

Berninger, Christoph; Stöcker, Almond; Rügamer, David

Article — Published Version

A Bayesian time-varying autoregressive model for improved short-term and long-term prediction

Journal of Forecasting

Provided in Cooperation with:

John Wiley & Sons

Suggested Citation: Berninger, Christoph; Stöcker, Almond; Rügamer, David (2021) : A Bayesian time-varying autoregressive model for improved short-term and long-term prediction, Journal of Forecasting, ISSN 1099-131X, Wiley, Hoboken, NJ, Vol. 41, Iss. 1, pp. 181-200, <https://doi.org/10.1002/for.2802>

This Version is available at:

<https://hdl.handle.net/10419/284759>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by/4.0/>

RESEARCH ARTICLE

A Bayesian time-varying autoregressive model for improved short-term and long-term prediction

Christoph Berninger¹ | Almond Stöcker² | David Rügamer¹

¹Department of Statistics, Ludwig Maximilian University of Munich, Munich, Germany

²School of Business and Economics, Humboldt University of Berlin, Berlin, Germany

Correspondence

Christoph Berninger, Department of Statistics, Ludwig Maximilian University of Munich, Munich, Germany.
Email: christoph.berninger@stat.uni-muenchen.de

Abstract

Motivated by the application to German interest rates, we propose a time-varying autoregressive model for short-term and long-term prediction of time series that exhibit a temporary nonstationary behavior but are assumed to mean revert in the long run. We use a Bayesian formulation to incorporate prior assumptions on the mean reverting process in the model and thereby regularize predictions in the far future. We use MCMC-based inference by deriving relevant full conditional distributions and employ a Metropolis-Hastings within Gibbs sampler approach to sample from the posterior (predictive) distribution. In combining data-driven short-term predictions with long-term distribution assumptions our model is competitive to the existing methods in the short horizon while yielding reasonable predictions in the long run. We apply our model to interest rate data and contrast the forecasting performance to that of a 2-Additive-Factor Gaussian model as well as to the predictions of a dynamic Nelson-Siegel model.

KEYWORDS

Bayesian time-varying autoregressive models, Gibbs sampler, interest rate models, long run regularization, MCMC metropolis-Hastings

1 | INTRODUCTION

To forecast an univariate time series the first model of choice is often a linear model. An archetype of this model class in the context of time series analysis is the autoregressive model of order 1 (AR(1)), which is defined as follows:

$$x_t = \alpha + \beta x_{t-1} + \epsilon_t, \quad (1)$$

where x_t represents the variable that is defined on $t \in \mathbb{Z}$ and was observed at time points $t = 1, \dots, T$. α and β are real valued constants, while $|\beta| < 1$ is assumed to ensure stationarity. The innovation process ϵ_t can be, e.g., a

Gaussian white noise process, i.e., an independent and identically (i.i.d.) normal distributed $\epsilon_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2)$ for all time points t . In this classical model, the one-step-ahead expectation $E(x_t | x_{t-1}) = \alpha + \beta x_{t-1}$ and variance $\text{Var}(x_t | x_{t-1}) = \sigma^2$ are closely linked to the marginal characteristics $E(x_t) = \frac{\alpha}{1-\beta}$ and $\text{Var}(x_t) = \frac{\sigma^2}{1-\beta^2}$ approached in the long run. In this sense, fitting the short-term behavior of a time series with a linear model has wide implications for its long-term behavior, and, conversely, controlling the long-term behavior of the model constrains its short-time fit. In practice, this close linkage may present an important limitation when short-term performance conflicts with long-term plausibility.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Journal of Forecasting* published by John Wiley & Sons Ltd.

This gets evident in the macroeconomic literature and, more specifically, in contrasting modeling approaches for interest rates, which motivate this work: Among others, Diebold and Li (2006) focus on predictions in the short horizon. The authors apply an AR(1) process to extracted factors of the yield curve. While they do make the long-term assumption that interest rates are principally mean reverting, the process exhibits an almost integrated behavior. Estimating the model parameters of a near integrated but stationary AR(1) model might give large estimation errors and lead to unrealistic long run mean estimates far beyond the range of the data. Duffee (2011) even discard the mean reversion/stationarity assumption. In their proposed random walk model, the prediction variance linearly increases in time leading to extreme values in the long run. The strong focus on the short horizon might lead to questionable and potentially implausible long-term behavior. Also the model of Caldeira and Torrent (2017), who apply a nonparametric functional data approach, has the focus on the short horizon. In contrast, Korn and Wagner (2019) apply a linear model to model long horizon features. Their proposal is a Gauss2++ model with a forward looking estimation approach, i.e., they calibrate the parameters of their model without using historical data but amongst current prices of interest rate derivatives and long run survey forecasts. The Gauss2++ model is a standard model in the insurance industry, where plausible forecasts in the longer run are required. As demonstrated later in this paper, their model achieves a realistic long run distribution, however, at the expense of inferior short term predictions.

The aim of this paper is to bridge the gap between short and long horizons, generalizing above mentioned approaches to a model with the flexibility to (a) sufficiently adapt to sample data to yield good short-term predictions and (b) apply suitable regularization to achieve plausible long- and middle-term forecasts at the same time. This is particularly, yet not exclusively, relevant in applications where a stable stationary distribution is assumed in the long run—despite observing a “temporary non-stationary behavior” in the available data where, e.g., unconstrained linear model fits or a Dickey-Fuller test (Dickey & Fuller, 1979) would indicate an integrated process.

As in general linearity is often a restrictive assumption in practice and many time series exhibit features that cannot be captured by a linear model (Hamilton, 1989), a lot of research has been conducted to introduce different types of nonlinear models in the last decades. In particular, nonlinear models offer more flexibility to account for both, short and long horizon.

A bi-linear model is an example of a nonlinear model, which assumes a nonlinear relationship between the

covariates and response variable (Granger & Andersen, 1978; Rao & Gabr, 2012), although not often used in macroeconomic applications (Morley, 2014). A more immediate approach is to allow one (or more) parameters of a linear model to change over time. This comprises the regime switching and time-varying parameter models.

Regime switching models can allow for a different mean reversion level in the short and the long horizon. This feature can be used to regularize the long run mean of the model. The first approaches to regime switching models were conducted by Quandt (1958), who considered a switching regression model extending a linear regression model by allowing the parameters to switch between different states according to a random variable. Bacon and Watts (1971) introduced a smooth transition model, which implements a smooth transition from one regime to another without a sudden jump. Goldfeld and Quandt (1973) introduced the Markov switching regression model and use a discrete latent Markov process to determine the current regime. These models were adapted to time series models by Lim and Tong (1980) and Chan and Tong (1986) introducing the threshold autoregressive model (TAR) and the smooth transition autoregressive model (STAR), respectively. Hamilton (1989) introduced the Markov switching autoregressive model for applications in economics, which have been investigated thoroughly together with different variants in the literature (Haggan & Ozaki, 1981; Jansen & Teräsvirta, 1996; Teräsvirta, 1994). Lanne and Saikkonen (2002) used a TAR-model, which only allows regime changes for the constant parameter α , and applied it to strongly autocorrelated time series data—which is very related to temporary non-stationary behavior of the time series and, therefore, to the aim of the paper. When there is, however, no concrete indication for the process dynamics to result from switches between discrete underlying states, we consider it more natural and promising to assume a continuous latent process.

