International Journal of Forecasting 39 (2023) 922–937

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

The RWDAR model: A novel state-space approach to forecasting

Giacomo Sbrana [a], Andrea Silvestrini [b,*]

[a] NEOMA Business School, 59 Rue Pierre Taittinger, 51100 Reims, France
[b] Bank of Italy, DG for Economics, Statistics and Research, Via Nazionale 91, 00184 Roma, Italy

ARTICLE INFO

Keywords:
Theta method
Kalman gain
Forecasting
M3 and M4 competitions
Approximate maximum likelihood estimation

ABSTRACT

This paper introduces the Random Walk with Drift plus AutoRegressive model (RWDAR) for time-series forecasting. Owing to the presence of a random walk plus drift term, this model shares some similarities with the Theta model of Assimakopoulos and Nikolopoulos (2000). However, the addition of a first-order autoregressive term in the state equation provides additional adaptability and flexibility. Indeed, it is shown that RWDAR tends to outperform the Theta model when forecasting both stationary and nearly non-stationary time series. This paper also proposes a simple estimation method for the RWDAR model based on the solution of the algebraic Riccati equation for the prediction error covariance of the state vector. Simulation results show that this estimator performs as well as the standard Kalman filter approach. Finally, using yearly data from the M3 and M4 competition datasets, it is found that RWDAR outperforms traditional forecasting methods.

1. Introduction

This paper presents a new time-series model called Random Walk with Drift plus AutoRegressive (RWDAR), in which multiple errors drive the whole system's dynamics. In its state-space representation, RWDAR features two separate components: a random walk with drift for the level or trend of the series, and a first-order autoregression for the residual components. As such, the RWDAR model resembles the Theta model of Assimakopoulos and Nikolopoulos (2000), which performed remarkably well in the M3 competition (Makridakis & Hibon, 2000). Besides introducing the RWDAR model, this paper proposes a simple method for estimating its parameters that relies on the analytical solution to the algebraic Riccati equation for the covariance matrix of the state vector's prediction error (Sbrana & Silvestrini, 2020). This analytical solution greatly simplifies both the Kalman filter recursions and the evaluation of the model's likelihood function. This estimation method is similar to that successfully introduced by Sbrana and Silvestrini (2019) for estimating the random switching exponential smoothing model. Using Monte Carlo simulation, the finite-sample properties of the newly proposed estimator are compared with those of the standard likelihood maximization procedure based on the Kalman filter. The results show that the performance of both methods is quite similar, irrespective of the sample size.

The closely related Theta model introduced by Assimakopoulos and Nikolopoulos (2000) is a forecasting method based on a decomposition of the second differences of a time series into short- and long-term components. More precisely, the Theta method modifies the local curvature of the series through the theta (θ) coefficients, which are applied to the second differences of the data. Depending on the values of the theta coefficients, the local curvature of the data is either decreased (i.e., the series is deflated) or increased (i.e., the series is dilated), resulting in a magnification

* Corresponding author.
E-mail addresses: Giacomo.SBRANA@neoma-bs.fr (G. Sbrana), andrea.silvestrini@bancaditalia.it (A. Silvestrini).

https://doi.org/10.1016/j.ijforecast.2022.03.003

of the long-term or short-term component, respectively; see Nikolopoulos and Thomakos (2019).[1] More recently, assuming a stochastic trend in the data-generating process, Thomakos and Nikolopoulos (2014) showed that the equally weighted theta-lines forecast corresponds to the optimal (i.e., the minimum mean squared error) forecast of the same data-generating process. Thus, the unit root model underpins the forecasts produced by the Theta method. Thomakos and Nikolopoulos (2015) proposed a multivariate extension of the original univariate Theta method and found that the bivariate Theta method performs very well and in most evaluations outperforms the univariate version of the method. Other important contributions on the Theta method include Fiorucci, Pellegrini, Louzada, Petropoulos, and Koehler (2016), Nikolopoulos, Assimakopoulos, Bougioukos, Litsa, and Petropoulos (2011), Spiliotis, Assimakopoulos, and Makridakis (2020), Spiliotis, Assimakopoulos, and Nikolopoulos (2019), among others.

Building on earlier work by Assimakopoulos and Nikolopoulos (2000), Hyndman and Billah (2003) worked out the underlying forecasting function of the Theta method employed by Assimakopoulos and Nikolopoulos (2000) in the M3 competition for monthly data, which is a simple exponential smoothing with drift, in which the drift parameter is set to half the slope of the fitted trend line through the original time series. Using a probabilistic framework, Hyndman and Billah (2003) also showed that it is possible to derive prediction intervals and likelihood-based estimation of the parameters. However, Nikolopoulos and Thomakos (2019) pointed out that a clear link between the basic Theta method and simple exponential smoothing with drift can only be established when focusing on the specific Theta formulation used by Assimakopoulos and Nikolopoulos (2000) in the M3 competition for monthly data, and not in general.[2]

From the above discussion, it appears clear that RWDAR shares some key similarities with Theta. Despite that, it is a different model. Specifically, the presence of an autoregressive component in the measurement equation, which is absent in the Theta model, provides additional flexibility well suited for time series with a trend plus cycle dynamics, or for stationary and nearly non-stationary data. This is evident from the Monte Carlo experiment reported in Section 4.1. Moreover, this finding is reinforced by an empirical application to forecasting US output growth and inflation (Section 4.2). Another difference is related to the fact that the RWDAR model is cast into a multiple-error-source state-space form, and therefore requires the use of the Kalman filter, while in the probabilistic interpretation of the Theta method provided by Hyndman and Billah (2003), the observation and state equations are driven by the same innovation.

The RWDAR model can be also related to the large strand of literature on the decomposition of time series into permanent and transitory components; see among others Beveridge and Nelson (1981), Canova (1998), Harvey (1985), Hodrick and Prescott (1997). This literature is clearly interlinked with the macroeconomic debate on the relative importance of permanent and transitory components in GDP growth (Nelson & Plosser, 1982; Watson, 1986). In the RWDAR specification, the trend of the time series is modeled as a random walk with drift, as, for instance, in the representation suggested by Watson (1986) and, more recently, by Oh, Zivot, and Creal (2008), to decompose the series into a trend and a cyclical component. The RWDAR model is used in this paper only for forecasting purposes, yet it is useful to interpret it in a broader perspective, also in light of efforts already made in the econometric literature to model trends.

The out-of-sample forecasting accuracy of the RWDAR model is assessed using simulated data from a Monte Carlo experiment and two forecast applications based on the M3 and M4 competition datasets. A set of relevant competing methods (ETS, Theta, auto-ARIMA, and a combination of these three methods, termed "EAT") is employed and different time frequencies are considered (i.e., yearly and quarterly). Overall, results show that RWDAR delivers more accurate forecasts than several traditional statistical models, and scores well even when compared to the Theta model, which is nowadays in the forecast practitioner's toolkit. Therefore, the RWDAR model is potentially of interest to forecast practitioners. Interestingly, additional gains in forecast accuracy can be achieved by combining RWDAR predictions with forecasts from the other individual models under study (ETS, Theta, and auto-ARIMA).

The rest of the paper proceeds as follows. Section 2 describes the RWDAR model and the newly proposed estimation method. Section 3 introduces a simulation exercise designed to study the finite-sample properties of the new estimator. Section 4 illustrates another Monte Carlo experiment designed for comparing the forecast accuracy of the RWDAR model relative to that of simple exponential smoothing with drift. It also examines the relative forecasting performance of RWDAR against simple exponential smoothing with drift using selected time series (US GDP and CPI). Section 5 presents the out-of-sample forecasting exercise and the results obtained on the M3 competition dataset. Section 6 discusses the results achieved using the M4 dataset, which extends the M3 competition in terms of the number of time series included and forecasting methods considered. Section 7 investigates whether any gains in forecast accuracy can be made by combining RWDAR predictions with forecasts from ETS, Theta, and auto-ARIMA. Section 8 draws conclusions and proposes suggestions for future research. The appendix contains the complete proofs of the propositions presented in the paper together with the R code to be used for estimating the RWDAR model.

2. The RWDAR model

Consider the following Random Walk with Drift plus AutoRegressive (RWDAR) model:

$$y_t = \lambda_{t-1} + \beta_{t-1}, \qquad t = 1, 2, \ldots, n,$$

[1] For a step-by-step description of the Theta method, see Petropoulos and Nikolopoulos (2017).

[2] A direct link exists only when the decomposition is undertaken assuming a theta line with $\Theta = 0$ and a theta line with $\Theta = 2$, as discussed in Nikolopoulos and Thomakos (2019), pp. 11–12.

$$\lambda_t = \tau + \lambda_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{NID}(0, \sigma_\epsilon^2), \qquad (1)$$
$$\beta_t = \phi\beta_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{NID}(0, \sigma_\eta^2),$$

where $\mathcal{NID}$ denotes normally and independently distributed innovations. There is no disturbance term on the measurement equation. It is also assumed that $E\left(\epsilon_{t-j}\eta_{t-g}\right) = 0 \;\; \forall j, g$ and that $0 \le \phi < 1$. We exclude the case when $\phi = 1$ consisting of the sum of two random walks.[3] The RWDAR is a flexible parameterization which encompasses several alternative specifications widely employed in empirical applications. For instance, when $\phi = 0$, model (1) is equivalent to the so-called random walk with drift plus noise. In contrast, when $\sigma_\epsilon^2 = 0$, model (1) collapses to an AR(1) plus constant. Therefore, although the RWDAR model is originally designed for non-stationary series, it is also suitable for stationary processes. We also note that, unlike the single-error-source state-space approach used by Hyndman and Billah (2003), two uncorrelated noises drive the system in (1). Hence, this formulation is also referred to as a multiple-error-source scheme. As a consequence, RWDAR belongs to the class of unobserved-components time series models, as extensively discussed in Harvey (1986), Durbin and Koopman (2012), and Proietti (2021).

The state-space representation of RWDAR can be written as follows[4]:

$$y_t = \mathbf{z}\boldsymbol{\mu}_{t-1}$$
$$\boldsymbol{\mu}_t = \boldsymbol{\delta} + \mathbf{T}\boldsymbol{\mu}_{t-1} + \mathbf{R}\mathbf{u}_t, \qquad (2)$$

where, by definition,

- $\mathbf{z} = \begin{bmatrix} 1 & 1 \end{bmatrix}$
- $\mathbf{T} = \begin{bmatrix} 1 & 0 \\ 0 & \phi \end{bmatrix}$
- $\boldsymbol{\delta} = \begin{bmatrix} \tau \\ 0 \end{bmatrix}$
- $\boldsymbol{\mu}_t = \begin{bmatrix} \lambda_t \\ \beta_t \end{bmatrix}$
- $\mathbf{R} = \begin{bmatrix} \sigma_\epsilon & 0 \\ 0 & \sigma_\eta \end{bmatrix}$
- $\mathbf{u}_t = \begin{bmatrix} u_{1t} \\ u_{2t} \end{bmatrix}$, with $\mathbb{E}(\mathbf{u}_t) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and $\mathbb{E}(\mathbf{u}_t\mathbf{u}_t') = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$.

