, unit root testing is pointless and can only be justified under severe restrictions which are probably implausible in most applications. Where does that leave us when we still want to answer questions raised by climatologists, economists or social scientists? In light of the above considerations, it would be naïve to go for p-values which would only give us a false sense of security. Instead, we take a more pragmatic approach. Knowing that we will never be able to unmask a time series of the type shown in Figure 1.g as being nonstationary, we content ourselves with the construction of a tool which helps to determine whether trend patterns or long swings (like those shown in Figures 1.a and 1.c) are consistent with stationarity or not. In some simple cases, taking first differences already does the trick. No new method is required. Figure 2.a shows the first differences of the nonstationary time series displayed in Figure 1.a and Figure 2.e shows the first differences of the stationary time series displayed in Figure 1.c. Again, there is hardly any difference. However, with a great deal of good will, we may find that the second differenced series looks slightly "more stationary" than the first one. This is due to a negative autocorrelation caused by overdifferencing which tends to increase the number of crossings of the center line. A suitable way to make this negative correlation visible is to plot the sums

$$S_t(k) = \frac{1}{n-k} \sum_{s=k+1}^{t} (y_s - \bar{y})(y_{s-k} - \bar{y}), \text{ t=2,…,n,} \quad (8)$$

against t. While the first-order sample autocovariance $S_n(1)$ is negative in both cases (see Figures 2.b and 2.f), only the second cumulative plot shows a clear downward trend. However, this visual significance is somewhat put into perspective when higher-order lags are also considered (see Figures 2.c and 2.g). All computations are carried out with the free statistical software R (R Core Team, 2018).

Unfortunately, overdifferencing can also occur in the case of a nonstationary time series. When we consider the general class of autoregressive fractionally integrated moving average (ARFIMA) processes

$$y_t = \left(1 - \phi_1 L - \cdots - \phi_p L^p\right)^{-1} (1-L)^{-d} (1 + \theta_1 L + \cdots + \theta_q L^q) u_t \quad (9)$$

(see Granger and Joyeux, 1980; Hosking, 1981), there are not only I(1) processes, i.e., processes that are integrated of order one ($d = 1$), and I(0) processes, i.e., processes that are integrated of order zero ($d = 0$), but there are also processes that are fractionally integrated. For stationarity, it is required that $d < 0.5$. Our goal is to distinguish stationary processes with $0 \le d < 0.5$, which include white noise processes as well as AR, MA and ARMA processes, from nonstationary processes with $0.5 \le d \le 1$, which include random walks as well as ARIMA processes. We have seen above that a negative autocorrelation can be an indication of overdifferencing. Taking first differences reduces the memory parameter d by 1. In the case of an I(1) process, the memory parameter decreases from one to zero and is therefore still nonnegative whereas in the case of an I(0)

process, it decreases from 0 to -1. In contrast, both in the case of a fractionally integrated process with $d = 0.4$, which is stationary, and in the case of a fractionally integrated process with $d = 0.6$, which is nonstationary, the memory parameter will become negative after differencing. Checking for negativity after differencing is therefore pointless. An obvious alternative is fractional differencing with the help of the fractional differencing operator, which is defined as a power series expansion in integer powers of L, i.e.,

$$\Delta^d = (1-L)^d = 1 - dL + \frac{d(d-1)L^2}{2} - \frac{d(d-1)(d-2)L^3}{3!} + \cdots \quad (10)$$

Choosing $d = 0.5$ in (10) will reduce the memory parameter by 0.5, hence we will observe overdifferencing exactly in the stationary case where the original order of integration is less than 0.5. Fractional differencing with $d = 0.5$ can also help to detect underdifferencing. For example, when a strong positive autocorrelation is not only present in the original, not differenced series (see, e.g., Figures 2.d and 2.h) but also after fractional differencing, albeit to a lesser extent.

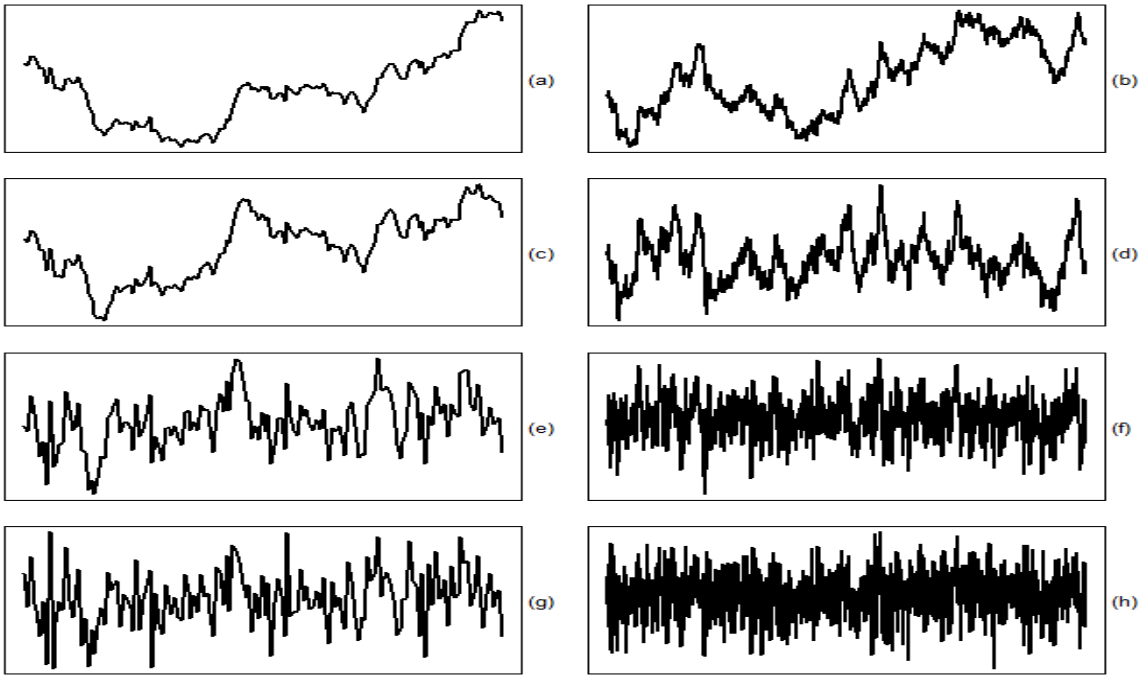

Figure 1. Realizations of lengths 150 (first column) and 1000 (second column), respectively, of a random walk (first row), an AR(1) process with $\phi = 0.95$ (second row), an AR(1) process with $\phi = 0.5$ (third row), and an ARMA(1,1) process with $\phi = 1$ and $\theta = -0.99$ (fourth row)