In contrast to regime switching models, time-varying parameter models allow one (or more) of the parameters in a linear model to be driven by its own continuous process (Morley, 2014). For example, if the parameter vector $(\alpha, \beta, \sigma^2)$ of the linear AR(1) model becomes a stochastic process, this results in a time-varying autoregressive model of order 1 (TV-AR(1))

$$x_t = \alpha_t + \beta_t x_{t-1} + \epsilon_t \quad (2)$$

with $\epsilon_t \sim \mathcal{N}(0, \sigma_t^2)$. Certain distribution assumptions for the underlying stochastic process of the parameter vector $(\alpha_t, \beta_t, \sigma_t)$ are made in practice to complete the TV-AR

(1) model specification (Teräsvirta et al. 2010). Similar to the TAR model in Lanne and Saikkonen (2002) the time variation of the TV-AR(1) model can be restricted to the constant parameter α_t , resulting in a time-varying constant autoregressive model of order 1 (TVC-AR(1)):

$$x_t = \alpha_t + \beta x_{t-1} + \epsilon_t. \quad (3)$$

If $|\beta| < 1$ and the latent process of α_t is stationary, the process x is also stationary. But due to random shifts in the mean reversion level—because of the time-varying constant parameter—realizations of the model can resemble those of a (close to) random walk process, when restricting to a limited time window.

Another time-varying parameter model is the shifting endpoint model introduced by Kozicki and Tinsley (2001). Similar to the TVC-AR(1) model their approach has a time-varying mean reversion level, referred to as shifting endpoints. In particular, Van Dijk et al. (2014) applied this model to interest rates presenting a slow moving trend using exponential smoothing or long survey forecasts. There is also a strand of literature, which connects the level of interest rates in the long run to the expected inflation dynamics, also referred to as trend inflation (Bekaert et al. 2010; Cieslak & Povala, 2015; Kozicki & Tinsley, 2001; Rudebusch & Wu, 2008). Associating the variable of interest with appropriate covariates might practically help in several scenarios, but does not directly address the core of the present problem.

In this paper, we propose a model approach competitive in terms of short horizon forecasts, yet controlled to obtain realistic predictive distributions for the long horizon. We propose a Bayesian TVC-AR(1) model, which is stationary but can resemble short-term properties of an integrated or nearly integrated linear process due to a stochastic mean reversion level. The model allows us to regularize the long run distribution of the time series without affecting short term distributions adversely. Different to Van Dijk et al. (2014) we do not use a deterministic shifting mean reversion level, but incorporate long run assumptions via prior information in a Bayesian approach, such that the latent coefficient process, and with it the mean reversion dynamics, are still estimated from the data.

In particular, the novelty of our approach lies in the model allowing us to (1) regularize long run predictions by using prior assumptions, (2) separate the modeling process into a data driven short horizon model-part and a long horizon model-part that is determined by prior (or expert) assumptions and (3) yield improved forecasting performance in the short horizon compared to the commonly used linear dynamic Nelson-Siegel model and Gauss2++ model, while retaining realistic long run

distributions. Moreover, we place particular emphasis on the interpretability of the model structure and prior parameters, preserving a close link to the common linear models. This allows to include expert knowledge or assumptions in accordance with economic theory about the long run behavior of a time series into the model in a sound mathematical way. We here specifically focus on the application to interest rates. As our model allows to regularize long run predictions, it is also of particular interest for insurance companies, where realistic long run interest rate forecasts are needed to evaluate the risk and performance of specific insurance products.

The remainder of this paper is arranged as follows. Section 2 specifies the Bayesian TVC-AR(1), including the derivation of required full conditional posterior distributions and the application of a Metropolis-Hastings within Gibbs sampling routine for statistical inference. In Section 3 we discuss an application of our model to interest rate data and compare the forecasting performance as well as the long run distribution of our nonlinear model with the dynamic Nelson-Siegel model (short-term focus) and the Gauss2++ model (long-term focus). We conclude with Section 4 and give a brief outlook on potential further research topics.

2 | A BAYESIAN TVC-AR(1) MODEL FOR LONG RUN REGULARIZATION

In this Section we introduce the Bayesian TVC-AR(1) (BTVC-AR(1)) model. The model incorporates assumptions about the long-term behavior of the time series and thereby regularizes the process in the long horizon. At the same time, the model is mainly driven by the given data in the short run and thus fosters a good short-term prediction.

2.1 | The BTVC-AR(1) model

The BTVC-AR(1) model is defined as follows:

$$x_t = \alpha_t + \beta x_{t-1} + \epsilon_t, \quad \text{for } t \in \mathbb{Z}, \quad (4)$$

where β represents the mean reversion speed and $|\beta| < 1$ to secure stationarity. ϵ_t is assumed to be a Gaussian white noise process, i.e., $\epsilon_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2)$. We further specify α_t as a stationary Gaussian process, which is defined to have the unconditional expectation $\boldsymbol{\theta} := \vartheta \cdot \mathbf{1}$ and covariance $\boldsymbol{\Sigma}$ on the observed time points $t = 1, \dots, T$, i.e.,

$$\boldsymbol{\alpha} := (\alpha_1, \alpha_2, \dots, \alpha_T) \sim \mathcal{N}_T(\boldsymbol{\theta}, \boldsymbol{\Sigma}). \quad (5)$$

As this time frame is most relevant, we focus on α for further investigations. The Bayesian approach considers the parameters of model (4) as random variables. For the conditional prior distribution of β conditional on σ^2 a truncated normal distribution with lower bound -1 and upper bound 1 is assumed as a prior, i.e.,

$$\beta|\sigma^2 \sim \mathcal{N}(\mu_\beta, \sigma^2 \cdot \sigma_\beta^2, -1, 1),$$

with conditional prior expectation μ_β and additional multiplicative variance parameter σ_β . The prior distribution for σ^2 is an inverse gamma distribution with shape and scale parameter, a and b , respectively,

$$\sigma^2 \sim \mathcal{IG}(a, b).$$

These two prior distributions are conjugate priors for model (4) if the respective other parameter is known and therefore allow for an analytical derivation of the corresponding full conditional distributions.

Using these priors the defined model can be seen as a Bayesian version of the TVC-AR(1) model. The mean θ and covariance Σ might be assumed fixed or defined as random variables with further attached prior distributions. In the latter case (5) describes the distribution of α conditional on θ and Σ . Placing priors on these parameters allows to incorporate assumptions about the long run distribution into the model as further elaborated in Section 2.2.

While this basic model setup is flexible in many ways and particular in terms of its covariance structure assumptions for α , further practical insights can be obtained from a more in-depth model characterization. In the following, we will shed light on useful properties of this framework when assuming an AR-covariance structure. This covariance structure has shown to provide good results in applications.

2.2 | Arbitrating between short and long run distribution

The goal of our work is to propose a new modeling framework, which can regularize the long run distribution of (nearly) integrated time series by keeping a good forecasting performance in the short horizon. Linear models often concentrate on the conditional distribution in the short horizon, but due to the near integration property of the time series this can lead to inappropriate long run distributions. For example, if the AR(1) model is estimated for a time series, which shows a (close to) random walk behavior, the parameter β of the model will take a value close to 1. This can lead to a large long run variance given by

$$\frac{\sigma^2}{1 - \beta^2},$$

potentially yielding unrealistic values in the long run that have never been observed in the past. On the other hand, calibrating β to a given long run variance is not straightforward without deteriorating the short run prediction performance. Figure 1 depicts this undesired behavior by showing the long run mean of a linear AR(1) model which is driven by the conditional short run distribution at the expense of an unrealistic long-term distribution.