Moreover, the corresponding Kalman filter recursions are (Harvey, 1989):

$$f_t = \mathbf{z}\mathbf{P}_{t-1}\mathbf{z}',$$
$$\mathbf{k}_t = \mathbf{T}\mathbf{P}_{t-1}\mathbf{z}'f_t^{-1},$$
$$v_t = y_t - \mathbf{z}\,\mathbf{m}_{t-1}, \qquad (3)$$
$$\mathbf{m}_t = \boldsymbol{\delta} + \mathbf{T}\mathbf{m}_{t-1} + \mathbf{k}_t v_t,$$
$$\mathbf{P}_t = \mathbf{T}\mathbf{P}_{t-1}\mathbf{T}' - \mathbf{T}\mathbf{P}_{t-1}\mathbf{z}'f_t^{-1}\mathbf{z}\mathbf{P}_{t-1}\mathbf{T}' + \mathbf{R}\mathbf{R}', \quad t = 2, \ldots, n$$

where the prediction $\mathbf{m}_t = \mathbb{E}(\boldsymbol{\mu}_t|Y_{t-1}) = \begin{bmatrix} l_t \\ b_t \end{bmatrix}$ is the expected value of the state vector given the information provided by data up to time $t-1$, that is, $Y_{t-1} = \{y_1, y_2, \ldots, y_{t-1}\}$, and $\mathbf{P}_t = \mathbb{E}((\boldsymbol{\mu}_t - \mathbf{m}_t)(\boldsymbol{\mu}_t - \mathbf{m}_t)')$ is the prediction error covariance matrix.

Since $\mathbf{T}$, $\mathbf{R}$, and $\mathbf{z}$ are constant, model (1) is time-invariant. If there exists a time-invariant covariance matrix such that $\mathbf{P}_t = \mathbf{P}_{t-1} = \bar{\mathbf{P}}$, the Kalman filter has a steady-state solution to which it converges exponentially fast. It can be shown (Harvey, 1989) that the prediction error's covariance matrix converges to the following solution of the well-known algebraic Riccati equation:

$$\bar{\mathbf{P}} = \mathbf{T}\bar{\mathbf{P}}\mathbf{T}' - \mathbf{T}\bar{\mathbf{P}}\mathbf{z}'f^{-1}\mathbf{z}\bar{\mathbf{P}}\mathbf{T}' + \mathbf{R}\mathbf{R}'$$

As a consequence, at the steady-state, one does not need to store the expressions for $\boldsymbol{P}_t$, f_t and $\boldsymbol{k}_t$ at each point in time, and the only two recursions of the Kalman filter that need to be executed are the innovation equation (i.e., v_t) and the prediction equation, or state vector (i.e., $\mathbf{m}_t$). See Harvey (1989), Sections 3.3.3 and 4.1.3 (expressions 4.1.38).

In Proposition 1, the algebraic mapping between the parameters of model (1) and the elements of $\bar{\mathbf{P}}$ is provided. This link is crucial for deriving the newly proposed estimator in Proposition 2.

Proposition 1. *Given $0 \le \phi < 1$ and the signal-to-noise ratio $q = \frac{\sigma_\epsilon^2}{\sigma_\eta^2} > 0$, there exists a unique positive-definite solution for $\bar{\mathbf{P}} = \sigma_\eta^2 \begin{bmatrix} p_1 & p_2 \\ p_2 & p_3 \end{bmatrix}$. This solution is:*

$$p_1 = \frac{-3q\phi + \sqrt{q}\sqrt{q(\phi+1)^2 + 4} + q}{2 - 2\phi} \qquad (4)$$

$$p_2 = -\frac{\phi\left(q\phi - \sqrt{q}\sqrt{q(\phi+1)^2+4} + q\right)}{2(\phi-1)} \qquad (5)$$

$$p_3 = \frac{\phi\left(q(\phi+1)\phi - \sqrt{q}\phi\sqrt{q(\phi+1)^2+4} + 2\right) - 2}{2(\phi-1)} \qquad (6)$$

Proof. See Appendix A. □

This result has important consequences. Indeed, the next proposition puts forward an estimation procedure that simplifies the maximization of the log-likelihood function.[5]

Proposition 2. *Given $0 \le \phi < 1$ and $q = \frac{\sigma_\epsilon^2}{\sigma_\eta^2} > 0$, consider the following two recursions:*

$$v_t = y_t - l_{t-1} - b_{t-1} \qquad (7)$$

and

$$\begin{bmatrix} l_t \\ b_t \end{bmatrix} = \begin{bmatrix} \tau \\ 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & \phi \end{bmatrix}\begin{bmatrix} l_{t-1} \\ b_{t-1} \end{bmatrix} + \begin{bmatrix} k_1 \\ k_2 \end{bmatrix} v_t \qquad (8)$$

---

3 This point was suggested by an anonymous reviewer, to whom we are grateful.

4 The following convention is used: $\boldsymbol{X}$ is a matrix and $\boldsymbol{X}'$ is its transpose; $\boldsymbol{x}$ is a vector such that $\boldsymbol{x}'$ is its transpose. A lower-case letter, such as x, represents a scalar.

5 A similar estimation approach was recently used by Sbrana and Silvestrini (2021).

with

$$\begin{bmatrix} k_1 \\ k_2 \end{bmatrix} = \begin{bmatrix} -\frac{2\sqrt{q}}{\sqrt{q}(\phi-1)-\sqrt{q(\phi+1)^2+4}} \\ \frac{\phi\left(\phi q+q-\sqrt{q(\phi+1)^2+4}\sqrt{q}+2\right)}{2q\phi+2} \end{bmatrix} \tag{9}$$

It follows that the log-likelihood function of model (1) can be maximized by choosing the set of parameters (ϕ, q, and τ) that minimizes the sum of squared innovations:

$$argmin_{(\phi\ ;\ q\ ;\tau)}\left(\sum_{t=2}^{n} {v_t}^2\right).$$

Proof. See Appendix B. □

A few remarks are in order. The Kalman filter is a recursive algorithm. At the start of Kalman filtering, the state vector in (8) has to be initialized before running the estimation procedure. For an AR(1) process the initialization is $b_1 = 0$, while for the random walk the standard diffuse initialization assumption is used: $\begin{bmatrix} l_1 \\ b_1 \end{bmatrix} = \begin{bmatrix} y_1 \\ 0 \end{bmatrix}$.

Just after the state vector is initialized, for $t = 2, 3, \ldots, n$ one needs to run only (7) and (8) using (9), choosing the set of parameters $\left(\hat{\phi}\ ;\ \hat{\tau}\ ;\ \hat{q}\right)$ that minimize the sum of squared prediction errors (see the discussion in Appendix B). This procedure is termed "approximate maximum likelihood estimation" (Harvey, 1989) and was recently used by Sbrana and Silvestrini (2019), to which we refer the reader for additional details.

These remarks, together with results in Propositions 1 and 2, encapsulate all one needs to know to estimate the RWDAR model. The R code to implement the estimator as in Proposition 2 is provided in Appendix C.

3. Monte Carlo simulation: Analyzing the finite-sample behavior of the RWDAR estimator

This section presents some Monte Carlo evidence on the finite-sample behavior of the RWDAR's approximate maximum likelihood estimator developed in Section 2. The Monte Carlo experiment aims at comparing its performance to that of the standard Kalman filter maximum likelihood estimator, using samples of different lengths. The simulation is implemented using R software, and the developed computer code is available from the authors upon request.

The data-generating process (DGP) is the RWDAR model in (1). A large number of replications (10,000) are used in order to produce accurate results. For each replication, parameters are randomly generated, allowing simulation results to be general rather than case-specific. The coefficients are drawn as follows. The ϵ_t noise is independently sampled from a Gaussian density function with zero mean and variance drawn from *Uniform*$(0.1, 0.5)$, where *Uniform*(a, b) represents a uniform distribution defined over the interval a and b. The η_t noise is instead sampled independently from a Gaussian density with zero mean and variance drawn from *Uniform*$(1.5, 2.5)$. In principle, this choice yields a signal-to-noise ratio $q = \frac{\sigma_\epsilon^2}{\sigma_\eta^2}$ that is positive and below one. The first-order autoregressive parameter ϕ is drawn from a uniform distribution over the interval $[0.65, 0.95]$, thus producing some persistence in the DGP, while the drift parameter τ is sampled from a uniform distribution over the interval $[0.1, 1]$. The number of observations in each Monte Carlo sample is $n = 30; 90; 150; 300$. This allows one to assess how the bias of the estimator (accuracy) and its variance (precision) are affected by the number of observations available in the sample.

Fig. 1 displays the Monte Carlo bias distribution of the estimates of ϕ, $q = \frac{\sigma_\epsilon^2}{\sigma_\eta^2}$ and τ across all considered sample sizes. Each chart presents boxplots for the distributions of the parameter estimates achieved by implementing the estimation method based on the steady-state Kalman filter as in Proposition 2 (steady state, SS) and the alternative estimation procedure based on the standard Kalman filter (KF) in . The KF method, which represents the standard likelihood maximization procedure, is carried out without concentrating the likelihood. Therefore, the two innovation variances σ_ϵ^2 and σ_η^2 are estimated separately.

Fig. 1 reveals that the ϕ parameter and the drift term (τ) are very precisely estimated, except in small samples ($n = 30$), with a bias converging monotonically to zero as the sample size increases. Furthermore, no significant differences emerge between the two estimation methods, with boxplots showing similar medians and interquartile ranges. When considering the signal-to-noise ratio, the bias also converges to zero, but at a slower pace. However, the median bias is close to zero except when $n = 30$, after which the interquartile range shrinks, indicating concentration of the bias towards zero. Still focusing on the signal-to-noise ratio q, the boxplots of the two estimators are almost overlapping, barring $n = 90{,}150$, for which the distributions of the SS estimates are more dispersed compared to the standard KF. However, based on the available evidence, it can be concluded that the two likelihood procedures yield similar results in terms of bias.

Fig. 2 turns to the variance distribution for the three parameters. It can be clearly seen that the variance distributions of the ϕ and τ estimates collapse to zero very quickly, with no remarkable differences between the two estimation methods. The same holds true for the signal-to-noise parameter, even if in this case the convergence seems to be slower, but still monotonic to zero. An improvement on the precision of the estimates is achieved when $n = 90$ and for larger sample sizes.