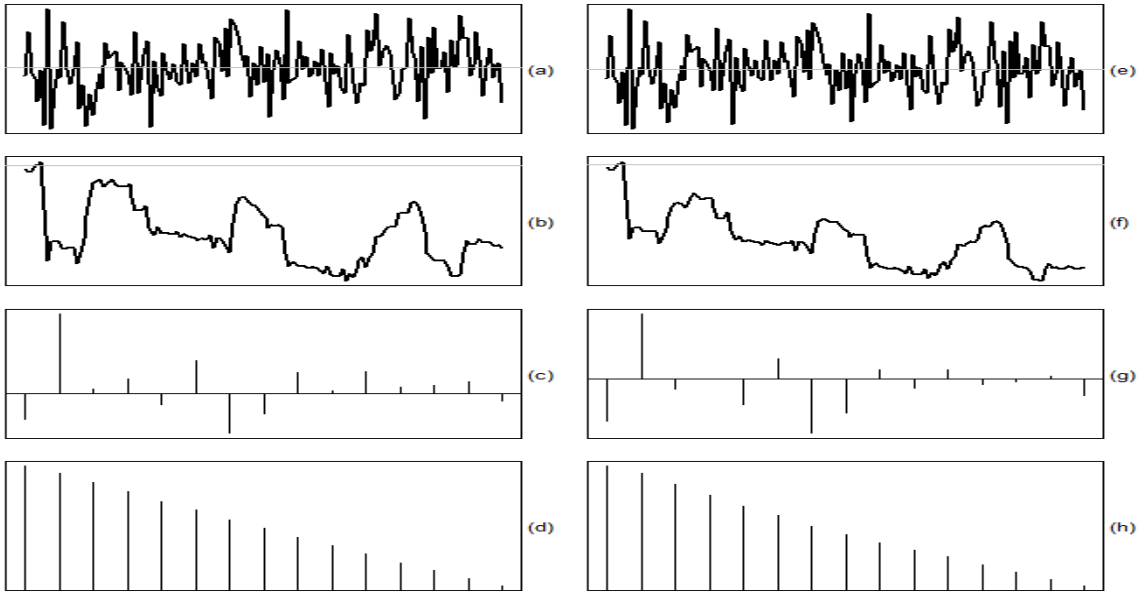

Figure 2. Random walk (first column) vs. AR(1) process with $\phi = 0.95$ (second column): (a), (e): First differences, (b), (f): First-order sample autocovariance plotted cumulatively, (c), (g) Sample autocovariances, (d), (h): Sample autocovariances of original series

The only problem with fractional differencing is that a truncated version of (10) must be used in practice which makes the first part of the obtained series unusable because it takes some time until the asymptotics kicks in. Clearly, giving up on a part of the data is not an option in many applications. E.g., for the investigation of climate change we definitely need the whole set of available historical temperature measurements and not just a part of it. In order to rectify this serious shortcoming of approximating the root differencing operator $\sqrt{\Delta} = \Delta^{0.5}$ by truncation, we will construct in Subsection 2.3 another approximation which is based on a parsimonious ARMA(1,1) model. But before that we will in the next subsection briefly leave the time domain and switch to the frequency domain where root differencing is a trivial exercise.

2.2. Analysis in the Frequency-Domain

The spectral density of an ARFIMA process is given by

$$f(\omega) = \frac{\sigma^2}{2\pi}\left|1 - e^{-i\omega}\right|^{-2d}\left|1 + \sum_{j=1}^{q}\theta_j e^{-i\omega j}\right|^2 \left|1 - \sum_{j=1}^{p}\phi_j e^{-i\omega j}\right|^{-2}. \quad (11)$$

In general, $f(\omega)$ either goes to infinity (if $d > 0$) or to zero (if $d < 0$) as $\omega \to 0$. Only in the case of a pure ARMA process ($d = 0$), it converges (horizontally) to a positive number. After root differencing, the spectral density goes to zero if and only if $d < 0.5$, i.e., exactly in the case of stationarity. In the frequency domain, root differencing can be accomplished simply by multiplying the spectral density by the factor

$$F(\omega) = \left|1 - e^{-i\omega}\right|, \quad (12)$$

hence we can easily assess the stationarity of a time series by estimating its spectral density, multiplying the estimate by the factor (12), and observing whether the transformed estimate decreases as $\omega \to 0$. Nonparametric estimates of the spectral density can be obtained by smoothing the periodogram

$$I(\omega_j) = \frac{1}{2\pi n}\left|\sum_{t=1}^{n} y_t e^{-i\omega_j t}\right|^2, \quad (13)$$

which is the sample spectral density evaluated at the Fourier frequencies $\omega_j = 2\pi j/n$.

A more elaborate method is the log periodogram regression which is based on the low-frequency approximation

$$\log f(\omega) \approx C + \log\left|1 - e^{-i\omega}\right|^{-2d} = C + d\left(-2\log\left|1 - e^{-i\omega}\right|\right). \quad (14)$$

To get an estimate of the memory parameter d, we must replace the spectral density in this approximation by the periodogram (see Geweke and Porter-Hudak, 1983) or a smoothed version of it (see Hassler, 1993, Peiris and Court, 1993, and Reisen, 1994) and choose a suitable neighborhood $\omega_1, \ldots, \omega_K$ of frequency zero. The parameter K determines how many of the lowest Fourier frequencies are included in the regression (for a procedure involving non-Fourier frequencies, see Reschenhofer and Mangat, 2021). As always, there is a trade-off between bias and variance. A small value of K increases the variance whereas a large value of K may introduce a bias caused by short-term autocorrelation.

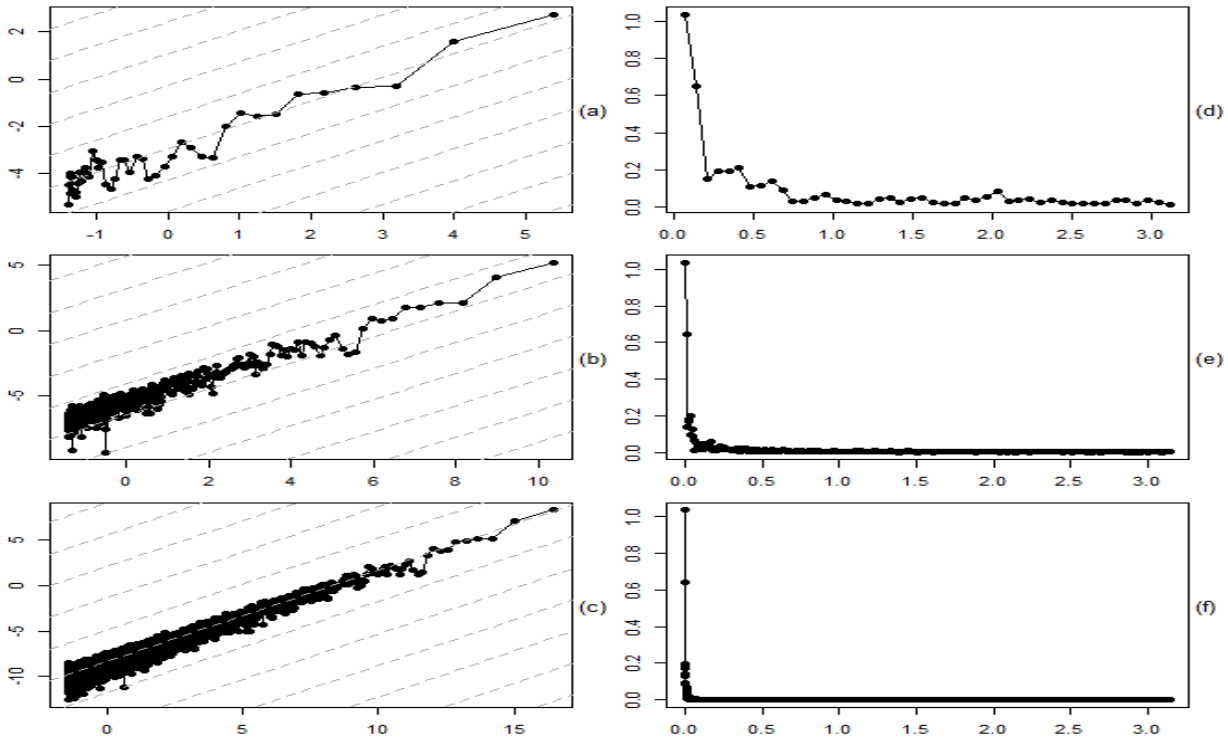

Figure 3. For each frequency (first row: annual, second row: monthly, third row: daily), we first plot the log periodogram of the log S&P 500 from 1928 to 2020 against $-2\log\left|1 - e^{-i\omega_j}\right|$ and compare it to dashed lines of slope 1 and then we plot the root differenced periodogram against $\omega_j = 2\pi j/n$