We address this issue by incorporating a time-varying mean reversion level, which locally preserves the good short term prediction and at the same time regularizes the long run distribution. The current mean reversion level valid in the short run can be different to the long run behavior accounting for the current market situation and therefore improving the short run prediction. We enable the model to stay in a reasonable range in the long run by assuming a stationary process for the time-varying mean reversion level and a stronger mean reversion to this time-varying level than a linear AR(1) model would induce to its constant mean reversion level. Such a behavior can be achieved by introducing a time-varying α parameter into a linear AR(1) model with additional prior assumptions. In particular, this does not change the (weak) stationarity property of the model if the assumed process for α is (weakly) stationary itself. This can be verified by calculating the unconditional mean, the unconditional variance and the unconditional covariance:

$$\begin{aligned} E(x_t) &= \frac{\theta}{1 - \beta} \\ \text{Var}(x_t) &= \frac{\sigma^2}{(1 - \beta^2)} + \frac{\text{Var}(\alpha_t) + 2\beta\text{Cov}(\alpha_t, x_{t-1})}{(1 - \beta^2)} \\ \text{Cov}(x_{t-h}, x_t) &= \sum_{i=0}^{h-1} \beta^i \text{Cov}(\alpha_{t-i}, x_{t-h}) + \beta^h \text{Var}(x_{t-h}) \end{aligned}$$

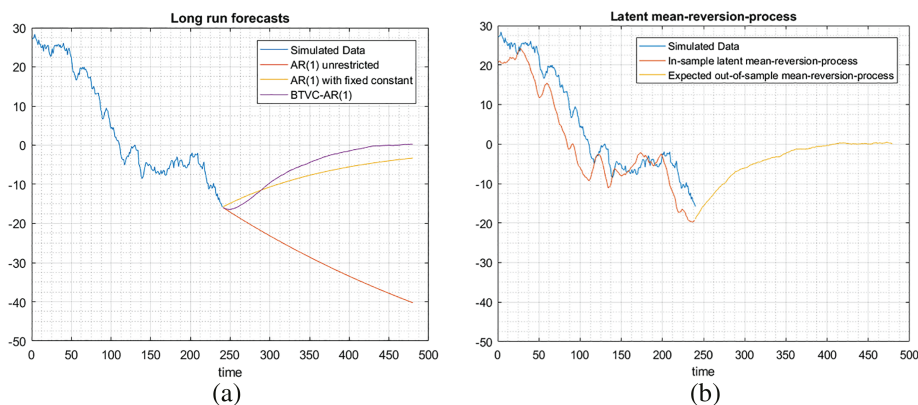
As the α -process is stationary and

$$\text{Cov}(\alpha_t, x_{t-h}) = \sum_{i=0}^{\infty} \beta^i \text{Cov}(\alpha_t, \alpha_{t-h-i}),$$

the BTVC-AR(1) is (weakly) stationary.

The time-varying α increases the flexibility of our model to account for short and long run distributional properties. As current observations have almost no influence in the very long run, a reasonable way to include information about the long run mean and long run variance in a Bayesian setting is via prior assumptions for

FIGURE 1 A comparison of a linear AR(1) model with no restrictions for the constant parameter, a linear AR(1) model with a fixed constant parameter ($\alpha = 0$) and a BTVC-AR(1) model applied on a simulated time series



α . We will further elaborate this in Sections 2.2.1 and 2.2.2. The time-varying α also increases the flexibility of the model such that the conditional distribution in the short run is consistent with the empirical data, i.e., $E(x_{t+h}|x_t, x_{t-1}, \dots)$ and $\text{Var}(x_{t+h}|x_t, x_{t-1}, \dots)$ still reflect the empirical distribution for a short horizon h .

Our BTVC-AR(1) model can therefore produce both a conditional short term distribution, which roughly corresponds to an unrestricted linear model, and a long run distribution with a reasonable range of values.

2.2.1 | The long run mean and time-varying mean reversion

The mean reversion level in a linear AR(1) model as specified in (1) amounts to

$$\frac{\alpha}{1 - \beta}$$

As the mean reversion level stays constant over time it is also the long run mean of the model. In contrast, the mean reversion level in the BTVC-AR(1) model changes over time and is given by

$$\frac{\alpha_t}{1 - \beta}$$

for time point t . This local mean reversion level is in general different to the long-term mean and can even pull the process away from it in expectation, i.e.,

$$|E(x_{t+h}|x_t, x_{t-1}, \dots) - \mu_{LR}| \geq |x_t - \mu_{LR}|,$$

where μ_{LR} is the unconditional mean of the time series and x_t denotes a fixed realisation of the process. This helps fitting the model to a time series exhibiting a (close to) random walk behavior. The long run mean of the

TABLE 1 Estimated parameters for the restricted and unrestricted AR(1) model

	constant	β	σ^2
restricted AR(1)	0	0.9935	0.7787
unrestricted AR(1)	-0.1645	0.9978	0.7554

BTVC-AR(1) depends on the unconditional mean of α and amounts to

$$\frac{\vartheta}{1 - \beta}$$

in our model. We assume the data to be centered around a prior specified long run mean. By setting $\theta = \mathbf{0}$, i.e., $\vartheta = 0$, this long run mean is reached in expectation after reshifting the simulated data.

The implications of the time-varying mean reversion level of the BTVC-AR(1) model are visualized in Figure 1. Two AR(1) models (with unrestricted and restricted constant parameter) and the BTVC-AR(1) model have been exemplary fitted to a simulated stationary time series, which shows a (nearly) integrated behavior and is visualised in Figure 1.

In the left graphic the “historical” time series can be seen as well as the expected future development according to the three models. The AR(1) model with no restrictions has a long-term mean far away from the historical domain, as its focus lies on the conditional short term distribution. The restricted AR(1) model sets the α parameter to 0 to regularize the long run mean, but at the same time the expected values in the short horizon are pushed in the direction of the long run level leading to an inferior forecasting performance. The parameters of the estimated unrestricted and restricted AR(1) model are given in Table 1.

If we assume that the (close to) random walk behavior stems from changes in the mean reversion level

determined by unobserved variables, the BTVC-AR (1) model has a more desired behavior. The time-varying constant parameter in the model leads to a time-varying mean reversion level and can therefore account for the changes induced by the unobserved variables. The long run mean can still be regularized to 0 while influencing the short term distribution less abruptly. This allows the time series to follow the current trend in expectation and veer away from the long run mean for a couple of time steps. The reason for this behavior is that the latent α -process induces a local mean reversion level that lies below the last observation, which can be seen in the right plot of Figure 1 showing the average latent mean reversion level extracted during the simulation process. In the long run the mean reversion level returns in expectation to the prespecified value of 0.

2.2.2 | Long run variance

The long run variance of a linear AR(1) model is given by

$$\frac{\sigma^2}{1-\beta^2}.$$

The closer the model behaves like a random walk, i.e., the closer β approaches 1, the larger the long run variance gets under the assumption of a fixed conditional variance σ^2 . In terms of the long run variance, the BTVC-AR(1) model is more flexible by incorporating two sources of variation, the residual term of the AR(1) model and variance of the latent α -process. The model's long run variance is given by

$$\text{Var}(x_t) = \frac{\sigma^2}{1-\beta^2} + \frac{\text{Var}(\alpha_t) + 2\beta \text{Cov}(\alpha_t, x_{t-1})}{1-\beta^2}. \quad (6)$$

The first term has the same form as the long run variance of a linear AR(1) model and can be interpreted as the “unconditional” variance around the time-varying mean reversion level, i.e., the variance conditional on the α -process. The second term incorporates the part of the variance stemming from the α -process and depends on both its unconditional variance and unconditional covariances. This allows the BTVC-AR(1) model to be more flexible and to control the long run variance of x_t , while reducing the opposing effect on the conditional distribution in the short horizon. The model thus still produces short term distributions consistent with the given data. If α is a constant process, the second term is 0 and the BTVC-AR(1) model reduces to a linear AR(1) model.