All in all, these simulation results highlight that the approximate maximum likelihood estimation method presented in Proposition 2 yields very similar results compared to the standard likelihood maximization procedure. As a result, this evidence confirms that using the steady-state Kalman filter just after initializing the state vector does not affect the estimation performance.

Finally, Table 1 reports the average times for executing both the standard maximum likelihood (Kalman filter) and the approximate maximum likelihood as in

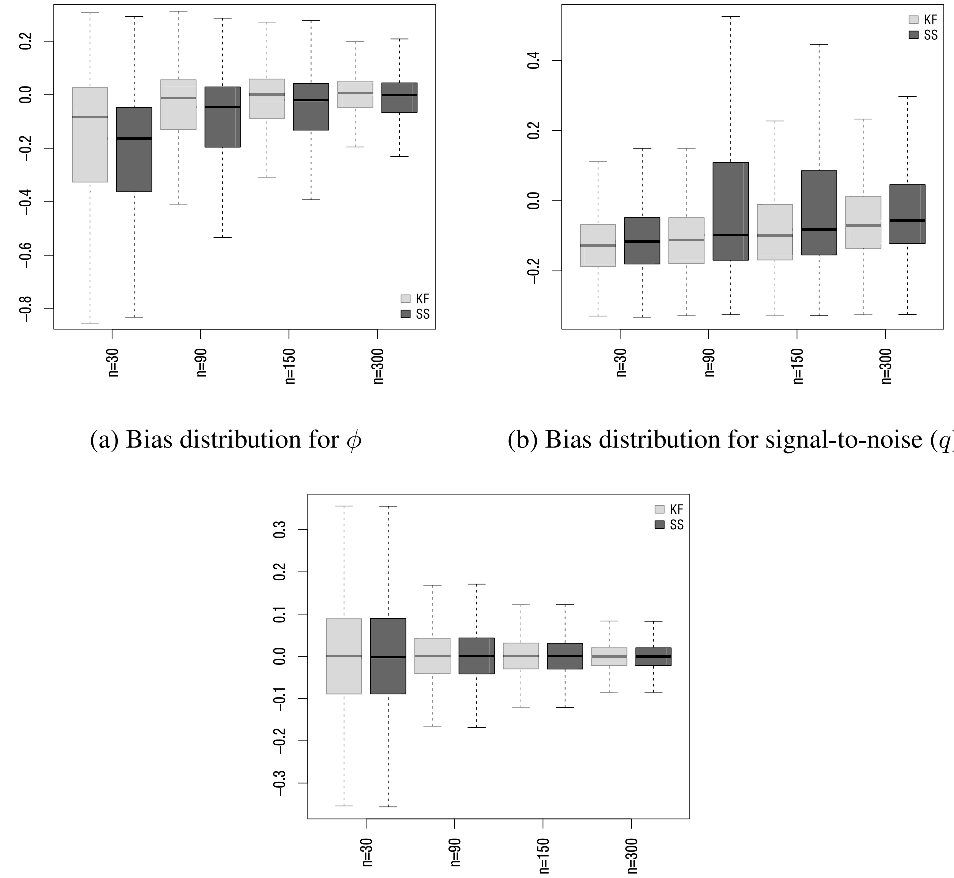

(a) Bias distribution for ϕ

(b) Bias distribution for signal-to-noise (q)

(c) Bias distribution for τ

Fig. 1. Monte Carlo bias of the RWDAR estimators. Notes. This figure shows for a generic parameter, say θ, the distribution of $\hat{\theta}_i - \theta$, where $\hat{\theta}$ is the estimated value, and $i = 1, 2, \ldots, 10,000$ is the generic replication in the Monte Carlo simulation.

Proposition 2. The computational advantages of using our approach (approximate ML) are clear for all sample sizes considered. Indeed, the approximate maximum likelihood method is on average more than 100 times faster than the standard Kalman filter, as can be seen from the "Ratio" column.

Table 1
Execution times (in seconds).

Sample size	Standard ML	Approximate ML	Ratio
n=30	0.793	0.007	117.0
n=90	2.079	0.017	121.1
n=150	3.453	0.027	127.0
n=300	6.314	0.065	122.2

4. Understanding RWDAR with the Theta: Comparing forecasting performance

4.1. Monte Carlo evidence

This section illustrates the results of a Monte Carlo simulation designed for comparing the forecasting performance of RWDAR relative to simple exponential smoothing with drift, referred to as the Theta–forecast adaptation model, as in Hyndman and Billah (2003) (or HB Theta–forecast adaptation model).[6] The same framework already adopted in the previous section is used. The simulation is implemented using R software and the R forecast package (Hyndman et al., 2020; Hyndman & Khandakar, 2008). The computer code is available from the authors upon request.

The first DGP is a random walk with drift plus autoregressive term in a multiple-error-source framework, as in (1). For each replication, parameters are randomly generated. More specifically, the ϵ_t noise is independently sampled from a Gaussian distribution with zero mean and variance drawn from *Uniform*$(0.10, 5.10)$. The η_t noise is instead independently sampled from Gaussian density with zero mean and variance drawn from *Uniform*$(5, 10)$, yielding a signal-to-noise ratio $q = \frac{\sigma_\epsilon^2}{\sigma_\eta^2}$ that is positive and usually smaller than one. The drift parameter τ is sampled

[6] As stated above, this equivalence is valid only in the specific case in which the Theta method is applied, assuming a theta line with $\Theta = 0$ and a theta line with $\Theta = 2$. Forecasts from simple exponential smoothing with drift are produced by using the "thetaf" function of the R forecast package.

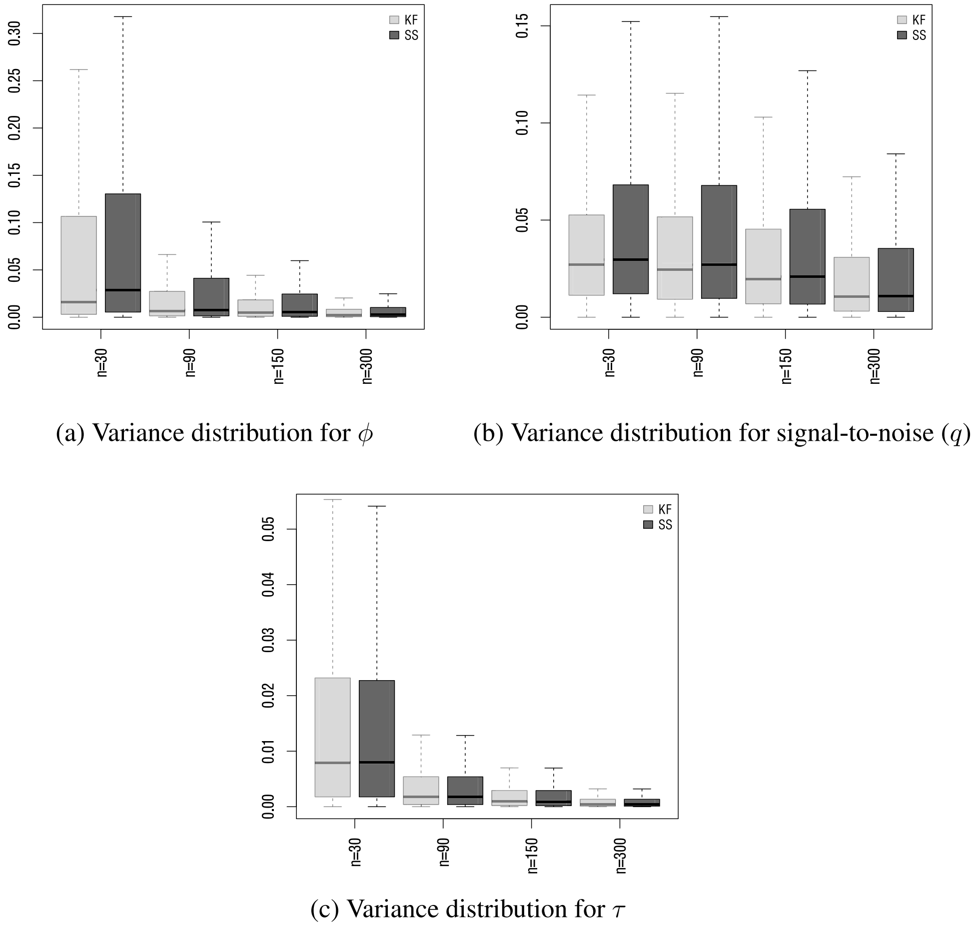

(a) Variance distribution for ϕ

(b) Variance distribution for signal-to-noise (q)

(c) Variance distribution for τ

Fig. 2. Monte Carlo variance of the RWDAR estimators. Notes. This figure shows for a generic parameter, say θ, the distribution of $(\hat{\theta}_i - \theta)^2$, where $\hat{\theta}$ is the estimated value and $i = 1, 2, \ldots, 10,000$ is the generic replication in the Monte Carlo simulation.

from a uniform distribution over the interval [0.01, 0.11], while the first-order autoregressive parameter ϕ assumes the following values: 0.1, 0.3, 0.5, 0.7, and 0.9, so that we can appreciate the influence of different intensities of the persistence parameter. The number of replications is set to 5000. The number of observations in each sample is 300. For each generated sample, the RWDAR and the Theta models are first estimated in-sample, and then one- to six-step-ahead forecasts are produced out-of-sample.

The forecasting results are presented in Table 2 using the root mean squared scaled error (RMSSE); see Hyndman and Koehler (2006). Specifically, each cell in Table 2 contains the ratio of the Theta's RMSSE relative to that of RWDAR. Results in the table indicate that, when the AR parameter $\phi = 0.1$ and $\phi = 0.3$, both models have almost identical RMSSEs, no matter the forecast horizon considered. This is expected since the AR has very little persistence (mimicking white noise) and the model tends to a random walk with drift plus noise (see Hyndman & Billah, 2003). When increasing the process persistence, RWDAR tends to perform slightly better than the Theta model, particularly when $\phi = 0.5$ and $\phi = 0.7$. This is also expected, since, contrary to the Theta model, RWDAR is able to detect the additional AR component. Overall, the RWDAR and Theta models have rather similar forecasting performance for different values of the persistence parameter, regardless the forecast horizon used, with no model clearly outperforming its competitor. Thus, in such a non-stationary setting, these simulation results appear to be in line with previous findings of Thomakos and Nikolopoulos (2014), and shed further light on the AR parameter's influence on the forecast accuracy of the two models under comparison.

So far we have considered a DGP with a stochastic trend in the form of a unit root. It is also interesting to examine the case of stationary time series being close to non-stationary (nearly unit-root) series. Therefore, we consider a second DGP consisting of an AR(1) plus constant, where the AR gets close to a unit-root process:

$$y_t = \lambda_{t-1} + \beta_{t-1},$$
$$\lambda_t = \tau,$$
$$\beta_t = \phi\beta_{t-1} + \eta_t$$

The η_t noise is sampled as specified above, with the first DGP. The autoregressive parameter ϕ assumes three values (0.85, 0.9, and 0.95), pointing to an increasingly high degree of persistence, which appears to be realistic in many economic and financial time series.