In order to illustrate the frequency-domain approach outlined above, we consider a financial application where d is known a priori. Although stock price series cannot adequately be modeled by a random walk (and not even by a conditionally heteroskedastic random walk with drift), there is a broad agreement that they are integrated of order 1, hence $d = 1$. A disadvantage of financial datasets is that they are usually very short. Stock prices are rarely available for hundred years or more. A large sample size is not helping in this regard. E.g., for the investigation of climate change, a long annual temperature series from 1850 to 2020 ($n = 171$) is certainly more appropriate than a short daily series from 2001 to 2020 ($n = 7305$). Indeed, Figure 3 shows that simply increasing the resolution of a time series does not change the shape of the periodogram in the low-frequency range. For each frequency, annual (3.a), monthly (3.b), and daily (3.c), the scatter plot based on the low-frequency relationship approximation (14) corroborates our suspicion that $d = 1$. Moreover, after multiplication by the factor (12), the periodograms still increase as $\omega \to 0$ (see 3.d-f), which is inconsistent with stationarity of the original series.

2.3. Approximating the Root Differencing Operator

Usually, we interpret an ARMA equation such as (6), which can in lag operator notation also be written as

$$y_t = \frac{1+\theta L}{1-\phi L} u_t, \tag{15}$$

as a description of the transformation of a simple white noise u to a more complex process y. In the frequency domain, this transformation is accomplished by multiplying the constant spectral density

$$f(\omega) = \frac{\sigma^2}{2\pi} \tag{16}$$

by the factor

$$\left|\frac{1+\theta e^{-i\omega}}{1-\phi e^{-i\omega}}\right|^2. \tag{17}$$

However, in the case of differencing, we prefer to interpret the equation

$$(1-L)y_t = u_t, \tag{18}$$

as a description of the transformation of a possibly nonstationary process y to a possibly stationary process u. The spectral density of u is obtained from the spectral density of y by multiplication with the factor

$$\left|1 - e^{-i\omega}\right|^2, \tag{19}$$

which vanishes at frequency zero. The steep decline towards zero can be reduced either by replacing the unit root in (18) by a near unit root, i.e.,

$$(1-\phi L)y_t = u_t,\ \phi = 1-\varepsilon, \tag{20}$$

or more flexibly by introducing an almost cancelling root on the right hand side of (20), i.e.,

$$(1-\phi L)y_t = (1+\theta u_t),\ \phi = 1-\varepsilon,\ \theta = -1+\delta. \tag{21}$$

Clearly, ε must be smaller than δ or else the decline will vanish completely. The result of the dampened differencing is then given by

$$u_t = \frac{1-\phi L}{1+\theta L} y_t. \tag{22}$$

In the frequency domain, this transformation comes down to a multiplication of the original spectral density by the factor

$$\left|1-\phi e^{-i\omega}\right|^2 \left|1+\theta e^{-i\omega}\right|^{-2}. \tag{23}$$

For an approximation of the root differencing operator $\Delta^{0.5}$, we need to find suitable values of ϕ and θ so that a plot of the log of (23) against $2\log\left|1-e^{-i\omega_j}\right|$ has approximately a slope of 0.5 in the neighborhood of frequency zero.

Table 1 gives pairs of values of ϕ and θ for various sample sizes which mimic the effect of root differencing. These values were found by minimization of

$$\begin{aligned}&\sum_{j=1}^{K}\left(\log\left(\left|1-\phi e^{-i\omega_j}\right|^2\left|1+\theta e^{-i\omega_j}\right|^{-2}\right) - \log\left(\left|1-e^{-i\omega_j}\right|^{2\cdot 0.5}\right) + c\right)^2\\ &= \sum_{j=1}^{K}\left(\log\left(\left|1+\theta e^{-i\omega_j}\right|^2\left|1-\phi e^{-i\omega_j}\right|^{-2}\right) - 0.5\left(-2\log\left(\left|1-e^{-i\omega_j}\right|\right)\right) + C\right)^2\end{aligned} \tag{24}$$

with respect to ϕ, θ, and c, where $K = \left[\sqrt{n/2}\right]$. The values in Table 1 can also be used for the approximation of the root integration operator $\Delta^{-0.5}$, which is just the inverse operator of the root differencing operator $\Delta^{0.5}$. Indeed, if u is obtained from y by approximate root differencing, i.e.,

$$u_t = -\theta u_{t-1} + y_t - \phi y_t, \tag{25}$$

then y can be obtained from u by approximate root integrating, i.e.,

$$y_t = \phi y_t + u_t + \theta u_{t-1}. \tag{26}$$

Using the values in Table 1, we plotted spectral densities of the form

$$f(\omega) \propto \left|1+\theta e^{-i\omega}\right|^2 \left|1-\phi e^{-i\omega}\right|^{-2} \tag{27}$$

against the transformed Fourier frequencies $-2\log\left(\left|1-e^{-i\omega_j}\right|\right)$ in Figures 4.b, e, h. Of special interest is the low frequency range (see Figures 4.a, d, g), where the graphs are approximately linear with slope 0.5. In this frequency range, the fit obtained by truncated series approximations is generally worse (see Figures 4.c, f, i). In the following, we will therefore change our setting slightly to make up for the shortcomings of the latter approximation. Firstly, we will introduce an initial settling period of length $w = [2\sqrt{n}]$ and thereby reduce the analysis period from n to $n-w$. Secondly, we will use an expanding cut-off lag instead of a fixed one.

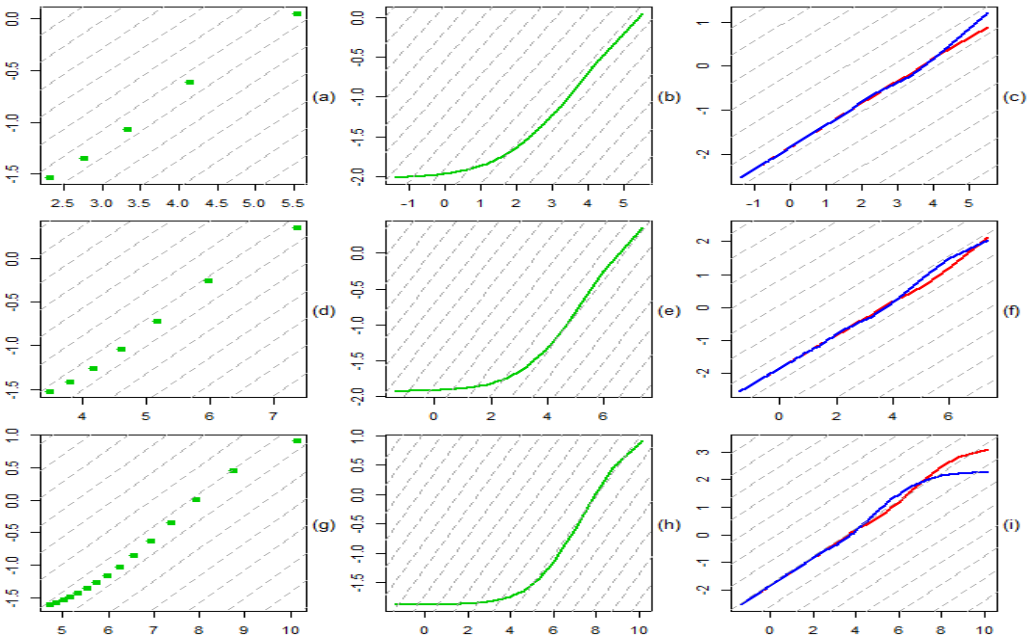

Figure 4. For sample sizes $n = 100$ (first row), 250 (second row), and 1000 (third row), selected log ARMA(1,1) spectral densities are plotted against $-2\log(|1 - e^{-i\omega_j}|)$ for $j = 1, \ldots, \lfloor\sqrt{n}/2\rfloor$ (first column) and $j = 1, \ldots, \lfloor(n-1)/2\rfloor$ (second column), respectively, and are then compared to dashed base lines with slope 0.5. Note that the low frequencies appear on the right side! The spectral shapes in the third column were obtained by truncated series approximations of the root differencing operator with cut-off lags 20 (blue) and 50 (red), respectively