Prior assumptions

With the goal in mind to control the long-term variance based on prior information, a more refined specification of the BTVC-AR(1) model is helpful in order to translate this information into the model. We will use a centered α -process with an AR-covariance structure for demonstrative purposes. In this case, α can be represented by a linear AR(1) model

$$\alpha_t = \rho\alpha_{t-1} + \eta_t, \quad \text{for } t \in \mathbb{Z},$$

where ρ represents the correlation between two successive time steps and η_t is an i.i.d. Gaussian white noise process, i.e., $\eta_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \tau^2)$. The long run variance of the BTVC-AR(1) model is then given by

$$\text{Var}(x_t) = \frac{\sigma^2}{(1-\beta^2)} + \frac{\tau^2(1+\rho\beta)}{(1-\rho\beta)(1-\beta^2)(1-\rho^2)}. \quad (7)$$

As indicated by the equation, τ^2 and σ^2 play an inter-related role in the model for the long run variance. Increasing one can be compensated by reducing the other one. To avoid identifiability issues, it is thus necessary to account for this interrelation through a meaningful prior parameter assumption. If the process x_t is supposed to reach a certain objective variance in the long run, the degrees-of-freedom in (7) reduce from four to three. For example, for given ρ , β and σ^2 and a prior value assumption for $\text{Var}(x_t)$, the variance of x_t has a one-to-one relationship with τ^2 and it is straightforward to solve (7) for τ^2 . Denote the solution by $\tilde{\tau}^2$. To ensure positivity the truncation limits for the prior distribution of β can be set to -1 and $\sqrt{\frac{\text{Var}(x_t) - \sigma^2}{\text{Var}(x_t)}}$. For this specific covariance structure, a possible prior distribution of τ^2 can thus be defined by the conditional distribution

$$\tau^2 | \rho, \beta, \sigma^2 \sim \delta_{\tilde{\tau}^2}, \quad (8)$$

where δ denotes a degenerated distribution with point mass 1 at $\tilde{\tau}^2$. This definition forces the process to reach its prespecified long run variance $\text{Var}(x_t)$ while controlling the speed of mean reversion of the α -process through ρ . A conjugate prior for ρ is a normal distribution truncated below by -1 and from above by 1, i.e.,

$$\rho \sim \mathcal{N}(\mu_\rho, \sigma_\rho^2, -1, 1), \quad (9)$$

with mean μ_ρ and variance σ_ρ^2 .

The previous prior specifications allow to introduce prior information into the model in a straightforward

manner while maintaining the properties of the BTVC-AR(1) model.

2.2.3 | The short run distribution

For the short run distribution of the BTVC-AR(1) model the goal is to balance between a consistent estimation with the observed data and the opposing effect of the prespecified long run distribution. For a linear AR(1) model with a restricted long run mean of 0 the conditional expectation and the conditional variance amount to

$$\begin{aligned} E(x_{t+1}|x_t, \dots) &= \beta x_t, \\ \text{Var}(x_{t+1}|x_t, \dots) &= \sigma^2. \end{aligned}$$

The model can get arbitrarily close to a centered random walk if β approaches 1, while the long run variance increases at the same time as shown in Section 2.2.2. For the BTVC-AR(1) model we get

$$\begin{aligned} E(x_{t+1}|x_t, \dots) &= E(\alpha_{t+1}|x_t, \dots) + \beta x_t, \\ \text{Var}(x_{t+1}|x_t, \dots) &= \text{Var}(\alpha_{t+1}|x_t, \dots) + \sigma^2. \end{aligned}$$

A random walk behavior, i.e., $E(x_{t+1}|x_t) \approx x_t$, can be reached without β necessarily being close to 1 due to the conditional expectation of the α -process that supports the random walk behavior in the short horizon. This increases the flexibility of the BTVC-AR(1) model compared to a linear AR(1) model in combining short and long run distributional characteristics.

We can further decompose the conditional expectation to see the similarities of the BTVC-AR(1) model to a linear AR(T) process if we consider a time series observed on time points $t=0, \dots, T$. Let $\check{\alpha} = (\alpha_1, \dots, \alpha_{T+1})$ denote the time-varying constant extended to $T+1$ in a consistent manner with the BTVC-AR(1) model definition, i.e., the same covariance parameterization is assumed. For a given data set $\mathbf{x} = (x_0, \dots, x_T)$, the conditional distribution of $\check{\alpha}|\mathbf{x}$ is multivariate normal (c.f. Appendix 2.1), i.e.,

$$\check{\alpha}|\mathbf{x} \sim \mathcal{N}(\check{\boldsymbol{\mu}}, \check{\boldsymbol{\Sigma}}).$$

As $\check{\boldsymbol{\mu}} = \frac{1}{\sigma^2} \check{\boldsymbol{\Sigma}} \check{\boldsymbol{\Delta}}$ with $\check{\boldsymbol{\Delta}} = (x_1 - \beta x_0, \dots, x_T - \beta x_{T-1}, 0)^\top$, the conditional expectation of α_{T+1} is given by the last entry of $\check{\boldsymbol{\mu}}$,

$$E(\alpha_{T+1}|\mathbf{x}) = \frac{1}{\sigma^2} \mathbf{s}_{T+1} \check{\boldsymbol{\Delta}},$$

where $\mathbf{s}_{T+1} = (s_{T+1,1}, \dots, s_{T+1,T+1})$ and $s_{i,j}$ represent the entries of $\check{\boldsymbol{\Sigma}}$. The one step ahead conditional expectation of the model therefore amounts to

$$\begin{aligned} E(x_{T+1}|x_T, \dots) &= \left(\frac{s_{T+1,T}}{\sigma^2} + \beta \right) x_T \\ &+ \sum_{i=1}^{T-1} \frac{s_{T+1,T-i} - \beta s_{T+1,T-(i-1)}}{\sigma^2} x_{T-i} \\ &- \frac{s_{T+1,1} \beta}{\sigma^2} x_0. \end{aligned}$$

This shows that the conditional expectation depends on all previous time points like in a linear AR(t) model, allowing the BTVC-AR(1) model to better account for current trends in the process. Due to the given covariance structure for $\boldsymbol{\alpha}$ the number of parameters are, however, much less than in an actual AR(t) process.

2.3 | Bayesian inference

The main parameters of interest in the BTVC-AR(1) model are $\tilde{\boldsymbol{\alpha}}$, β and σ^2 with $\tilde{\boldsymbol{\alpha}}$ extending $\boldsymbol{\alpha}$ by future time points up to the modeling horizon h , i.e.,

$$\tilde{\boldsymbol{\alpha}} = (\alpha_1, \dots, \alpha_T, \dots, \alpha_{T+h}).$$

This extension is necessary to sample from the predictive posterior distribution of the parameters and to generate forecasts. The prior distribution of $\tilde{\boldsymbol{\alpha}}$ incorporates the same assumptions as the prior distribution of $\boldsymbol{\alpha}$, i.e.,

$$\tilde{\boldsymbol{\alpha}} \sim \mathcal{N}_{T+h}(\tilde{\boldsymbol{\theta}}, \tilde{\boldsymbol{\Sigma}})$$

where

$$\tilde{\boldsymbol{\theta}} = \vartheta \cdot \mathbf{1} \text{ and } \tilde{\boldsymbol{\Sigma}} = \begin{pmatrix} \boldsymbol{\Sigma} & \boldsymbol{\Pi}_{T+1} & \dots & \boldsymbol{\Pi}_{T+h} \\ \boldsymbol{\Pi}_{T+1}^\top & \sigma_\alpha^2 & & \\ \vdots & & \ddots & \\ \boldsymbol{\Pi}_{T+h}^\top & & & \sigma_\alpha^2 \end{pmatrix}$$

with $\boldsymbol{\Pi}_{T+j} = \{\text{Cov}(\alpha_{T+j}, \alpha_1), \dots, \text{Cov}(\alpha_{T+j}, \alpha_{T+j-1})\}^\top$, i.e., the vector of covariances of α_{T+j} and all previous time points $1, \dots, (T+j-1)$. For these time points the same (autoregressive) covariance parameterization as for $\boldsymbol{\alpha}$ is assumed for consistency reasons. σ_α^2 represents the unconditional variance of the latent α -process. In the case of an AR-covariance structure as assumed before $\sigma_\alpha^2 = \frac{\tau^2}{1-\rho^2}$ and $\text{Cov}(\alpha_s, \alpha_{s+k}) = \rho^k$ for any $s, k \in \mathbb{Z}$.