Table 2
Out-of-sample root mean squared scaled error (RMSSE) of Theta vs. RWDAR (non-stationary DGP).

No. steps ahead	$\phi = 0.1$	$\phi = 0.3$	$\phi = 0.5$	$\phi = 0.7$	$\phi = 0.9$
$h = 1$	0.998	1.005	1.018	1.022	1.011
$h = 2$	0.998	1.006	1.021	1.026	1.013
$h = 3$	0.998	1.006	1.023	1.029	1.014
$h = 4$	0.997	1.005	1.022	1.029	1.014
$h = 5$	0.996	1.004	1.021	1.027	1.011
$h = 6$	0.995	1.004	1.020	1.023	1.006
Average	0.997	1.005	1.021	1.026	1.012

Notes. For $h = 1, 2, 3, 4, 5, 6$ (number of steps ahead) and five different values of the AR parameter ϕ, each cell contains the ratio of the Theta's RMSSE relative to the RWDAR's RMSSE; values greater than one, for instance, indicate that the RMSSE of Theta is larger than that of RWDAR.

Table 3
Out-of-sample root mean squared scaled error (RMSSE) of Theta vs. RWDAR (stationary DGP).

No. steps ahead	$\phi = 0.85$	$\phi = 0.9$	$\phi = 0.95$
$h = 1$	1.055	1.033	1.013
$h = 2$	1.072	1.043	1.017
$h = 3$	1.082	1.050	1.019
$h = 4$	1.087	1.053	1.019
$h = 5$	1.089	1.053	1.017
$h = 6$	1.086	1.050	1.014
Average	1.078	1.047	1.017

Notes. For $h = 1, 2, 3, 4, 5, 6$ (number of steps ahead) and two different values of the AR parameter ϕ, each cell contains the ratio of the Theta's RMSSE relative to the RWDAR's RMSSE; values greater than one, for instance, indicate that the RMSSE of Theta is larger than that of RWDAR.

The results are presented in Table 3. The message here is clear: the less the persistence of the AR process, the better the RWDAR performance. However, the edge of RWDAR over Theta vanishes when the AR parameter approaches 1. This is again expected, and perfectly coherent with previous results from several other studies (Nikolopoulos & Thomakos, 2019; Thomakos & Nikolopoulos, 2014). In summary, when stationary series are considered, the RWDAR outperforms the Theta model. On the other hand, when non-stationarity is present, as often happens in business, economics, and finance, the forecasting performance of the two models become very similar.

4.2. Selected time series

Having examined the forecasting performance of the RWDAR model with simulated data, this section further investigates this issue by using some selected macroeconomic time series. Specifically, the focus is on US real gross domestic product (billions of US dollars, quarterly frequency, seasonally adjusted annual rate) and the US Consumer Price Index for all urban consumers (index, 1982–1984=100, monthly frequency, seasonally adjusted), two very important macroeconomic variables for the United States. GDP data are available from 1947:Q1 through 2021:Q3 (299 observations), while CPI is available from 1947:01 through 2021:11 (899 observations). All data are obtained from the Federal Reserve Bank of Saint Louis database (FRED).

For each of these two series, four different transformations of the data are considered: the levels of the data, their logarithms, the first differences of the logarithms, and the seasonal differences of the logarithms (lag-4 differences for quarterly GDP and lag-12 differences for monthly CPI). As in the previous section, RWDAR is compared to simple exponential smoothing with drift, referred to as Theta–forecast adaptation model (Hyndman & Billah, 2003). These two models are first estimated in-sample, and then forecasts are produced out-of-sample. For both series, the sample is split into in-sample and out-of-sample portions, by letting two-thirds of the data points in the in-sample and the remaining third of the data values in the out-of-sample. In order to appreciate whether the choice of the estimation window size affects the model's forecast accuracy, both rolling- and recursive-window estimation schemes are considered in the analysis. One-step-ahead through four-step-ahead predictions are computed to assess whether the models' forecast accuracy changes at different forecast horizons.

Results are presented in Table 4. For simplicity, as in the previous tables, for each series (GDP and CPI), forecast transformation (levels, logs, y-o-y growth rate, and q-o-q/m-o-m growth rate), forecast scheme (recursive and rolling), and forecast horizon ($h = 1, 2, 3, 4$), the ratio of the RMSSE of Theta relative to that of RWDAR is reported in each cell. Thus, values larger than one indicate that the RMSSE of Theta is higher than the RMSSE of RWDAR.

The upper part of Table 4 refers to US real GDP, whose plots after the four different transformations have been applied are presented in Fig. 3 (upper charts). By focusing on the GDP in levels, it appears that the RWDAR model significantly outperforms the Theta, and improves on this later for all forecast horizons considered, no matter the scheme (recursive or rolling) used. The estimated RWDAR parameters are such that the signal-to-noise ratio approaches 0, and the values of the AR parameters are close to 1. That is, the RWDAR combines a persistent AR process with the addition of deterministic drift. The inclusion of the AR term allows for tracking the slowly changing cyclical component of the GDP in levels, and this delivers more accurate forecasts compared to the Theta. Most importantly, since the first differences of the GDP in levels have a unit root, the drift of the Theta model is not appropriate and this explains its results.

A different picture emerges when considering the GDP in log levels. In this case, the Theta model outperforms the RWDAR by a large extent, no matter the horizon and estimation scheme. A visual inspection of the log of GDP shows that it follows a random walk with a slowly changing growth rate. Indeed, after taking first differences, it turns out that the growth rate of GDP fluctuates very little around its average, which is nicely captured by the drift term of the Theta model. On the contrary, the estimated signal-to-noise ratio of the RWDAR model is in most cases very volatile (ranging between 0 and 15), while the estimated AR parameter is close to 1. The volatility of the signal-to-noise ratio clearly indicates that the RWDAR model fails to identify whether the random

Table 4
Out-of-sample root mean squared scaled error (RMSSE) of Theta vs. RWDAR (US real GDP and CPI).

No. steps ahead	GDP in levels		Log of GDP	
	Recursive	Rolling	Recursive	Rolling
$h = 1$	1.168	1.192	0.823	0.819
$h = 2$	1.188	1.220	0.773	0.769
$h = 3$	1.201	1.242	0.744	0.739
$h = 4$	1.206	1.252	0.714	0.707
	GDP growth (y-o-y rate)		GDP growth (q-o-q rate)	
	Recursive	Rolling	Recursive	Rolling
$h = 1$	1.038	1.006	1.048	1.026
$h = 2$	1.040	0.996	1.036	1.023
$h = 3$	1.059	1.005	1.054	1.029
$h = 4$	1.078	1.008	1.064	1.023
	CPI in levels		Log of CPI	
	Recursive	Rolling	Recursive	Rolling
$h = 1$	1.101	1.109	0.873	0.830
$h = 2$	1.132	1.141	0.860	0.810
$h = 3$	1.158	1.175	0.834	0.776
$h = 4$	1.179	1.200	0.808	0.743
	CPI growth (y-o-y rate)		CPI growth (m-o-m rate)	
	Recursive	Rolling	Recursive	Rolling
$h = 1$	1.018	1.016	1.073	1.075
$h = 2$	1.038	1.034	1.052	1.054
$h = 3$	1.058	1.053	1.043	1.041
$h = 4$	1.072	1.067	1.039	1.037

Notes. For $h = 1, 2, 3, 4$ (number of steps ahead), each cell contains the ratio of the Theta's RMSSE relative to the RWDAR's RMSSE; values greater than one, for instance, indicate that the RMSSE of Theta is larger than that of RWDAR.

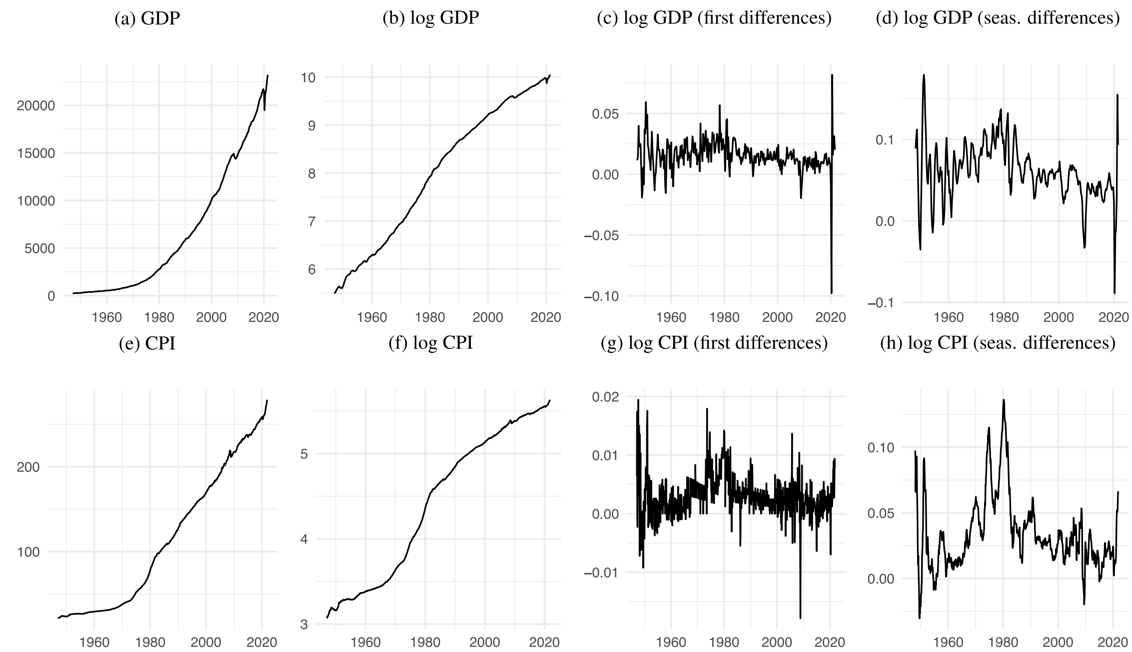

Fig. 3. Time series plots of US GDP in real terms (1947:Q1–2021:Q3) and of US CPI (1947:01–2021:11, 1982–1984=100): four different transformations.

walk component dominates the AR term or vice versa. This has a detrimental impact on its forecast accuracy.

Focusing on the GDP y-o-y growth rate (i.e., the growth rate over the previous year), the RWDAR model features an estimated signal-to-noise ratio close to zero and an estimated AR parameter fluctuating around 0.9. Visually, the y-o-y growth rate series looks like a persistent AR process plus a constant term. In this context, the RWDAR model performs as well as Theta in forecasting, even if it is slightly more accurate than Theta when recursive

Table 5
Types of M3 series per frequency and domain.