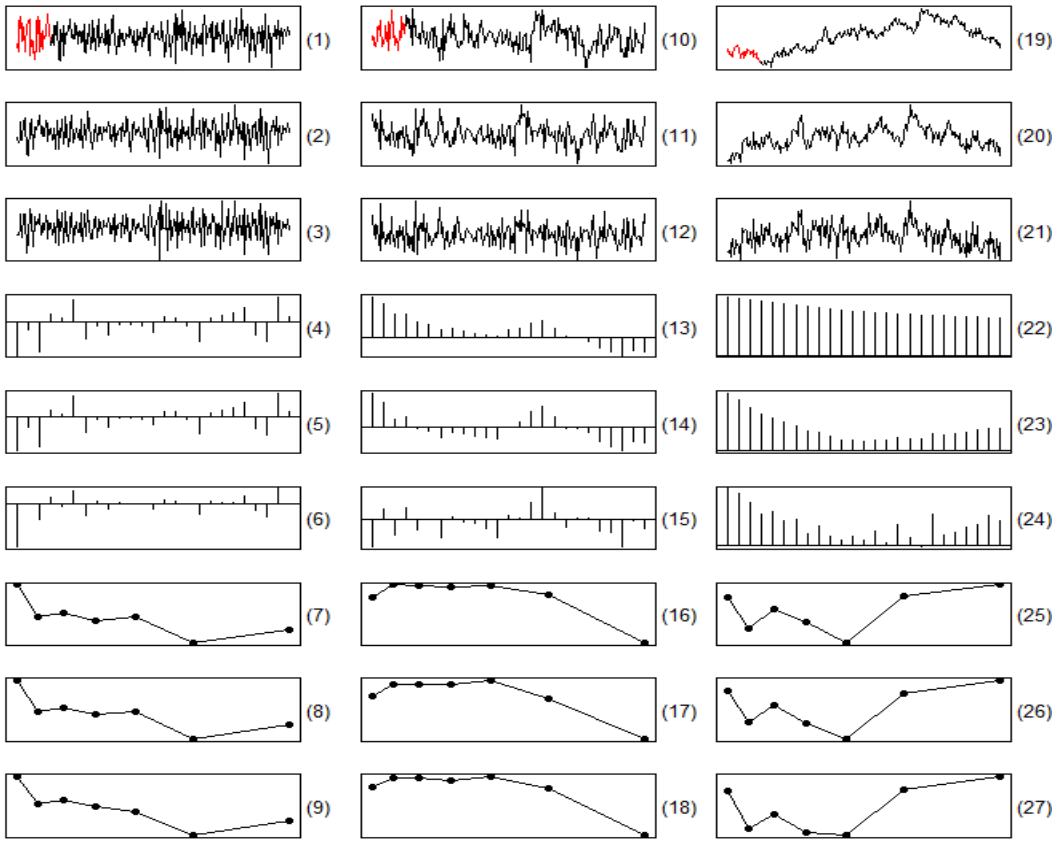

Figure 5. Analysis of fractional series with $d = -0.3$ (1st), $d = 0.3$ (2nd), $d = 0.8$ (3rd column): 1st row: Realizations with initial settling period (red, n=30) and analysis period (black, n=220); 2nd row: Root differenced series obtained with ARMA(1,1) approximation; 3rd row: Root differenced series obtained with truncated series (expanding cut-off lag); Rows 4-6: Sample autocorrelations of the original series and the two root differenced versions; Rows 7-9: Log periodogram plots ($K = 13$) for the root differenced series; (7: root differencing in the frequency domain, 8: ARMA(1,1), 9: truncated series)

Table 1. Approximating root differencing $u_t = (1-L)^{0.5} y_t$ by $u_t = -\theta u_{t-1} + y_t - \phi y_t$ (sample size, number of lowest Fourier Frequencies used for fitting, ϕ, θ)

90	4	0.926	-0.767	**180**	6	0.956	-0.846	**630**	12	0.983	-0.929
91	4	0.927	-0.770	**185**	6	0.957	-0.850	**635**	12	0.983	-0.929
92	4	0.928	-0.773	**190**	6	0.958	-0.853	**640**	12	0.983	-0.930
93	4	0.928	-0.773	**195**	6	0.960	-0.859	**645**	12	0.983	-0.930
94	4	0.929	-0.775	**200**	7	0.958	-0.849	**650**	12	0.983	-0.930
95	4	0.930	-0.778	**205**	7	0.959	-0.852	**655**	12	0.983	-0.930
96	4	0.931	-0.781	**210**	7	0.960	-0.856	**660**	12	0.984	-0.933
97	4	0.931	-0.781	**215**	7	0.961	-0.859	**665**	12	0.984	-0.933
98	4	0.932	-0.784	**220**	7	0.962	-0.862	**670**	12	0.984	-0.933
99	4	0.933	-0.787	**225**	7	0.963	-0.866	**675**	12	0.984	-0.933
100	5	0.928	-0.765	**230**	7	0.963	-0.866	**680**	13	0.983	-0.929
101	5	0.928	-0.765	**235**	7	0.964	-0.869	**685**	13	0.983	-0.929
102	5	0.929	-0.768	**240**	7	0.965	-0.873	**690**	13	0.984	-0.932
103	5	0.930	-0.771	**245**	7	0.966	-0.876	**695**	13	0.984	-0.932
104	5	0.930	-0.771	**250**	7	0.966	-0.876	**700**	13	0.984	-0.932
105	5	0.931	-0.774	**255**	7	0.967	-0.880	**705**	13	0.984	-0.932
106	5	0.932	-0.777	**260**	8	0.966	-0.873	**710**	13	0.984	-0.933
107	5	0.932	-0.777	**265**	8	0.966	-0.873	**715**	13	0.984	-0.933
108	5	0.933	-0.780	**270**	8	0.967	-0.877	**720**	13	0.984	-0.933
109	5	0.934	-0.783	**275**	8	0.968	-0.880	**725**	13	0.984	-0.933
110	5	0.934	-0.783	**280**	8	0.968	-0.880	**730**	13	0.984	-0.933
111	5	0.935	-0.786	**285**	8	0.969	-0.884	**735**	13	0.985	-0.936
112	5	0.935	-0.786	**290**	8	0.969	-0.884	**740**	13	0.985	-0.936
113	5	0.936	-0.789	**295**	8	0.970	-0.887	**745**	13	0.985	-0.936
114	5	0.936	-0.789	**300**	8	0.970	-0.887	**750**	13	0.985	-0.937
115	5	0.937	-0.792	**305**	8	0.971	-0.891	**755**	13	0.985	-0.937
116	5	0.937	-0.792	**310**	8	0.971	-0.891	**760**	13	0.985	-0.937
117	5	0.938	-0.795	**315**	8	0.972	-0.894	**765**	13	0.985	-0.937
118	5	0.938	-0.795	**320**	8	0.972	-0.895	**770**	13	0.985	-0.937
119	5	0.939	-0.798	**325**	9	0.971	-0.889	**775**	13	0.985	-0.937
120	5	0.939	-0.798	**330**	9	0.971	-0.889	**780**	13	0.985	-0.937
121	5	0.940	-0.801	**335**	9	0.972	-0.892	**785**	14	0.985	-0.936
122	5	0.940	-0.801	**340**	9	0.972	-0.892	**790**	14	0.985	-0.936
123	5	0.941	-0.804	**345**	9	0.973	-0.896	**795**	14	0.985	-0.936
124	5	0.941	-0.804	**350**	9	0.973	-0.896	**800**	14	0.985	-0.936
125	5	0.942	-0.807	**355**	9	0.973	-0.896	**805**	14	0.985	-0.936
126	5	0.942	-0.807	**360**	9	0.974	-0.899	**810**	14	0.985	-0.936
127	5	0.943	-0.810	**365**	9	0.974	-0.900	**815**	14	0.986	-0.939
128	5	0.943	-0.810	**370**	9	0.974	-0.900	**820**	14	0.986	-0.939
129	5	0.944	-0.813	**375**	9	0.975	-0.903	**825**	14	0.986	-0.940
130	5	0.944	-0.813	**380**	9	0.975	-0.903	**830**	14	0.986	-0.940
131	5	0.944	-0.813	**385**	9	0.975	-0.904	**835**	14	0.986	-0.940
132	5	0.945	-0.816	**390**	9	0.976	-0.907	**840**	14	0.986	-0.940
133	5	0.945	-0.816	**395**	9	0.976	-0.907	**845**	14	0.986	-0.940
134	5	0.946	-0.819	**400**	10	0.975	-0.901	**850**	14	0.986	-0.940
135	5	0.946	-0.819	**405**	10	0.976	-0.905	**855**	14	0.986	-0.940
136	5	0.946	-0.819	**410**	10	0.976	-0.905	**860**	14	0.986	-0.940
137	5	0.947	-0.822	**415**	10	0.976	-0.905	**865**	14	0.986	-0.940