The goal of Bayesian inference is to find the joint posterior distribution, $p(\tilde{\boldsymbol{\alpha}}, \beta, \sigma^2|\mathbf{x})$, conditional on the

observed data $\mathbf{x} = (x_0, \dots, x_T)$. If the full conditional distribution of all parameters is known, the Gibbs sampler (Gelman et al., 2013) can be used to draw samples from this joint posterior distribution and inference can be based on Monte Carlo approximation (Chib, 2001). By regularizing the long run variance under the assumption of an AR-covariance structure and choosing a degenerated prior distribution for τ^2 as in (8), the full conditional distributions of ρ, β and σ^2 depend on the prior of τ^2 and can not be derived analytically. We therefore apply a Metropolis-Hastings within Gibbs sampling routine (Millar & Meyer, 2000). We will state the algorithmic details in the following section and here only derive the necessary distributions.

As the model defined in Section 2.2.2 can be considered under a different parameterization where τ^2 is given by the function

$$\tau^2 = f(\rho, \beta, \sigma^2, \text{Var}(x_t))$$

and thus fixed for given ρ, β, σ^2 and a specified long run variance $\text{Var}(x_t)$, we will focus on deriving two conditional distributions in order to be able to employ a two-step Gibbs sampling procedure. The goal is to iteratively sample α and the vector (ρ, β, σ^2) based on the respective other full conditional distribution. As it is not straightforward to derive the conditional distribution for the latter vector, we will here derive conditional distributions for all parameters involved as if the parameter τ^2 was fixed and later employ these distributions to derive a suitable proposal distribution in a Metropolis-Hastings procedure. In the following subsections we just state the (full) conditional distributions. A more detailed derivation can be found in Appendices 2.1–2.4.

2.3.1 | Full conditional distributions of α

In the following we derive the full conditional distribution of $\tilde{\alpha}$. It holds

$$p(\tilde{\alpha}|\beta, \sigma^2, \mathbf{x}) \propto p(\mathbf{x}|\tilde{\alpha}, \beta, \sigma^2) \cdot p(\tilde{\alpha}) = \mathcal{L}(\tilde{\alpha}, \beta, \sigma^2) \cdot p(\tilde{\alpha}). \quad (10)$$

Due to the conditional independence induced by the Markov assumption in the AR(1) model the likelihood of the parameters is given by

$$\begin{aligned} \mathcal{L}(\tilde{\alpha}, \beta, \sigma^2) &= p(\mathbf{x}|\alpha, \beta, \sigma^2) \\ &= \prod_{j=0}^{T-1} \phi(x_{T-j}|\alpha_{T-j} + \beta x_{T-j-1}, \sigma^2), \end{aligned} \quad (11)$$

where $\phi(\cdot|\mu, \sigma^2)$ denotes the density function of a normal distribution with expectation μ and variance σ^2 . Note,

that we have assumed a degenerated distribution with point mass 1 for the first entry in \mathbf{x} . An alternative option is to estimate the unconditional distribution. For increasing length of the time series the difference between these two approaches will however vanish.

With (10) and (11) and the prior distributions specified in Section 2.1 the full conditional distributions of $\tilde{\alpha}$, can be derived analytically. Under the assumption that $\tilde{\theta} = \mathbf{0}$ as specified in Section 2.2.1 to regularize the long run mean, the full conditional distribution of $\tilde{\alpha}$ is given by

$$\tilde{\alpha}|\beta, \sigma^2, \mathbf{x} \sim \mathcal{N}_T(\tilde{\boldsymbol{\mu}}_{\text{post}}, \tilde{\boldsymbol{\Sigma}}_{\text{post}}).$$

with

$$\begin{aligned} \tilde{\boldsymbol{\mu}}_{\text{post}} &= \tilde{\boldsymbol{\Sigma}}_{\text{post}} \tilde{\Delta} \frac{1}{\sigma^2} \text{ and} \\ \tilde{\boldsymbol{\Sigma}}_{\text{post}} &= \left(\tilde{\boldsymbol{\Sigma}}^{-1} + \frac{1}{\sigma^2} \begin{pmatrix} \mathbf{I}_T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \right)^{-1}. \end{aligned}$$

$\tilde{\Delta}$ in this case denotes

$$\tilde{\Delta} = (x_2 - \beta x_1, \dots, x_T - \beta x_{T-1}, 0, \dots, 0).$$

As $\tilde{\Delta}$ incorporates data information up to time point (vector entry) T , is 0 for time points $>T$ and $\text{Cov}(\alpha_{T+j}, \alpha_T) \rightarrow 0$ with increasing j , the mean of the full conditional distribution tends to 0, corresponding to the unconditional mean of the prior distribution. The covariance structure of the full conditional distribution behaves analogously. Therefore, the distribution of $\alpha_{T+j}|\mathbf{x}, \beta, \sigma^2$ in the long run tends to the prior distribution. This means that the prior distribution of α effectively regularizes the distribution of \mathbf{x} in the long horizon towards the prespecified long run mean and long run variance.

Note that the derivations are independent of the specific choice of $\tilde{\boldsymbol{\Sigma}}$. If prior distribution assumptions for the parameters in $\tilde{\boldsymbol{\Sigma}}$ are used, we need to further condition on the hyper-parameters for the full conditional distribution of $\tilde{\alpha}$. In the following, we assume an AR-covariance structure for $\tilde{\alpha}$ and therefore need to condition on the parameters ρ and τ^2 .

We note that, in general, regarding the α process as a vector of parameters can lead to a computational burden when calculating the inverse of $\tilde{\boldsymbol{\Sigma}}$ for $\tilde{\boldsymbol{\Sigma}}_{\text{post}}$. In this case it might be preferable to represent the model in state space and to use a forward filtering backward sampling algorithm as proposed by Carter and Kohn (1994) and Frühwirth-Schnatter (1994). In our specific case, however, we initialize a specific parametric covariance structure and did not experience any computational problems.

2.3.2 | Full conditional distributions of ρ, β, σ^2

If we assume an AR-covariance structure with prior distributions for its parameters as specified in Section 2.2.2, the conditional distribution of ρ is given by

$$\rho | \boldsymbol{\alpha}, \tau^2 \sim \mathcal{N}(\mu_{\rho, \text{post}}, \sigma_{\rho, \text{post}}^2, -1, 1)$$

where

$$\begin{aligned} \sigma_{\rho, \text{post}}^2 &= \left(\frac{\sum_{j=0}^{T-1} \alpha_{T-j-1}^2}{\tau^2} + \sigma_{\rho}^{-2} \right)^{-1} \\ \mu_{\rho, \text{post}} &= \left(\frac{\sum_{j=0}^{T-1} \alpha_{T-j} \alpha_{T-j-1}}{\tau^2} + \frac{\mu_{\rho}}{\sigma_{\rho}^2} \right) \sigma_{\rho, \text{post}}^2. \end{aligned}$$

The conditional distribution of β is given by

$$\beta | \mathbf{x}, \boldsymbol{\alpha}, \sigma^2 \sim \mathcal{N} \left(\mu_{\beta, \text{post}}, \sigma_{\beta, \text{post}}^2, -1, \sqrt{\frac{\text{Var}(x_t) - \sigma^2}{\text{Var}(x_t)}} \right)$$

where

$$\begin{aligned} \sigma_{\beta, \text{post}}^2 &= \left(\frac{\sum_{j=0}^{T-1} x_{T-j-1}^2}{\sigma^2} + (\sigma \sigma_{\beta})^{-2} \right)^{-1} \\ \mu_{\beta, \text{post}} &= \left(\frac{\sum_{j=0}^{T-1} \check{d}_{T-j} x_{T-j-1}}{\sigma^2} + \frac{\mu_{\beta}}{\sigma^2 \sigma_{\beta}^2} \right) \sigma_{\beta, \text{post}}^2. \end{aligned}$$

\check{d}_{T-j} is defined by $\check{d}_{T-j} := x_{T-j} - \alpha_{T-j}$.