Frequency	Micro	Industry	Macro	Finance	Demographic	Other	Total
Yearly	146	102	83	58	245	11	645
Quarterly	204	83	336	76	57	0	756
Monthly	474	334	312	145	111	52	1428
Other	4	0	0	29	0	141	174
Total	828	519	731	308	413	204	3003

forecasts are produced. The results change when examining the GDP q-o-q growth rate (i.e., the growth rate over the previous quarter). This transformation leads to a stationary time series, whose pattern is correctly identified by the RWDAR model, with an estimated AR parameter fluctuating around 0.4. In this case, the RWDAR model delivers more accurate forecasts than Theta, both for rolling and recursive schemes, and particularly for short forecast horizons. This is in line with the Monte Carlo findings presented in Section 4.1.

The bottom part of Table 4 refers to US CPI. The time series plots of the four different transformations are presented in Fig. 3 (lower charts). Overall, the results are broadly in line with those obtained when working with real GDP. The RWDAR model is more accurate at forecasting than Theta when considering the series in levels, while the opposite holds true when the log transformation is used. When focusing on seasonal growth rates, the RWDAR forecasts lead on average to lower RMSSE values. These results are robust across forecast horizons and estimation schemes considered.

5. Application to the M3 competition

This section examines the out-of-sample forecast performance of the RWDAR model using the M3 competition dataset (Makridakis & Hibon, 2000). In this database there are 3003 time series belonging to the following domains: micro, industry, macro, finance, demographic, other sectors. The time series have different frequencies (yearly, quarterly, monthly, and other); see Table 5, taken from Makridakis and Hibon (2000).

Each time series is divided into a training part, on which models are estimated, and a holdout period, on which forecasts are produced. Due to space constraints, the focus is on yearly and quarterly data. The forecast horizon is six for the yearly series and eight for the quarterly series.

Four competing models are considered, in addition to the RWDAR. The first one is the ETS, which is the exponential smoothing state-space model proposed by Hyndman, Koehler, Snyder, and Grose (2002). The second is simple exponential smoothing with drift, referred to as the Theta–forecast adaptation model (Hyndman & Billah, 2003), akin to the Theta model introduced by Assimakopoulos and Nikolopoulos (2000), which performed very well in the M3 competition. The third model is the automatic ARIMA modeling algorithm (auto-ARIMA), which corresponds to the best ARIMA model automatically selected using information criteria; see Hyndman and Khandakar (2008). Lastly, the fourth model is EAT, which is a combination of an ETS model, an ARIMA model, and the Theta method with equal weights. Forecasts produced by the RWDAR are compared against the predictions delivered by the other four methods, obtained with the R forecast package (Hyndman et al., 2020; Hyndman & Khandakar, 2008).

Two scale-independent error measures are employed in the forecast evaluation, namely, the symmetric mean absolute percentage error (or sMAPE, introduced by Makridakis (1993)) and the mean absolute scaled error (or MASE, proposed by Hyndman & Koehler, 2006):

$$MASE = \frac{1}{h}\frac{\sum_{t=n+1}^{n+h}|y_t-\hat{y}_t|}{\frac{1}{n-m}\sum_{t=m+1}^{n}|y_t-y_{t-m}|}$$

$$sMAPE = \frac{2}{h}\sum_{t=n+1}^{n+h}\frac{|y_t-\hat{y}_t|}{|y_t|+|\hat{y}_t|}\times 100, \qquad (10)$$

where y_t is the observed value of the series at time t, $\hat{y}_t$ is the forecast made at time t, h is the forecast horizon, n is the number of data points in-sample, and m is the frequency of the data (e.g., $m = 1$ for yearly data, $m = 4$ for quarterly, etc.). These two performance measures for point forecasts were used in the M4 competition (Makridakis, Spiliotis, & Assimakopoulos, 2020) and therefore are also applied in this paper.

A full description of the M3 results is provided in the sequel. The results are discussed separately for yearly and quarterly data.[7]

5.1. Yearly M3 series

Results for yearly (non-seasonal) M3 series are presented in Table 6.

In order to improve the readability of the table, the results are presented in terms of the MASE and sMAPE ratios instead of the individual MASE and sMAPE values. Thus, each cell in the top part of the table displays the ratio of the average MASE value for RWDAR relative to the average MASE values for the competing models: ETS, Theta, auto-ARIMA, and EAT.

Each cell in the bottom part of the table contains the ratio of the average sMAPE value for RWDAR relative to the average sMAPE values for the competing models. Values smaller than one indicate that the average MASE or sMAPE value for RWDAR is smaller than the corresponding average for any of the competing models.

When focusing on mean ratios, the RWDAR model has smaller MASE and sMAPE values overall than the ETS,

[7] Results for monthly data are available from the authors upon request.

Table 6
Out-of-sample forecast comparison using yearly M3 series.

No. steps ahead	Mean ratio RWDAR vs. ETS	Mean ratio RWDAR vs. Theta	Mean ratio RWDAR vs. ARIMA	Mean ratio RWDAR vs. EAT	Med. ratio RWDAR vs. ETS	Med. ratio RWDAR vs. Theta	Med. ratio RWDAR vs. ARIMA	Med. ratio RWDAR vs. EAT
	MASE							
$h=1$	0.951	0.922	0.931	0.988	1.006	0.904	1.042	1.047
$h=2$	0.939	0.918	0.917	0.987	0.941	0.841	0.990	0.956
$h=3$	0.938	0.927	0.911	0.983	0.970	0.850	0.984	0.979
$h=4$	0.927	0.936	0.900	0.977	0.991	0.886	1.041	0.999
$h=5$	0.921	0.943	0.892	0.976	0.994	0.912	1.044	1.075
$h=6$	0.919	0.947	0.889	0.977	0.979	0.941	1.002	1.019
Average	0.932	0.932	0.906	0.980	0.976	0.889	1.017	1.012
	sMAPE							
$h=1$	0.906	0.919	0.907	0.954	0.935	0.825	0.960	0.928
$h=2$	0.918	0.923	0.922	0.972	0.945	0.915	0.975	1.022
$h=3$	0.923	0.924	0.922	0.974	0.943	0.865	0.932	0.967
$h=4$	0.922	0.926	0.919	0.975	0.958	0.930	0.939	0.977
$h=5$	0.920	0.929	0.914	0.974	0.941	0.957	0.951	1.015
$h=6$	0.919	0.933	0.914	0.975	0.993	0.990	0.971	1.048
Average	0.918	0.926	0.916	0.971	0.952	0.914	0.955	0.993

Notes. Columns "Mean ratio" and "Med. ratio" report, for each forecast horizon (number of steps ahead), $\frac{Mean\ RWDAR}{Mean\ ETS}$, $\frac{Mean\ RWDAR}{Mean\ Theta}$, $\frac{Mean\ RWDAR}{Mean\ ARIMA}$, $\frac{Mean\ RWDAR}{Mean\ EAT}$ and $\frac{Median\ RWDAR}{Median\ ETS}$, $\frac{Median\ RWDAR}{Median\ Theta}$, $\frac{Median\ RWDAR}{Median\ ARIMA}$, $\frac{Median\ RWDAR}{Median\ EAT}$, respectively, of the corresponding MASEs (top half of the table) and sMAPEs (bottom half of the table).

Theta, auto-ARIMA, and EAT methods. The general message is consistent across the two error measures. While the RWDAR model scores better than ETS, Theta, and auto-ARIMA, the differences are less pronounced between the RWDAR and the ensemble EAT method. The better performance of RWDAR compared to the Theta model might be explained by the fact that most of the yearly series used in the M3 and M4 competitions are in levels (and not in logs), in line with the results presented in Section 4.2.

When examining the median ratios, the picture is broadly unchanged, even though the differences across models are less discernible in some cases. However, the RWDAR model appears to be more accurate than the other competing methods most of the time, with very few exceptions. The second best performing model is EAT. Overall, the RWDAR model turns out to be the most accurate forecasting method, with very good performance in the case under analysis.

5.2. Quarterly M3 series

This section examines the quarterly M3 series. In this case, seasonality has to be taken into account and adjusted for if it is present. The RWDAR model is not suitable in the presence of a seasonal component. One can include seasonality by using a dummy or a trigonometric variable, as described in Casals, Garcia-Hiernaux, Jerez, Sotoca, and Trindade (2018). However, in this case one cannot use the results with Propositions 1 and 2, given the lack of algebraic mapping, and the full Kalman filter must be run.

Another approach relies on an "external" seasonal adjusting procedure, as in the Theta method of Assimakopoulos and Nikolopoulos (2000), and seasonality is removed by performing multiplicative classical decomposition (Fiorucci et al., 2016).[8] This approach is also used for the RWDAR model and allows our results with Propositions 1 and 2 to be applied. Specifically, whenever seasonality is detected,[9] RWDAR is estimated on a seasonally adjusted (or deseasonalized) time series. Seasonal adjustment is achieved through the following multiplicative decomposition (see Harvey, 1989): defining s as the seasons, and assuming s is even, the smoothed series is given by $y_t^* = \sum_{j=-\frac{s}{2}}^{\frac{s}{2}} w_j y_{t-j}$ with $t = \frac{s}{2}+1, \dots, n-\frac{s}{2}-1$ where $w_j = \frac{1}{s}$ for $j = 0, \pm 1, \dots, \pm(\frac{s}{2}-1)$ and $w_{-j} = w_j = \frac{1}{2s}$ for $j = \pm\frac{s}{2}$. The ratios $\frac{y_t}{y_t^*}$ are then averaged over each season to give a set of seasonal factors. Dividing the original series by these seasonal factors gives the seasonally adjusted series. The RWDAR forecasts made on the deseasonalized series are finally multiplied by the seasonal factors to obtain the seasonal forecasts.

The results for quarterly M3 series are presented in Table 7. When looking at mean ratios, the RWDAR model usually has smaller MASE and sMAPE values than ETS, Theta, and auto-ARIMA, except for very short forecast horizons ($h = 1, 2, 3$). However, it is outperformed by the EAT method, which—being a combination of an ETS model, an auto-ARIMA model, and the Theta method—is a very strong benchmark. These results hold true according to both the MASE and sMAPE measures.

8 See Nikolopoulos and Thomakos (2019), p. 10, for additional details on the steps to be followed for creating forecasts from the basic Theta model.

9 An ETS model is fitted using the "ets" function in R, and if the chosen model has a seasonal component, then the data are considered to be seasonal and seasonality is removed as described above. For additional details, see https://robjhyndman.com/hyndsight/detecting-seasonality/.