138	5	0.947	-0.823	**420**	10	0.976	-0.905	**870**	14	0.986	-0.940
139	5	0.947	-0.823	**425**	10	0.977	-0.909	**875**	14	0.987	-0.944
140	5	0.948	-0.826	**430**	10	0.977	-0.909	**880**	14	0.987	-0.944
141	5	0.948	-0.826	**435**	10	0.977	-0.909	**885**	14	0.987	-0.944
142	5	0.949	-0.829	**440**	10	0.977	-0.909	**890**	14	0.987	-0.944
143	5	0.949	-0.829	**445**	10	0.978	-0.912	**895**	14	0.987	-0.944
144	6	0.946	-0.814	**450**	10	0.978	-0.913	**900**	15	0.986	-0.940
145	6	0.946	-0.814	**455**	10	0.978	-0.913	**905**	15	0.987	-0.943
146	6	0.946	-0.814	**460**	10	0.978	-0.913	**910**	15	0.987	-0.943
147	6	0.947	-0.817	**465**	10	0.979	-0.916	**915**	15	0.987	-0.943
148	6	0.947	-0.817	**470**	10	0.979	-0.916	**920**	15	0.987	-0.943
149	6	0.947	-0.817	**475**	10	0.979	-0.917	**925**	15	0.987	-0.943
150	6	0.948	-0.820	**480**	10	0.979	-0.917	**930**	15	0.987	-0.943
151	6	0.948	-0.820	**485**	11	0.979	-0.915	**935**	15	0.987	-0.943
152	6	0.948	-0.820	**490**	11	0.979	-0.915	**940**	15	0.987	-0.943
153	6	0.949	-0.823	**495**	11	0.979	-0.915	**945**	15	0.987	-0.943
154	6	0.949	-0.824	**500**	11	0.979	-0.915	**950**	15	0.987	-0.943
155	6	0.949	-0.824	**505**	11	0.979	-0.915	**955**	15	0.987	-0.944
156	6	0.950	-0.827	**510**	11	0.980	-0.919	**960**	15	0.987	-0.944
157	6	0.950	-0.827	**515**	11	0.980	-0.919	**965**	15	0.987	-0.944
158	6	0.950	-0.827	**520**	11	0.980	-0.919	**970**	15	0.987	-0.944
159	6	0.951	-0.830	**525**	11	0.980	-0.919	**975**	15	0.987	-0.944
160	6	0.951	-0.830	**530**	11	0.980	-0.919	**980**	15	0.988	-0.947
161	6	0.951	-0.830	**535**	11	0.981	-0.922	**985**	15	0.988	-0.947
162	6	0.952	-0.833	**540**	11	0.981	-0.923	**990**	15	0.988	-0.947
163	6	0.952	-0.833	**545**	11	0.981	-0.923	**995**	15	0.988	-0.947
164	6	0.952	-0.833	**550**	11	0.981	-0.923	**1000**	15	0.988	-0.947
165	6	0.952	-0.833	**555**	11	0.981	-0.923	**1005**	15	0.988	-0.947
166	6	0.953	-0.836	**560**	11	0.981	-0.923	**1010**	15	0.988	-0.947
167	6	0.953	-0.836	**565**	11	0.982	-0.926	**1015**	15	0.988	-0.948
168	6	0.953	-0.837	**570**	11	0.982	-0.927	**1020**	15	0.988	-0.948
169	6	0.953	-0.837	**575**	11	0.982	-0.927	**1025**	16	0.988	-0.946
170	6	0.954	-0.840	**580**	12	0.981	-0.922	**1030**	16	0.988	-0.947
171	6	0.954	-0.840	**585**	12	0.981	-0.922	**1035**	16	0.988	-0.947
172	6	0.954	-0.840	**590**	12	0.982	-0.925	**1040**	16	0.988	-0.947
173	6	0.955	-0.843	**595**	12	0.982	-0.925	**1045**	16	0.988	-0.947
174	6	0.955	-0.843	**600**	12	0.982	-0.925	**1050**	16	0.988	-0.947
175	6	0.955	-0.843	**605**	12	0.982	-0.926	**1055**	16	0.988	-0.947
176	6	0.955	-0.843	**610**	12	0.982	-0.926	**1060**	16	0.988	-0.947
177	6	0.956	-0.846	**615**	12	0.982	-0.926	**1065**	16	0.988	-0.947
178	6	0.956	-0.846	**620**	12	0.982	-0.926	**1070**	16	0.988	-0.947
179	6	0.956	-0.846	**625**	12	0.983	-0.929	**1075**	16	0.988	-0.947

Figure 5 shows that these measures are quite effective. There are no major discrepancies between the two approximation methods. Both agree that the root differenced fractional series with $d = -0.3$ and $d = 0.3$, respectively, are overdifferenced as indicated by a log periodogram that decreases as the frequency decreases (see Figures 5.8, 9, 17, 18), and that a root differenced fractional series with $d = 0.8$ is underdifferenced as indicated by a log periodogram that still increases as the frequency decreases (see Figures 5.25, 26). These findings are also in line with the results obtained by root differencing in the frequency domain (see Figures 5.7, 16, 25). Moreover, there is also agreement that the strong positive autocorrelation, which is present in the fractional series with $d = 0.8$ (see Figure 5.22), is (to a lesser extent) still present after root differencing (see Figures 5.23, 24). In contrast, the analogous plots for the stationary

series are inconclusive (see Figures 5.4-6, 13-15). However, the truncated series approximation has at least managed to produce a negative first-order autocorrelation in each case (Figure 5.6, 15). In general, this method appears to remove a possible (stochastic) trend slightly more aggressively than the ARMA(1,1) approximation (see Figures 5.1-3, 20-12, 19-21).