The conditional distribution of σ^2 is given by an inverse gamma distribution with parameters

$$\begin{aligned} \tilde{a} &= \frac{T+1}{2} + a \quad \text{and} \\ \tilde{b} &= \frac{\sum_{j=0}^{T-1} \epsilon_{T-j}^2}{2} + b + \frac{(\beta - \mu_{\beta})^2}{2\sigma_{\beta}}, \end{aligned}$$

where ϵ_t is the error term of the BTVC-AR(1) model (c.f. (4)). This means

$$\sigma^2 | \tilde{\boldsymbol{\alpha}}, \beta, \mathbf{x} \sim \mathcal{IG}(\tilde{a}, \tilde{b}).$$

Note that this only holds if the prior of $\beta | \sigma^2$ is a normal distribution instead of a truncated normal distribution as assumed in Section 2.1. When using a truncated distribution assumption, the derivation of the full conditional of σ^2 is more intricate as the prior distribution of β also conditions on σ^2 . Since our approach will make use of the full conditionals as proposal distributions in

the Metropolis-Hastings part of our sampling routine, this simplification allows a more straightforward implementation while we observe that values outside the given truncation are highly unlikely and practically occur with zero probability in our application.

2.4 | Markov Chain Monte Carlo inference

In the following we assume again an AR-covariance structure for $\boldsymbol{\Sigma}$ determined by the parameters ρ and τ^2 with prior distributions as specified in Section 2.2.2. To conduct inference, we use the Metropolis-Hastings within Gibbs sampler. More specifically, we generate samples from the posterior distribution by iteratively sampling from the full conditional distribution of $\tilde{\boldsymbol{\alpha}}$ given a sample of (ρ, β, σ^2) and vice versa. Based on the derivation of the full conditional distribution for $\boldsymbol{\alpha}$ in the previous section we are able to directly sample from a multivariate normal distribution to generate values for $\boldsymbol{\alpha}$. To obtain a sample from $p(\rho, \beta, \sigma^2 | \boldsymbol{\alpha}, \mathbf{x})$ conditional on $\boldsymbol{\alpha}$, we apply the Metropolis-Hastings algorithm as neither the joint distribution of ρ, β, σ^2 nor each single full conditional distribution is available. A suitable and already available proposal distribution q for these parameters is given by

$$q(\rho, \beta, \sigma^2 | \boldsymbol{\alpha}, \mathbf{x}) = q(\rho | \boldsymbol{\alpha}, \mathbf{x}) q(\beta | \boldsymbol{\alpha}, \sigma^2, \mathbf{x}) q(\sigma^2 | \boldsymbol{\alpha}, \beta, \mathbf{x}). \quad (12)$$

In other words, we use the product of all full conditional distributions under the assumption of a fixed τ^2 .

In the BTVC-AR(1) model we use this approach in a first step to draw from the joint posterior distribution $p(\tilde{\boldsymbol{\alpha}}, \beta, \sigma^2 | \mathbf{x})$. A detailed description of the sampling routine can be found in Appendix A1. In a second and final step, we use these samples to generate paths of the x -process as follows:

$$x_{T+j}^{(m)} = \alpha_{T+j}^{(m)} + \beta^{(m)} x_{T+j-1}^{(m)} + \epsilon_{T+j}, \quad j > 0,$$

where m denotes the simulated path. In Appendix B1, we give some insights about the performance and further details about the MCMC algorithm applied to our dataset.

3 | APPLICATION TO INTEREST RATE DATA

We now apply the BTVC-AR(1) model to the first principal component (PC) of a principal component analysis (PCA) on interest rate data to predict the term structure

of interest rates and compare it to the *2-Additive-Factor Gaussian (Gauss2++) model* (Brigo & Mercurio, 2007) and the *dynamic Nelson-Siegel model* (Diebold & Li, 2006) with respect to the forecasting performance and the long run distribution.

3.1 | Motivation and background

The Gauss2++ model is a popular short-rate model in the insurance industry, used, e.g., to classify certified pension contracts into risk classes. Because its mean reversion level is calibrated to external interest rate forecasts, it generates realistic interest rates in the long horizon, which is a necessary model feature for insurance companies, as they are obliged to calculate risk measures and performance scenarios for specific insurance contracts for up to 40 years (European Union, 2017).

Nevertheless, Diebold and Li (2006) point out that short-rate models perform poorly in forecasting. Their dynamic Nelson-Siegel model shows a better forecasting performance than the Gauss2++ model in the short horizon, but can produce unrealistic interest rates in the very long horizon. Our model, which we call the *BTVC-AR(1)-Factor model* in the following as it applies the BTVC-AR(1) model to the first PC of a PCA, combines both: a good forecasting performance in the short horizon and realistic interest rates in the long horizon. It further accounts for the strong autocorrelation and the (close to) random walk behavior of interest rates.

3.2 | Data

We use data of the German term structure of interest rates estimated by the Deutsche Bundesbank from prices of German government bonds. The exact estimation procedure can be found in Schich (1997). The time span ranges from September 1997 to August 2016. Figure 2 shows the monthly evolution of the interest rate curves.

In the last ten to fifteen years a decrease of the interest rates can be observed. Each maturity represents a dimension in the data set. We use PCA to reduce the dimension of the data set for the following reason. According to Litterman and Scheinkman (1991) a three factor model can explain for each interest rate with a specific maturity a minimum of 96% of the variability in the data. We here extract these (principle) factors but only use the first two to facilitate a fair comparison with the Gauss2++ model, which is a two factor model. Furthermore, the first two PCs already account for more than 99% of the variability in the given data. Figure 3 shows the loadings and the time series of the two extracted PCs.

The loadings of the first PC are similar for all 20 maturities, while the loadings of the second PC are positive for

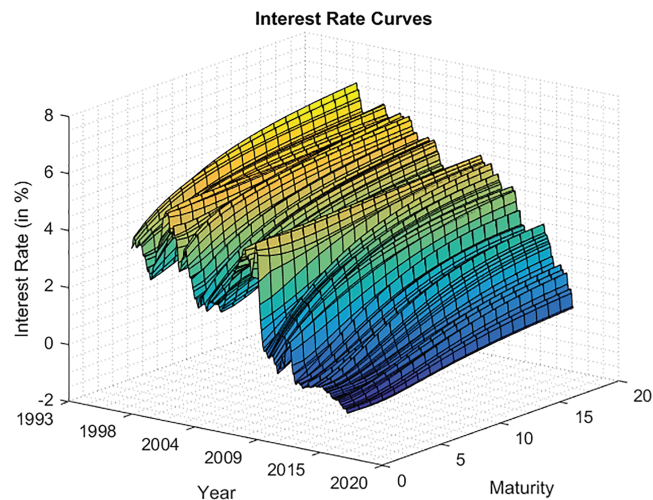


FIGURE 2 Time series of the term structure of German government bond yields




FIGURE 3 The scores and the loadings of the first two PCs

short and negative for long maturities. The first and the second PC are therefore often interpreted as level and slope of the term structure, respectively.

The decrease of the interest rates in the last years is also visible in the level factor, showing a downward trend. There is an ongoing discussion in the literature about mean reversion of interest rates. Economic theory predominantly assumes that interest rates are (in the long run) mean reverting. But statistical evidence is not so clear (van den End, 2011). The mainstream literature says that unit roots can not be rejected, which would imply that interest rates are not mean reverting (Campbell & Shiller, 1991; Rose, 1988; Siklos & Wohar, 1997; Stock & Watson, 1988). More recent literature investigates the unit root hypothesis by fractional integrated techniques that apply differencing to time series by an order smaller than or greater than one (Baum et al. 2000; Gil-Alana, 2004). These studies find that shocks to interest rates have a long memory, which explains their (close to) random walk behavior.