Table 7
Out-of-sample forecast comparison using quarterly M3 series.

No. steps ahead	Mean ratio RWDAR vs. ETS	Mean ratio RWDAR vs. Theta	Mean ratio RWDAR vs. ARIMA	Mean ratio RWDAR vs. EAT	Med. ratio RWDAR vs. ETS	Med. ratio RWDAR vs. Theta	Med. ratio RWDAR vs. ARIMA	Med. ratio RWDAR vs. EAT
	MASE							
$h=1$	1.059	0.961	1.021	1.103	1.142	0.942	1.123	1.231
$h=2$	1.040	0.987	1.028	1.099	1.112	1.029	1.054	1.194
$h=3$	1.020	0.988	0.994	1.081	1.060	1.029	0.985	1.127
$h=4$	0.992	0.978	0.964	1.055	1.040	1.001	0.979	1.127
$h=5$	0.976	0.979	0.947	1.052	1.021	1.011	0.990	1.111
$h=6$	0.962	0.976	0.939	1.045	0.977	0.991	0.982	1.089
$h=7$	0.945	0.977	0.925	1.035	1.009	1.022	1.021	1.108
$h=8$	0.928	0.972	0.913	1.021	0.946	0.973	0.975	1.047
Average	0.990	0.977	0.967	1.061	1.038	0.999	1.013	1.129
	sMAPE							
$h=1$	1.022	0.986	1.003	1.066	1.130	0.957	1.116	1.147
$h=2$	1.029	1.009	1.007	1.072	1.119	1.032	1.083	1.172
$h=3$	1.019	1.005	0.973	1.055	1.061	1.071	0.990	1.127
$h=4$	0.998	1.004	0.954	1.039	0.971	1.012	0.948	1.081
$h=5$	0.984	1.002	0.941	1.036	0.952	1.006	0.917	1.076
$h=6$	0.971	0.998	0.929	1.029	0.970	1.011	0.912	1.100
$h=7$	0.956	0.994	0.918	1.020	0.994	1.029	0.940	1.056
$h=8$	0.938	0.987	0.909	1.008	0.914	0.965	0.922	1.008
Average	0.989	0.998	0.954	1.041	1.014	1.011	0.978	1.096

Notes. Columns "Mean ratio" and "Med. ratio" report, for each forecast horizon (number of steps ahead), $\frac{Mean\ RWDAR}{Mean\ ETS}$, $\frac{Mean\ RWDAR}{Mean\ Theta}$, $\frac{Mean\ RWDAR}{Mean\ ARIMA}$, $\frac{Mean\ RWDAR}{Mean\ EAT}$ and $\frac{Median\ RWDAR}{Median\ ETS}$, $\frac{Median\ RWDAR}{Median\ Theta}$, $\frac{Median\ RWDAR}{Median\ ARIMA}$, $\frac{Median\ RWDAR}{Median\ EAT}$, respectively, of the corresponding MASEs (top half of the table) and sMAPEs (bottom half of the table).

Table 8
Types of M4 series per frequency and domain.

Frequency	Micro	Industry	Macro	Finance	Demographic	Other	Total
Yearly	6538	3716	3903	6519	1088	1,236	23,000
Quarterly	6020	4637	5315	5305	1858	865	24,000
Monthly	10,975	10,017	10,016	10,987	5728	277	48,000
Weekly	112	6	41	164	24	12	359
Daily	1,476	422	127	1559	10	633	4,227
Hourly	0	0	0	0	0	414	414
Total	25,121	18,798	19,402	24,534	8708	3437	100,000

Considering the median ratios, the RWDAR model again performs rather well, and tends to be as accurate as the ETS, Theta, and auto-ARIMA methods, while it is overcome by the EAT ensemble method. Also, in this case no significant differences emerge according to the two metrics of forecast accuracy used (MASE and sMAPE).

6. Application to the M4 competition

This section explores the out-of-sample predictive performance of the RWDAR model using the M4 competition database (Makridakis et al., 2020). The M4 competition dataset increases the sample size to 100,000 series. In addition, besides yearly, quarterly, and monthly series, even higher-frequency data are included, in the form of weekly, daily, and hourly data. The number of forecasting methods is expanded, and prediction intervals, as well as point forecasts, are incorporated in the evaluation process in order to address the issue of forecast uncertainty. The time series belong to six different domains (see Table 8) taken from Makridakis et al. (2020).

As for the M3 dataset, each time series is divided into a training part, on which models are estimated, and a holdout period, on which forecasts are produced and prediction errors computed. Yearly and quarterly data are examined. The forecast horizon is six for yearly series and eight for quarterly ones. As in the previous application, two scale-independent error measures are employed in the forecast evaluation: the sMAPE and the MASE.

A detailed description and a discussion of the M4 results is provided in what follows, separately for each time frequency (yearly and quarterly).[10]

6.1. Yearly M4 series

The focus is initially on the 23,000 yearly (non-seasonal) series. Results are presented in Table 9.

When focusing on mean ratios, the RWDAR model has overall smaller MASE and sMAPE values than the ETS, Theta, auto-ARIMA, and EAT methods. Specifically, the

10 Results for monthly data are available from the authors upon request.

Table 9
Out-of-sample forecast comparison using yearly M4 series.

No. steps ahead	Mean ratio RWDAR vs. ETS	Mean ratio RWDAR vs. Theta	Mean ratio RWDAR vs. ARIMA	Mean ratio RWDAR vs. EAT	Med. ratio RWDAR vs. ETS	Med. ratio RWDAR vs. Theta	Med. ratio RWDAR vs. ARIMA	Med. ratio RWDAR vs. EAT
	MASE							
$h = 1$	0.948	0.989	0.931	1.002	0.966	0.982	0.957	1.020
$h = 2$	0.925	0.976	0.909	0.987	0.930	0.983	0.921	0.984
$h = 3$	0.916	0.950	0.909	0.983	0.914	0.958	0.910	0.969
$h = 4$	0.907	0.929	0.908	0.980	0.902	0.921	0.900	0.962
$h = 5$	0.896	0.917	0.902	0.973	0.897	0.911	0.897	0.953
$h = 6$	0.888	0.906	0.900	0.971	0.892	0.898	0.897	0.954
Average	0.913	0.944	0.909	0.982	0.917	0.942	0.914	0.974
	sMAPE							
$h = 1$	0.952	0.994	0.931	1.001	0.954	0.985	0.946	1.016
$h = 2$	0.937	0.992	0.916	0.992	0.920	0.958	0.908	0.981
$h = 3$	0.924	0.975	0.913	0.988	0.917	0.928	0.914	0.979
$h = 4$	0.912	0.961	0.912	0.982	0.905	0.904	0.907	0.972
$h = 5$	0.900	0.950	0.906	0.975	0.906	0.887	0.907	0.966
$h = 6$	0.891	0.940	0.903	0.971	0.903	0.877	0.911	0.968
Average	0.920	0.970	0.913	0.985	0.917	0.923	0.916	0.980

Notes. Columns "Mean ratio" and "Med. ratio" report, for each forecast horizon (number of steps ahead), $\frac{Mean\ RWDAR}{Mean\ ETS}$, $\frac{Mean\ RWDAR}{Mean\ Theta}$, $\frac{Mean\ RWDAR}{Mean\ ARIMA}$, $\frac{Mean\ RWDAR}{Mean\ EAT}$ and $\frac{Median\ RWDAR}{Median\ ETS}$, $\frac{Median\ RWDAR}{Median\ Theta}$, $\frac{Median\ RWDAR}{Median\ ARIMA}$, $\frac{Median\ RWDAR}{Median\ EAT}$, respectively, of the corresponding MASEs (top half of the table) and sMAPEs (bottom half of the table).

RWDAR model clearly outperforms both ETS and auto-ARIMA, while differences are smaller compared to the Theta and EAT methods. Thus, the performance of the Theta method seems to improve in the case under examination compared to the M3 dataset.

The results are overall consistent when looking at median ratios. Also, in this case the RWDAR model is in general more accurate than the other competing methods; the second best performing model is again the ensemble EAT method, while the third best performing model is the Theta method. This evidence confirms the results obtained with the yearly M3 series in Section 5.1.

6.2. Quarterly M4 series

The results for quarterly M4 series are presented in Table 10. When working with quarterly data, the RWDAR performs slightly better than the Theta method based on mean MASE and sMAPE values. However, it tends to be outperformed by the other methods, even if differences between the mean MASE values are rather small across all forecast horizons. Similar conclusions can be reached by inspecting the median sMAPE values.

Overall, these results show that the forecast performance is similar across all the individual methods considered. This is largely expected, given that RWDAR does not include seasonality and, as with the other methods, relies on an "external" seasonal adjustment procedure. As a result, unlike with the yearly frequency, with quarterly data, differences in the forecast performance across models are wiped out.

7. Forecast combinations

We have seen that the RWDAR model often outperforms the individual counterparts (ETS, Theta, and auto-ARIMA), but not the combination of all of them (EAT). This result is in line with the empirical literature on forecast combinations, which have frequently been found to produce on average more accurate predictions than individual forecasting models (Timmermann, 2006). Thus, according to this literature, one way to improve forecast accuracy is through the combination of several forecasts of the same variable into a single forecast (Bates & Granger, 1969). This section further elaborates on this issue.[11]

Specifically, we investigate whether any gains in forecast accuracy over RWDAR can be achieved by combining RWDAR predictions with forecasts from the other individual models under examination (ETS, Theta, and auto-ARIMA). For the sake of simplicity, equally weighted forecast combinations are considered. Yearly (non-seasonal) data from the M3 and M4 competition datasets are used, and forecasts are produced from one up to six steps ahead. Average and median MASE and sMAPE values are again employed as scoring measures.

Two forecasting models are assessed:

1. RWDAR;
2. The following forecast combination, which gives equal weight to RWDAR and to the EAT ensemble method; this latter, in turn, represents an equally weighted combination of ETS, Theta, and auto-ARIMA:

$$\begin{aligned} COMB &= \frac{1}{2}RWDAR + \frac{1}{2}EAT \\ &= \frac{1}{2}RWDAR + \frac{1}{6}(ETS + Theta + ARIMA) \end{aligned}$$

The results are presented in Table 11. Values larger than one can be frequently inspected in the table, suggesting a poorer performance of the RWDAR model compared

[11] This experiment was suggested by an anonymous reviewer, to whom we are grateful.

Table 10
Out-of-sample forecast comparison using quarterly M4 series.