3. Empirical Results

3.1. Global Warming

There has been a broad and often controversial discussion about the true nature of the apparent upward trend in global surface temperature (see, e.g., Stern and Kaufmann, 1999; Fomby and Vogelsang, 2002; Kaufmann and Stern, 2002; Kaufmann et al., 2006; Kaufmann et al., 2010; Reschenhofer, 2013; Estrada et al., 2013; Estrada et al., 2017; Estrada and Perron, 2019; Lai and Yoon, 2018; Mangat and Reschenhofer, 2020; Chang et al., 2020). We add to this debate with an empirical study based on root differencing. Following the recommendation of the Climatic Research Unit (CRU) of the University of East Anglia (UEA), we use the HadCRUT5 dataset (Morice et al., 2021) for our investigation of the change of the global temperature since 1850. This dataset contains the global annual means from 1850 to 2020 of the combined land and marine temperature anomalies. Temperature anomalies are defined relative to the 1961–1990 temperature mean. The dataset is available at the website https://sites.uea.ac.uk/cru/data of the CRU. Although the observation period is not very long, we will again introduce an initial settling period of length $[2\sqrt{n}]$. The resulting reduction of the analysis period from 171 to 145 years can to some extent be justified by the fact that the global means of the first years are based on a significantly smaller number of measurements. In particular, this is true for the years before 1880. Note that a similar surface temperature dataset, namely the GISTEMP v4 (see GISTEMP Team, 2021; Lenssen, 2019) provided by the NASA (https://data.giss.nasa.gov/gistemp/), is only available from 1880.

The central question is whether the recent rise in temperature is just a transient phenomenon or an indication of nonstationarity. Applying the methods used for the production of Figure 5 to our global surface temperature series, we find that all indications point to nonstationarity. Firstly, the root differenced series (obtained from the mean-corrected original series) still exhibit some kind of an upward trend (see Figures 6.d, g). Secondly, there is also strong positive autocorrelation left after root differencing (see Figures 6.e, h). Thirdly, the log periodogram of the root differenced series always increases as the frequency decreases, regardless whether root differencing is carried out in the frequency domain (see Figure 6.c) or in the time domain (see Figures 6.f, i).

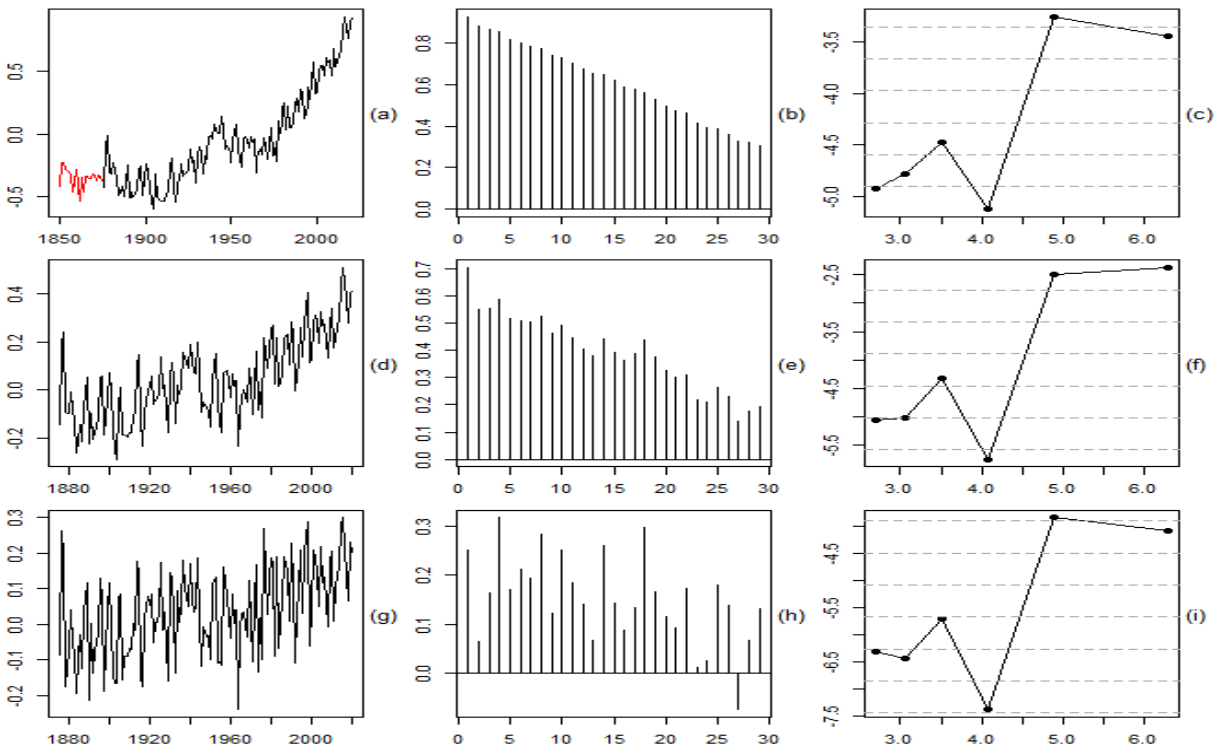

Figure 6. Analysis of global surface temperature from 1850 to 2020 (HadCRUT5): 1st column: Plots of original (a) and two root differenced series (d: ARMA, g: truncated); 2nd column: Sample autocorrelations of original and root differenced series; 3rd column: Log periodogram plots after root differencing (c: frequency domain differencing)

3.2. Economic Growth

We downloaded the real UK GDP per capita since 1252 (in 2011$) from the Maddison Project Database (see Bolt and van Zanden, 2020; Scheidel and Friesen, 2009; Stohr, 2016), which is maintained by the Groningen Growth and Development Centre at the University of Groningen. The 2020 version of this database covers the period up to 2018. In contrast to the previous section, we take the logarithm of the data before we carry out the analysis. Again, the empirical results are in favor of nonstationarity, only much stronger. There is stunning evidence of underdifferencing (see Figure 7). Of course, this does not come as a surprise. In the case of economic growth, it is a priori clear that there is an upward trend. The only question is whether this trend is deterministic or stochastic. In the latter case, it is safe to assume that $d \geq 1$, hence root differencing is definitely not enough.