3.3 | Estimation of model parameters

In this subsection the estimation of the BTVC-AR(1)-Factor model and the two benchmark models is described.

3.3.1 | Modeling interest rates with the BTVC-AR(1)-Factor Model

The factors of our BTVC-AR(1)-Factor model are the first two PCs extracted by a PCA and interpreted as level and slope of the interest rate curve. The level factor shows a (close to) random walk behavior, which can not be adequately captured by a stationary linear model. Following the economic theory view that interest rates (and therefore also the level) are mean reverting (in the long run) and assuming that the random walk behavior results from changes in the mean reversion level, we use therefore the BTVC-AR(1) model for this PC. It allows us to account for the (close to) random walk behavior as well as to regularize the level of the interest rate curve in the long horizon via prior assumptions. The slope factor is more stable over time. As an augmented Dickey Fuller test suggests that the existence of a unit root can be rejected, a linear AR(1) model is used for this factor. By modeling the level and the slope factor interest rate forecasts $\hat{r}_t(\tau)$ with maturity τ can be calculated via

$$\hat{r}_t(\tau) = \mu(\tau) + \xi_1(\tau)\hat{l}_t + \xi_2(\tau)\hat{s}_t, \quad (13)$$

where \hat{l}_t and \hat{s}_t denote the forecasts of the level and the slope factor, respectively. $\xi_1(\tau)$ and $\xi_2(\tau)$ denote the loading of the first and second PC for maturity τ . Before applying the PCA the data has been centered and therefore $\mu(\tau)$ is the mean interest rate of the data set for maturity τ . We now specify the prior assumptions of the BTVC-AR(1) model for the level factor and the estimation procedure of the AR(1) model for the slope factor.

The level factor

Latent AR1 parameter: For this application, we assume an AR-covariance structure for the α -process of the BTVC-AR(1) model with the parameters ρ and τ^2 representing the correlation of two successive time points and the conditional variance, respectively. The unconditional mean of the α -process is set to 0, which implies the assumption that the long run mean of the level factor is 0. Because we also assume that the slope factor is a centered process this means that the long run interest rate curve converges in expectation to the average interest rate curve of the dataset.

Autocorrelation parameter: As specified in Section 2.2.2, we assume for ρ a truncated normal distribution with the parameters $\mu_\rho = 0.98$ and $\sigma_\rho^2 = 0.001^2$ with lower truncation -1 and upper truncation 1 as a hyper prior, i.e.,

$$\rho \sim \mathcal{N}(0.98, 0.001, -1, 1).$$

The truncation ensures the stationarity of the process. The parameters of this hyper-prior rely on expert judgment and incorporate the assumption of a weak mean reverting α -process into the model and therefore allow the mean reversion level of the level factor to deviate from the long run mean for longer periods. This yields the (close to) random walk behavior present in (our) interest rate data.

Variance of the latent process: According to Section 2.2.2, the parameter τ^2 is set in each iteration of the sampling procedure such that the long run variance of the level factor amounts to a prespecified value. We here use the value 120, which is inferred from a quantile of the unconditional distribution. By giving consideration of the rather unusual market situation of extremely low interest rates we make the assumption that the last observation is equal to the 7.5%-quantile. Due to the model assumptions, the unconditional distribution is normal with mean 0 and the corresponding unconditional variance can be calculated easily.

Slope parameter of the AR(1) model: For β , we assume that $\mu_\beta = 0.95$ and $\sigma_\beta^2 = 0.015^2$. This expert judgment represents a weak mean reversion to the time-varying mean reversion level. The lower and upper

truncation of the truncated normal distribution amount to -1 and $\sqrt{\frac{\text{Var}(x_t) - \sigma^2}{\text{Var}(x_t)}}$ to ensure the stationarity of the model as well as the positivity of τ^2 , i.e.,

$$\beta|\sigma^2 \sim \mathcal{N}\left(0.95, \sigma^2 0.015^2, -1, \sqrt{\frac{\text{Var}(x_t) - \sigma^2}{\text{Var}(x_t)}}\right)$$

Residual variance: For the prior distribution of σ^2 the shape and scale parameter a and b are set to 0.5 and 2 respectively, representing an uninformative prior.

By specifying the parameters of the prior (and hyper-prior) distributions the full conditional distribution of $\tilde{\alpha}$ as well as the conditional distributions of the other parameters can be analytically derived as described in Section 2.3. Combining the Gibbs Sampler and the Metropolis-Hastings algorithm as explained in Section 2.4, paths of the level factor can be generated. Forecasts of the level factor are then represented by the average of the simulated paths.

The slope factor

The linear AR(1) model for the slope factor is given by

$$s_t = c + \gamma s_{t-1} + \eta_t,$$

where γ is a real valued constant between -1 and 1 and η_t is a Gaussian white noise process, i.e., $\eta_t \sim \mathcal{N}(0, \sigma^2)$. The constant parameter c is set to 0. The other parameters are estimated by a standard ordinary least squares approach.

3.3.2 | Modeling interest rates with the Gauss2++ model

The Gauss2++ model—in a different representation also known as the 2-Factor-Hull-White model—is a popular interest rate model in the insurance industry used for pricing interest rate derivatives as well as for risk management and forecasting purposes. The model assumes that the short-rate $r(t)$, which is the interest rate with an infinitesimal small maturity, is given by the sum of two latent processes $(x(t))_{t \geq 0}$ and $(y(t))_{t \geq 0}$, and a deterministic function φ :

$$r(t) = x(t) + y(t) + \varphi(t).$$

The latent processes are modeled by dependent Ornstein-Uhlenbeck processes, which are the continuous version of a linear AR(1) process. Interest rates with longer maturities are then derived from the short-rate via pricing the corresponding zero-coupon bonds, which is

analytically possible due to the model's distributional assumptions.

The estimation process is materially different from the one of the other two models as it does not use historic data but calibrates the model to current future market assumptions (implicitly) provided by the current interest rate curve, interest rate derivatives as well as interest rate forecasts. By applying the downhill simplex algorithm the parameters of the model are chosen in such a way that forward rates—implicitly given by the current interest rate curve—and swaption prices are met in expectation. The relevant data has been extracted from Bloomberg. Additionally the mean reversion level of the two latent factors are analytically set such that two interest rate forecasts with a maturity of 3 months and 10 years, which are published by the Organisation for Economic Co-operation and Development (OECD), are met in expectation. This approach is in line with the standard calibration procedure in the insurance industry (Korn & Wagner, 2019).

3.3.3 | Modeling interest rates with the dynamic Nelson-Siegel model

The dynamic Nelson-Siegel model of Diebold and Li applies specific time series models to extracted latent factors. Diebold and Li tested several time series models on the level, slope and curvature factors of the Nelson-Siegel interest rate curve and compared the forecasting performance Diebold and Li (2006). In this paper, we follow one of their approaches, in which they apply a PCA on interest rate data and use an univariate linear AR(1) process for each of the first three PCs. Because of comparison reasons to the other two two-factor models in this paper, we just use the first two PCs. The parameters of the AR(1) model are estimated by the ordinary least squares method.

3.4 | Backtest

We now compare the forecasting performance of the BTVC-AR(1)-Factor model, the Gauss2++ model and the dynamic Nelson-Siegel model and analyze their long run distributions of the 10-year interest rate (Figure 4).