No. steps ahead	Mean ratio RWDAR vs. ETS	Mean ratio RWDAR vs. Theta	Mean ratio RWDAR vs. ARIMA	Mean ratio RWDAR vs. EAT	Med. ratio RWDAR vs. ETS	Med. ratio RWDAR vs. Theta	Med. ratio RWDAR vs. ARIMA	Med. ratio RWDAR vs. EAT
	MASE							
$h=1$	1.047	0.984	1.055	1.078	1.095	0.960	1.087	1.104
$h=2$	1.045	0.982	1.043	1.073	1.076	0.975	1.059	1.093
$h=3$	1.046	0.979	1.040	1.071	1.061	0.965	1.053	1.080
$h=4$	1.038	0.974	1.030	1.061	1.061	0.956	1.044	1.069
$h=5$	1.033	0.970	1.024	1.056	1.051	0.952	1.039	1.067
$h=6$	1.026	0.965	1.018	1.051	1.048	0.948	1.035	1.059
$h=7$	1.020	0.960	1.014	1.048	1.041	0.945	1.028	1.053
$h=8$	1.015	0.956	1.010	1.043	1.037	0.944	1.026	1.048
Average	1.034	0.971	1.029	1.060	1.059	0.956	1.046	1.072
	sMAPE							
$h=1$	1.013	1.003	1.013	1.054	1.107	0.970	1.104	1.114
$h=2$	1.019	1.009	1.004	1.056	1.090	0.971	1.066	1.097
$h=3$	1.021	1.009	1.004	1.056	1.073	0.963	1.058	1.074
$h=4$	1.019	1.009	1.000	1.051	1.072	0.965	1.053	1.075
$h=5$	1.014	1.006	0.996	1.047	1.059	0.962	1.042	1.066
$h=6$	1.008	1.003	0.992	1.044	1.043	0.956	1.032	1.057
$h=7$	1.004	1.000	0.990	1.041	1.036	0.953	1.031	1.052
$h=8$	1.000	0.997	0.988	1.038	1.022	0.945	1.018	1.044
Average	1.012	1.004	0.999	1.048	1.063	0.961	1.050	1.072

Notes. Columns "Mean ratio" and "Med. ratio" report, for each forecast horizon (number of steps ahead), $\frac{Mean\ RWDAR}{Mean\ ETS}$, $\frac{Mean\ RWDAR}{Mean\ Theta}$, $\frac{Mean\ RWDAR}{Mean\ ARIMA}$, $\frac{Mean\ RWDAR}{Mean\ EAT}$ and $\frac{Median\ RWDAR}{Median\ ETS}$, $\frac{Median\ RWDAR}{Median\ Theta}$, $\frac{Median\ RWDAR}{Median\ ARIMA}$, $\frac{Median\ RWDAR}{Median\ EAT}$, respectively, of the corresponding MASEs (top half of the table) and sMAPEs (bottom half of the table).

to the forecast combination (COMB). The results are quite consistent across datasets and forecast horizons, with a single exception ($h = 1$, M3, sMAPE). Focusing on the M3 dataset, gains from forecast combination tend to be greater when considering median ratios; this is no longer true when examining the M4 dataset on the right-hand side of the table.

Even if gains arising from forecast combinations are not particularly large, it is clear from this simple empirical application that averaging across models produces a pooled forecast measure of the target variable that is more accurate than the predictions delivered by the single models. This is perfectly consistent with previous findings from the M3 competition dataset (Makridakis & Hibon, 2000), which highlighted the benefits of the combination of various forecast methods, often outperforming the individual methods being combined. In addition, it appears evident from Table 11 that adding RWDAR in the set of models used to form the forecast combination leads to an improvement in terms of forecast accuracy. This result is noteworthy and might be usefully exploited in other forecast competitions.

Table 11
Out-of-sample forecast combination comparison using yearly M3 and M4 series.

	M3		M4	
No. steps ahead	Mean ratio RWDAR vs. COMB	Med. ratio RWDAR vs. COMB	Mean ratio RWDAR vs. COMB	Med. ratio RWDAR vs. COMB
	MASE		MASE	
$h=1$	1.013	1.061	1.018	1.030
$h=2$	1.018	1.012	1.012	1.015
$h=3$	1.015	1.010	1.012	1.007
$h=4$	1.012	1.033	1.012	1.005
$h=5$	1.012	1.059	1.010	0.998
$h=6$	1.015	1.031	1.010	1.001
Average	1.014	1.034	1.012	1.009
	sMAPE		sMAPE	
$h=1$	0.992	0.990	1.015	1.026
$h=2$	1.008	1.120	1.012	1.006
$h=3$	1.012	1.024	1.013	1.018
$h=4$	1.015	1.017	1.013	1.022
$h=5$	1.017	1.014	1.011	1.018
$h=6$	1.021	1.036	1.011	1.015
Average	1.011	1.034	1.013	1.018

Notes. Columns "Mean ratio" and "Med. ratio" report, for each forecast horizon (number of steps ahead) and each forecast competition (M3 and M4), $\frac{Mean\ RWDAR}{Mean\ COMB}$ and $\frac{Median\ RWDAR}{Median\ COMB}$ of the corresponding MASEs (top half of the table) and sMAPEs (bottom half of the table).

8. Conclusions

This paper presented a new state-space model for producing time-series forecasts. The model features two separate components: a random walk with drift, akin to the Theta method of Assimakopoulos and Nikolopoulos (2000); and a first-order autoregressive term, which helps to model the stationary part of the series. This model is called Random Walk with Drift plus AutoRegressive (RWDAR), and belongs to the class of unobserved components models pioneered by Harvey (1986) and extensively discussed in Durbin and Koopman (2012) and Proietti (2021). It can be easily estimated by implementing the estimation method proposed in this paper, which hinges upon the analytical solution to the algebraic Riccati equation for the covariance matrix of the state vector's prediction error. A Monte Carlo simulation study compared the performance of the newly proposed estimator with that of the standard likelihood maximization procedure based on the Kalman filter, focusing on small and moderate-sized samples. This comparison showed that both methods have similar performance, irrespective of the sample size considered.

The out-of-sample forecast performance of the RWDAR model was firstly assessed in a Monte Carlo experiment. The Monte Carlo study compared the forecast performance of RWDAR and simple exponential smoothing with drift, referred to as the Theta–forecast adaptation model (Hyndman & Billah, 2003), and considered both non-stationary and stationary (or nearly non-stationary) data-generating processes. Secondly, two empirical applications were presented and discussed. The first empirical application was based on the M3 competition dataset (Makridakis & Hibon, 2000), and compared the RWDAR performance to that of several relevant competing methods (ETS, Theta, auto-ARIMA, and a combination of these three methods, termed "EAT"'). The second application focused on the M4 competition dataset (Makridakis et al., 2020), which extends the M3 database in terms of the number of series used and forecasting methods considered. The results were presented for different time frequencies (yearly and quarterly) and competing models.

As far as the empirical contribution of the paper is concerned, the results demonstrate that the forecasts produced by the RWDAR model are on average more accurate than those obtained when using several traditional statistical models, such as ETS and auto-ARIMA. Furthermore, RWDAR scored well even when compared to simple exponential smoothing with drift. Additional gains were made by combining forecasts derived from RWDAR with predictions delivered by the other methods considered in this study. The good performance of RWDAR is probably due to the combination of the random walk plus drift with the first-order autoregressive component, which provides both modeling adaptability and flexibility. Overall, these results suggest that the RWDAR model is well suited for producing forecasts for a large number of time series sampled at different frequencies, and can thus be used effectively in a wide variety of forecasting problems.

Further research should explore additional datasets in order to appreciate whether the RWDAR model is suitable to forecast time series in specific domains. Another aspect to be examined is whether we might enhance the accuracy of the resulting forecasts by combining RWDAR predictions with those produced by other time-series models not considered in this paper. All these issues are left for future work.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Proof of Proposition 1

Proof. It is well known that a univariate model can be reparameterized by concentrating out one parameter (see, for instance, Harvey, 1989). This procedure reduces the number of parameters to be estimated, therefore simplifying the maximization of the log-likelihood function. Following this route, let us consider the following representation of the RWDAR model in (1):

$$\begin{aligned} y_t &= \lambda_{t-1} + \beta_{t-1}, & t &= 1, 2, \ldots, n, \\ \lambda_t &= \tau + \lambda_{t-1} + \epsilon_t, & \epsilon_t &\sim \mathcal{NID}(0, \sigma_\eta^2 q), \\ \beta_t &= \phi\beta_{t-1} + \eta_t, & \eta_t &\sim \mathcal{NID}(0, \sigma_\eta^2), \end{aligned} \tag{A.1}$$

with $q = \frac{\sigma_\epsilon^2}{\sigma_\eta^2}$. Concentrating out σ_η^2, and imposing $\sigma_\eta^2 = 1$, the number of parameters reduces from four to three (ϕ ; τ ; q). The Kalman filter recursions for model (A.1) can then be written as before, with some small changes. In particular, the innovation variance becomes $f_t^c = \mathbf{z}\mathbf{P}_t^c\mathbf{z}'$, where the superscript "c" stands for "concentrated". Moreover, the last recursion (the equation for the covariance matrix of the state vector's prediction error) becomes:

$$\mathbf{P}_t^c = \mathbf{T}\mathbf{P}_{t-1}^c\mathbf{T}' - \mathbf{T}\mathbf{P}_{t-1}^c\mathbf{z}'(f^c)^{-1}\mathbf{z}\mathbf{P}_{t-1}^c\mathbf{T}' + \begin{bmatrix} q & 0 \\ 0 & 1 \end{bmatrix}$$

Defining its steady-state matrix $\bar{\mathbf{P}}^c = \begin{bmatrix} p_1 & p_2 \\ p_2 & p_3 \end{bmatrix}$, the Riccati matrix equation for (A.1) is:

$$\begin{bmatrix} \frac{(p_1+p_2)^2}{p_1+2p_2+p_3} - q & \frac{p_2(p_3-(\phi-2)p_2)+p_1(p_2+\phi p_3)}{p_1+2p_2+p_3} \\ \frac{p_2(p_3-(\phi-2)p_2)+p_1(p_2+\phi p_3)}{p_1+2p_2+p_3} & \frac{\left(p_2^2-p_1p_3\right)\phi^2}{p_1+2p_2+p_3} + p_3 - 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \tag{A.2}$$

Firstly, we demonstrate the existence of two analytic solutions to the matrix Eq. (A.2). Secondly, we show that only the solution in Proposition 1 allows $\bar{\mathbf{P}}^c$ to be positive-definite, and therefore it is the unique legitimate solution.

Recalling Sylvester's criterion in matrix algebra, there are two necessary and sufficient conditions to guarantee that $\bar{\mathbf{P}}^c$ is positive-definite. The first condition is that p_1 must be strictly positive. The second condition is that the determinant of $\bar{\mathbf{P}}^c$ must be positive.

Solving the upper-left equation of (A.2) with respect to p_2 and the extra-diagonal equation of (A.2) with respect to p_3, and rearranging, one finds (5) and (6), respectively.