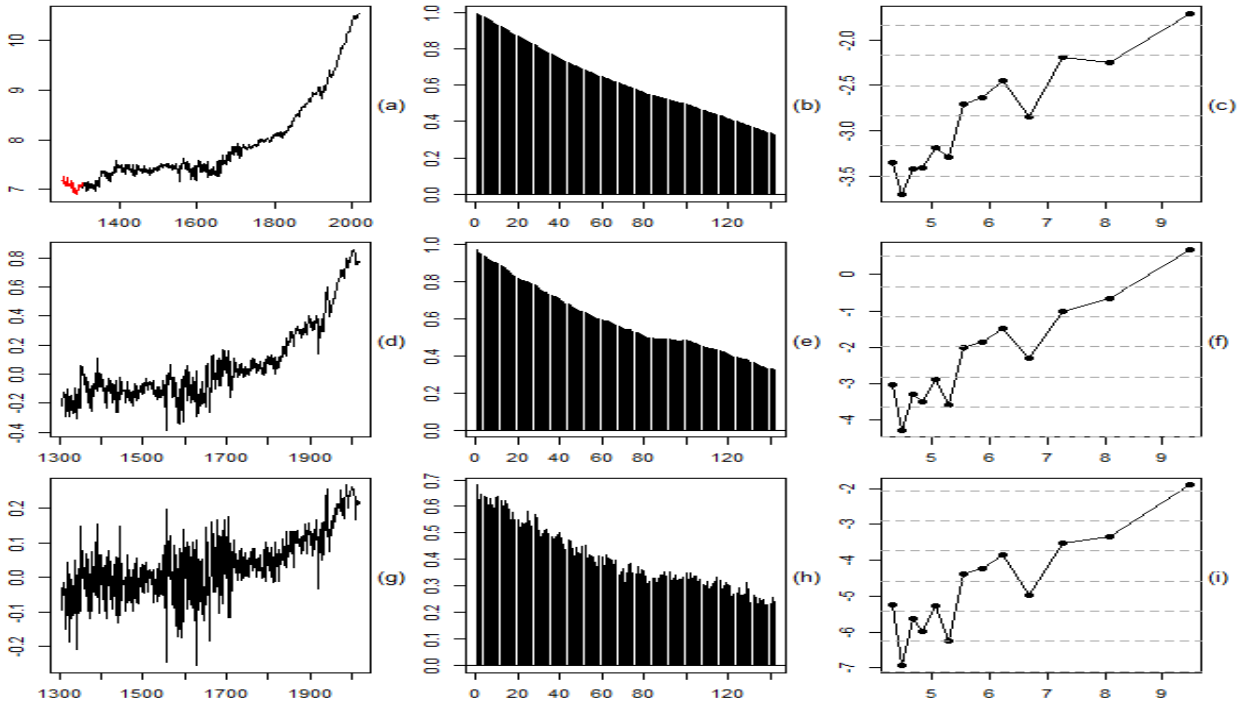

Figure 7. Analysis of log UK GDP per capita from 1252 to 2018 (Maddison Project Database): 1st column: Plots of original (a) and two root differenced series (d: ARMA, g: truncated); 2nd column: Sample autocorrelations of original and root differenced series; 3rd column: Log periodogram plots after root differencing (c: frequency domain differencing)

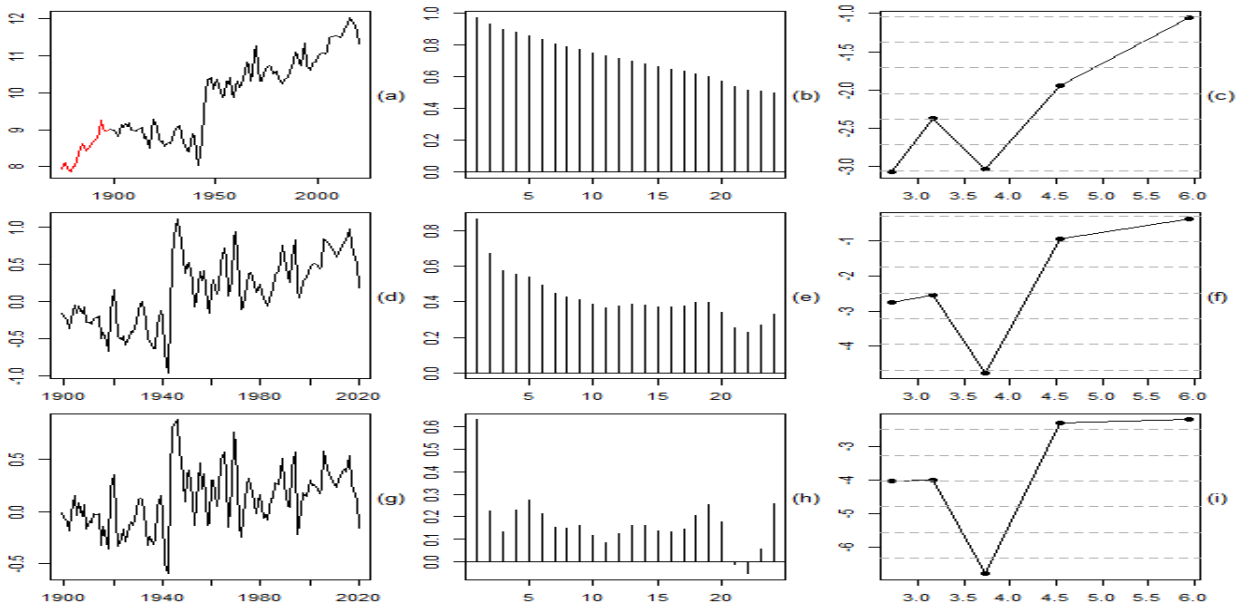

Figure 8. Analysis of log Swedish immigration numbers from 1875 to 2020: 1st column: Plots of original (a) and two root differenced series (d: ARMA, g: truncated); 2nd column: Sample autocorrelations of original and root differenced series; 3rd column: Log periodogram plots after root differencing (c: frequency domain differencing)

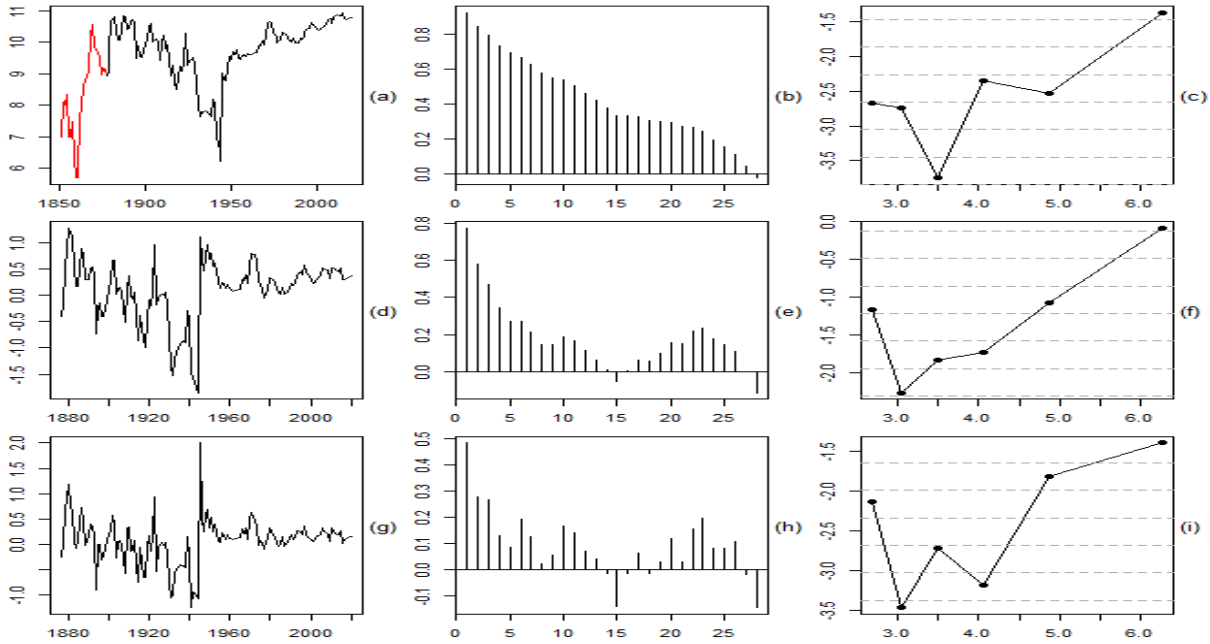

Figure 9. Analysis of log Swedish emigration numbers from 1851 to 2020: 1st column: Plots of original (a) and two root differenced series (d: ARMA, g: truncated); 2nd column: Sample autocorrelations of original and root differenced series; 3rd column: Log periodogram plots after root differencing (c: frequency domain differencing)

3.3. Migration

We obtained Swedish migration data from 1851 until 1969 from Mitchell (2003) and from 1970 until 2020 from the website of Statistics Sweden (https://www.scb.se/). The main reasons for selecting Sweden in our analysis of migration dynamics included the existing availability of long historical time series as well as the recent upsurge in academic interest in the subject of migration and flight by Swedish scholars and organizations (Lindfors and Alfonsdotter, 2016, see also Warnqvist, 2018). Furthermore, compared to other European countries, Sweden is known to have had generous asylum laws until 2016. Much earlier already, Sweden became an immigrant country in the 1930s, after initial mass migratory movements toward the United States. This mass migration unfolded throughout different waves between the 1840s and the late 1920s.