Substituting the solutions for p_2 and p_3 in the bottom-right equation, one obtains the following equation:

$$\frac{p_1(\phi-1)\left(p_1(\phi-1) - 3q\phi + q\right) + q(2q(\phi-1)\phi - 1)}{q} = 0$$

This equation has two different solutions for p_1. However, it can be numerically shown that the only strictly positive solution is (4).

In addition, substituting (4) in the expressions for p_2 and p_3, it can be shown numerically that (4), (5), and (6) jointly guarantee that the determinant of $\bar{\mathbf{P}}^c$ is strictly positive. This, together with the positivity of p_1, satisfies Sylvester's criterion and therefore confirms that the solution is legitimate. Finally, one can also show that the remaining solution does not satisfy the conditions of Sylvester's criterion. This rules out the choice of this alternative solution. Therefore, it can be concluded that (4), (5), and (6) provide the only solution to the matrix Eq. (A.2) allowing $\bar{\mathbf{P}}^c$ to be positive-definite. Finally, this also implies that $\bar{\mathbf{P}} = \sigma_\eta^2 \times \bar{\mathbf{P}}^c$ is positive-definite. This concludes the proof of Proposition 1.[12] □

Appendix B. Proof of Proposition 2

Proof.

As noted by Harvey (1989) and Durbin and Koopman (2012), the knowledge of a steady-state solution such as Proposition 1 leads to some computational savings, since some of the Kalman filter recursions become redundant. That is, the analytical expressions in (4), (5), and (6) also provide the algebraic solutions for f and $\mathbf{k}$.

Indeed, using the exact solution of the Riccati equation, it follows immediately that the innovation variance f is:

$$f = \mathbf{z}\bar{\mathbf{P}}\mathbf{z}'$$
$$= \sigma_\eta^2 \left(\frac{1}{2} \left(-\sqrt{q}\sqrt{q(\phi+1)^2 + 4}(\phi - 1) + q\left(\phi^2 + 1\right) + 2 \right) \right) \tag{B.1}$$

Moreover, the steady-state Kalman gain can be derived as $\mathbf{k} = \mathbf{T}\bar{\mathbf{P}}\mathbf{z}'f^{-1}$. That is, using $\bar{\mathbf{P}}$, as in Proposition 1 and Eq. (B.1), one obtains the expression as in (9).

Therefore, the only Kalman recursions that need to be run, once the steady state is achieved, are the updating equation for the state vector and the innovation equation, that is, Eqs. (7) and (8).

We now need to prove that the minimization of the sum of squared innovations maximizes the concentrated log-likelihood function. This has been already shown by Harvey (1989), p. 129. This makes the estimation process trivial. Indeed, given the data, the choice of (ϕ ; τ ; q) directly determines the concentrated log-likelihood function, which can be maximized using standard optimization routines (for example, for those using R, the function "optim"). This procedure has also been recently employed by Sbrana and Silvestrini (2019).

Noted that Harvey (1989) argued that the asymptotic properties of an estimator minimizing $\sum_t^n v_t^2$ are equivalent to these of the maximum likelihood estimator. This concludes the proof of Proposition 2. □

Appendix C. R code to estimate the RWDAR model as in Proposition 2

This section presents the developed R code for estimating the model as in Proposition 2. It is worth noting that (p, co, q) stands for (ϕ, τ, q) as in Propositions 1 and 2. In addition, in order to ensure that $0 \leq \phi \leq 1$, the transformation $1 - exp(-abs(x))$ is used.

```
LogLikelihood <- function(Kg){q<-abs(Kg[1]);
p<-1-exp(-abs(Kg[2]));co<-abs(Kg[3])
mu <-beta <-v<-c()
v[1]=0; mu[1]=y[1]; beta[1]<-0
k1<- -(2*sqrt(q))/(sqrt(q)*(-1+p)-sqrt(4+q* (1+p)
     ^2))
k2<-(p*(2+q+q*p-sqrt(q)*sqrt(4+q *(1+p)^2)))/(2+2 *
     q*p)
for (t in 2:(obs)) {
  v[t]<-y[t]-mu[t-1]-beta[t-1]
  mu[t]<-co+mu[t-1]+k1*v[t]
  beta[t]<-p*beta[t-1]+k2*v[t]
}
sum(v^2)
}
Estimation<-optim(c(.8,.4,2),LogLikelihood)
```

In the above optimization stage (that is, in the "optim" function), the following starting values are employed: $phi = 1 - exp(-0.8)$, $q = 0.4$ and $\tau = 2$. Please note that "y" is the name to be assigned to the time series.

[12] All the results, as well as the numerical proofs, were derived using *Mathematica* 9 by Wolfram. In order to reduce the algebra, the solutions that do not satisfy Sylvester's criterion are not provided. If needed, these are available from the authors upon request.

References

Assimakopoulos, V., & Nikolopoulos, K. (2000). The theta model: a decomposition approach to forecasting. *International Journal of Forecasting*, *16*(4), 521–530, The M3- Competition.

Bates, J. M., & Granger, C. W. J. (1969). The combination of forecasts. *Journal of the Operational Research Society*, *20*(4), 451–468.

Beveridge, Stephen, & Nelson, Charles (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary Economics*, *7*(2), 151–174.

Canova, Fabio (1998). Detrending and business cycle facts. *Journal of Monetary Economics*, *41*(3), 475–512.

Casals, Jose, Garcia-Hiernaux, Alfredo, Jerez, Miguel, Sotoca, Sonia, & Trindade, A. Alexandre (2018). *State-space methods for time series analysis: Theory, applications and software*. CRC Press.

Durbin, James, & Koopman, Siem Jan (2012). *Time series analysis by state space methods*. Oxford University Press.

Fiorucci, Jose A., Pellegrini, Tiago R., Louzada, Francisco, Petropoulos, Fotios, & Koehler, Anne B. (2016). Models for optimising the theta method and their relationship to state space models. *International Journal of Forecasting*, *32*(4), 1151–1161.

Harvey, Andrew (1985). Trends and cycles in macroeconomic time series. *Journal of Business & Economic Statistics*, *3*(3), 216–227.

Harvey, Andrew C. (1986). Analysis and generalisation of a multivariate exponential smoothing model. *Management Science*, *32*(3), 374–380.

Harvey, Andrew C. (1989). *Forecasting, structural time series analysis, and the kalman filter*. Cambridge.

Hodrick, Robert, & Prescott, Edward (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, *29*(1), 1–16.

Hyndman, Rob, Athanasopoulos, George, Bergmeir, Christoph, Caceres, Gabriel, Chhay, Leanne, O'Hara-Wild, Mitchell, et al. (2020). Forecast: forecasting functions for time series and linear models. R package version 8.13.

Hyndman, Rob J., & Billah, Baki (2003). Unmasking the theta method. *International Journal of Forecasting*, *19*(2), 287–290.

Hyndman, Rob J., & Khandakar, Yeasmin (2008). Automatic time series forecasting: the forecast package for R. *Journal of Statistical Software*, *26*(3), 1–22.

Hyndman, Rob J., & Koehler, Anne B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, *22*(4), 679–688.

Hyndman, Rob J., Koehler, Anne B., Snyder, Ralph D., & Grose, Simone (2002). A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of Forecasting*, *18*(3), 439–454.

Makridakis, Spyros (1993). Accuracy measures: theoretical and practical concerns. *International Journal of Forecasting*, *9*(4), 527–529.

Makridakis, Spyros, & Hibon, Michele (2000). The M3-competition: results, conclusions and implications. *International Journal of Forecasting*, *16*(4), 451–476.

Makridakis, Spyros, Spiliotis, Evangelos, & Assimakopoulos, Vassilios (2020). The M4 competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, *36*(1), 54–74, M4 Competition.

Nelson, Charles R., & Plosser, Charles I. (1982). Trends and random walks in macroeconmic time series: Some evidence and implications. *Journal of Monetary Economics*, *10*(2), 139–162.

Nikolopoulos, Konstantinos, Assimakopoulos, Vassilios, Bougioukos, Nikolaos, Litsa, Akrivi, & Petropoulos, Fotios (2011). The theta model: An essential forecasting tool for supply chain planning. In *Advances in automation and robotics, vol. 2* (pp. 431–437). Springer.

Nikolopoulos, Kostas I., & Thomakos, Dimitrios D. (2019). *Forecasting with the Theta Method: Theory and Applications*. John Wiley and Sons, Ltd.

Oh, Kum Hwa, Zivot, Eric, & Creal, Drew (2008). The relationship between the beveridge-nelson decomposition and other permanent-transitory decompositions that are popular in economics. *Journal of Econometrics*, *146*(2), 207–219.

Petropoulos, Fotios, & Nikolopoulos, Konstantinos (2017). The theta method. *Foresight: The International Journal of Applied Forecasting*, (46), 11–17.

Proietti, Tommaso (2021). Predictability, real time estimation, and the formulation of unobserved components models. *Econometric Reviews*, *40*(5), 433–454.

Sbrana, Giacomo, & Silvestrini, Andrea (2019). Random switching exponential smoothing: A new estimation approach. *International Journal of Production Economics*, *211*, 211–220.

Sbrana, Giacomo, & Silvestrini, Andrea (2020). Forecasting with the damped trend model using the structural approach. *International Journal of Production Economics*, *226*(C).

Sbrana, Giacomo, & Silvestrini, Andrea (2021). Random coefficient state-space model: Estimation and performance in M3–M4 competitions. *International Journal of Forecasting*, *38*(1), 352–366.

Spiliotis, Evangelos, Assimakopoulos, Vassilios, & Makridakis, Spyros (2020). Generalizing the theta method for automatic forecasting. *European Journal of Operational Research*, *284*(2), 550–558.

Spiliotis, Evangelos, Assimakopoulos, Vassilios, & Nikolopoulos, Konstantinos (2019). Forecasting with a hybrid method utilizing data smoothing, a variation of the theta method and shrinkage of seasonal factors. *International Journal of Production Economics*, *209*, 92–102, The Proceedings of the 19th International Symposium on Inventories.

Thomakos, D., & Nikolopoulos, K. (2014). Fathoming the theta method for a unit root process. *IMA Journal of Management Mathematics*, *25*(1), 105–124.

Thomakos, Dimitrios D., & Nikolopoulos, Konstantinos (2015). Forecasting multivariate time series with the theta method. *Journal of Forecasting*, *34*(3), 220–229.

Timmermann, Allan (2006). In G. Elliott, C. W. J. Granger, & A. Timmermann (Eds.), *Handbook of economic forecasting*: *vol. 1*, *Chapter 4 forecast combinations* (pp. 135–196). Elsevier.

Watson, Mark W. (1986). Univariate detrending methods with stochastic trends. *Journal of Monetary Economics*, *18*(1), 49–75.