Applying the same methods as in the previous subsections, we again find evidence of nonstationarity, albeit somewhat weaker than before. The empirical findings are shown in Figure 8 for the immigration data and in Figure 9 for the emigration data. In both cases, it is safe to assume that the observed nonstationarity is mainly due to structural breaks which separate different waves of migration.

4. Discussion

Carrying out fractional differencing with the help of a truncated version of the power series expansion of the fractional differencing operator is only meaningful when the sample size is very large. For more realistic sample sizes occurring in practice, we therefore propose a simple approximation which is based on a simple ARMA(1,1) model. We focus on the important special case where the fractional differencing parameter is equal to 0.5 (root differencing). This choice allows the use of the new method to distinguish between stationarity and nonstationarity of a given time series. For this purpose, we look at the time series plot, the sample autocorrelations, and the periodogram of the root differenced time series. A trend or long cycles in the plot, positive and slowly decaying sample autocorrelations, and a peak of the periodogram at frequency zero are interpreted as indications of nonstationarity. The new method is applied to long annual time series, the global surface temperature from 1850 to 2020 (HadCRUT5), the UK GDP per capita from 1252 to 2018 (Maddison Project Database), Swedish immigration numbers from 1875 to 2020 and Swedish emigration numbers from 1851 to 2020 (Mitchell, 2003; Statistics Sweden). In each case, we expect a priori that the hypothesis of stationarity is rejected. The temperature series and the GDP series have an upward trend because of global warming and long-term economic growth, respectively. The nonstationarity of the migration series is due to structural breaks which separate different waves of migration. It is reassuring that our empirical analysis of these series indeed produces evidence in favor of nonstationarity. However, this evidence is not always as strong as anticipated, which is particularly sobering in view of the fact that most historical time series are much shorter than the examples investigated in this paper. For example, studies of the long-term properties of macroeconomic time are usually based only on post-war data (see, e.g., Christiano and Eichenbaum, 1990; Hauser et al., 1999). Not helping in this matter would be to increase the frequency from annual to quarterly or monthly, which increases only the number of observations but not the length of the observation period. We conclude that even after

refraining from carrying out a formal statistical test (because of inherent theoretical issues; see Pötscher, 2002) and turning to a more informal approach as described in this paper, it may still be hard to find reliable evidence in favor or against stationarity unless the observation period is extremely long.

REFERENCES

[1] J. Bolt and J. L. van Zanden, Maddison style estimates of the evolution of the world economy. A new 2020 update, Maddison Project Database (2020).

[2] H. Buhaug, Climate not to blame for African civil wars, PNAS 107 (2010), 16477-16482.

[3] M. B. Burke, E. Miguel, S. Satyanath, J. A. Dykema and D. B. Lobell, Warming increases the risk of civil war in Africa, PNAS 106 (2009), 20670-20674.

[4] L. J. Christiano and M. S. Eichenbaum, Unit roots in real GNP: do we know, and do we care?, in Carnegie-Rochester Conference Series on Public Policy, Elsevier (1990), pp. 7–61.

[5] B. Drake, Changes in North Atlantic Oscillation drove population migrations and the collapse of the Western Roman Empire, Scientific Reports 7 (2017), 1227.

[6] C. W. J. Granger and R. Joyeux, An introduction to long-memory time series models and fractional differencing, Journal of Time Series Analysis 1 (1980), 15-29.

[7] GISTEMP Team, GISS Surface Temperature Analysis (GISTEMP), version 4, NASA Goddard Institute for Space Studies (2021).

[8] U. Hassler, Regression of spectral estimators with fractionally integrated time series, Journal of Time Series Analysis 14 (1993), 369-380.

[9] M. A. Hauser, B. M. Pötscher and E. Reschenhofer, Measuring persistence in aggregate output: ARMA models, fractionally integrated ARMA models and nonparametric procedures, Empirical Economics 24 (1999), 243-269.

[10] J. R. M. Hosking, Fractional differencing, Biometrika 68 (1981), 165-176.

[11] S. M. Hsiang, K. C. Meng and M. A. Cane, Civil conflicts are associated with the global climate, Nature 476 (2011), 438–441.

[12] C. P. Kelley, S. Mohtadi, M. A. Cane, R. Seager and Y. Kushnir, Climate change in the fertile crescent and implications of the recent Syrian drought, PNAS 112 (2015), 3241–3246.

[13] N. Lenssen, G. Schmidt, J. Hansen, M. Menne, A. Persin, R. Ruedy and D. Zyss, Improvements in the GISTEMP uncertainty model, Journal of Geophysical Research: Atmospheres 124 (2019), 6307-6326.

[14] R. Lindfors and K. Alfonsdotter, Silent books: A handbook on wordless picture books packed with narrative power, IBBY Sweden (2016), 1-16.

[15] M. K. Mangat and E. Reschenhofer, Frequency-domain evidence for climate change. Econometrics 8 (2020), 28, 1-15.

[16] E. Miguel, S. Satyanath and E. Sergenti, Economic shocks and civil conflict: an instrumental variables approach, Journal of Political Economy 112 (2004), 725-753.

[17] B. R. Mitchell, International Historical Statistics: Europe, 1750-2000, 5th ed., Palgrave Macmillan, New York (2003).

[18] C. P. Morice, J. J. Kennedy, N. A. Rayner, J. P. Winn, E. Hogan, R. E. Killick, R. J. H. Dunn, T. J. Osborn, P. D. Jones and I. R. Simpson, An updated assessment of near-surface temperature change from 1850: the HadCRUT5 dataset, Journal of Geophysical Research (2021).

[19] B. M. Pötscher, Lower risk bounds and properties of confidence sets for ill-posed estimation problems with applications to spectral density and persistence estimation, unit roots, and estimation of long memory parameters, Econometrica 70 (2002), 1035-1065.

[20] R Core Team, R: A language and environment for statistical computing, R Foundation for Statistical Computing, Vienna (2018).

[21] V. A. Reisen, Estimation of the fractional difference parameter in the ARIMA (p, d, q) model using the smoothed periodogram, Journal of Time Series Analysis 15 (1994), 335-350.

[22] E. Reschenhofer, Robust testing for stationarity of global surface temperature, Journal of Applied Statistics 40 (2013), 1349–61.

[23] E. Reschenhofer and M. K. Mangat, Fast computation and practical use of amplitudes at non-Fourier frequencies, Computational Statistics (2021), 1755–1773.

[24] W. Scheidel and S. J. Friesen, The size of the economy and the distribution of income in the Roman Empire, Journal of Roman Studies 99 (2009), 61–91.

[25] A. R. Solow, A call for peace on climate and conflict, Nature 497 (2013), 179–180.

[26] C. Stohr, Trading gains: new estimates of Swiss GDP, 1851-2008, Economic History Working Papers, London School of Economics and Political Science, Economic History Department (2016), 245.

[27] O. M. Theisen, H. Holtermann and H. Buhaug, Climate wars? Assessing the claim that drought breeds conflict, International Security 36 (2011), 79-106.

[28] A. Warnqvist, Family first in homes away from home: Depictions of refugee experiences and flight from war in picturebooks published in Sweden 2014-2018, Bookbird 56 (2018), 4, 60-71.