3.4.1 | Comparison of the forecasting performance

For the out-of-sample backtest, we apply an expanding window approach. The data of the first 10 years of the

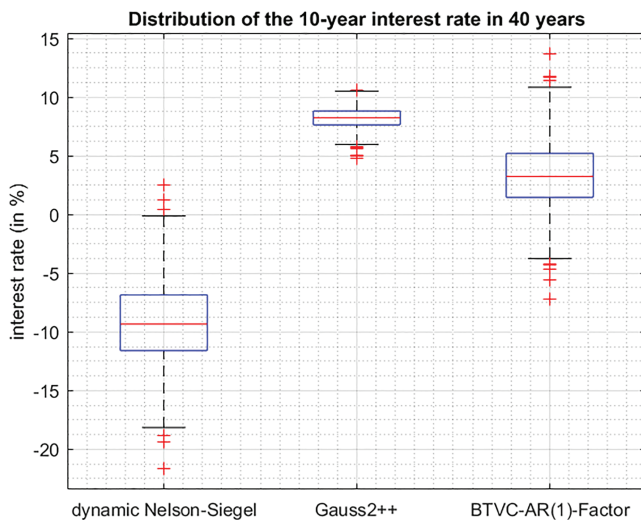


FIGURE 4 Comparison of the distributions of the 10-year interest rate in 40 years modeled by the dynamic Nelson-Siegel model, the Gauss2++ model and the BTVC-AR(1)-Factor model

observations are used to estimate the parameters of the BTVC-AR(1)-Factor model and the dynamic Nelson-Siegel model as described in the Section 3.3. The Gauss2++ model is calibrated to the current market data. We then forecast the interest rates for the maturities of 1, 3, 5, and 10 years (representing the yield curve) for the horizons of 1, 3, 6, and 12 months. We expand the training sample by one month and repeat the procedure again. This is done until 12 months before the last observation in the data set. To evaluate the forecasting performance the error between the predicted interest rate $\hat{r}_\tau(t)$ and the actual interest rate $r_\tau(t)$ with the maturity τ is calculated, i.e.,

$$\text{error}_\tau(t) = r_\tau(t) - \hat{r}_\tau(t).$$

Tables D1–D4 in the Appendix C1 show the mean and the standard deviation of this error for each model. In addition, the root mean squared error

$$\text{RMSE}(\tau) = \sqrt{\frac{1}{N} \sum_{k=1}^N (r_\tau(k) - \hat{r}_\tau(k))^2} \quad (14)$$

for the given deviation is calculated, where N is the number of forecasts conducted in the backtest.

The RMSE for the 1-month ahead forecasts is similar for all three models. For longer forecasting horizons the Gauss2++ model shows the highest RMSE. For example, the 6-month ahead forecast of the 10-year interest rate of the Gauss2++ model has a RMSE approximately twice as high as the RMSE of the other two models and more

than three times as high for the 12-month ahead forecast. This supports the statement of Diebold and Li (2006) that short-rate models perform poorly in forecasting. However, it should be mentioned that the performance of the Gauss2++ model highly depends on the interest rate forecasts used in the calibration process. Regarding the predominant negative mean error suggests that the OECD forecasts have been too optimistic in the past.

The results of the BTVC-AR(1)-Factor model and the dynamic Nelson-Siegel model are more consistent. For the forecasting horizon of 1-month the BTVC-AR(1)-Factor model shows a slightly lower RMSE except for the 10-year interest rate. For the 3, 6, and 12 months forecasting horizons the BTVC-AR(1) model shows a lower RMSE for the short maturities, but a higher RMSE for the longer maturities compared to the dynamic Nelson-Siegel model. Note that the dynamic Nelson-Siegel model anticipated the downward trend present in the last years, which might have been beneficial in terms of the forecasting performance in the past, but also produces unrealistic interest rates in the long horizon. In contrast the BTVC-AR(1)-Factor model forces the model to mean revert to a prespecified level to regularize the interest rates in the long horizon. It can therefore follow the current trend only for a couple of time steps, which might explain the slightly worse performance for the 6 and 12 months forecasting horizon. The fact that the RMSE error is still similar to the dynamic Nelson-Siegel model suggests that this does not affect the forecasting performance in the short horizon much.

3.4.2 | Comparison of the distribution in the long run

We further investigate the interest rate distribution in the long horizon. This is especially important for insurance companies as risk measures and performance scenarios for their products have to be calculated for up to 40 years (European Union, 2017). We therefore fit all three models on all data points up to the last observation date of the data set. We then simulate paths of the 10-year interest rate and visualize the distribution in 40 years. The median of the dynamic Nelson-Siegel model amounts to approximately -10% . A value that is not realistic for the 10-year interest rate. In comparison, the distribution of the BTVC-AR(1)-Factor model and the Gauss2++ model seem to be more realistic as the range of their distributions is (mainly) positive between 0% and 10% . It can be observed that the standard deviation of the Gauss2++ model is much smaller than of the BTVC-AR(1)-Factor model and as the median is quite high negative values are not reached by this model. This is due to the

fact that the Gauss2++ model assumes a stronger mean reversion than historic data would suggest. The (close to) random walk behavior is better captured by the BTVC-AR(1)-Factor model leading to a prediction range which fits historical observations quite well. This is due to the regularization of the mean and the standard deviation of the BTVC-AR(1)-Factor model induced by appropriate prior assumptions, which represents the main difference to other interest rate models.

4 | CONCLUSION

In this paper, we introduced a new Bayesian framework for the TVC-AR(1) model particularly suitable for nearly integrated time series which can not be estimated by a linear model consistent with economic theory or historical observations. In these cases a (close to) random walk behavior can be an indication for a missing variable, for which we account for by the usage of a nonlinear model. The time-varying constant of the BTVC-AR(1) allows a stochastic mean reversion level leading to realizations, which exhibit a random walk behavior although being stationary and do not have an exploding long run variance. Additionally, with the Bayesian approach it is possible to incorporate prior assumption about the long run distribution into the model without affecting the short-term predictions adversely. This gives the possibility to include expert knowledge or well known economic facts about the long-term behavior of the time series into the model that is otherwise fully data-driven in the short term forecast.

We apply the proposed approach to interest rate data. We find that the BTVC-AR(1)-Factor model, which applies a BTVC-AR(1) model to the first PC of a PCA, shows a similar forecasting performance as the dynamic Nelson-Siegel model in the short horizon but in contrast produces realistic interest rates in the very long horizon and also yields better forecasts compared to the Gauss2++ model.

The presented framework allows for many different specifications and is, in particular, flexible in terms of the assumed covariance structure of the latent α process in the model. In this paper we propose an AR-covariance structure and explain how model parameters can be inferred in this special case. Investigating other covariance structures may further improve the forecasting performance in the short horizon while still regularizing the distribution in the long run.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available at <https://www.bundesbank.de/de/statistiken/>

zeitreihen-datenbanken. These data were derived from the following resources available in the public domain: BBK01.WZ9801, BBK01.WZ9802, BBK01.WZ9803, BBK01.WZ9804, BBK01.WZ9805, BBK01.WZ9806.

REFERENCES

- Bacon, D. W., & Watts, D. G. (1971). Estimating the transition between two intersecting straight lines. *Biometrika*, 58(3), 525–534.
- Baum, C. F., Bekdache, B., et al. (2000). A re-evaluation of empirical tests of the fisher hypothesis. *Working Papers in Economics*, 1–24.
- Bekaert, G., Cho, S., & Moreno, A. (2010). New Keynesian macroeconomics and the term structure. *Journal of Money, Credit and Banking*, 42(1), 33–62.
- Brigo, D., & Mercurio, F. (2007). *Interest rate models—Theory and practice: With smile, inflation and credit*. Springer Science & Business Media.
- Caldeira, J., & Torrent, H. (2017). Forecasting the us term structure of interest rates using nonparametric functional data analysis. *Journal of Forecasting*, 36(1), 56–73.
- Campbell, J. Y., & Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies*, 58(3), 495–514.
- Carter, C. K., & Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81(3), 541–553.
- Chan, K. S., & Tong, H. (1986). On estimating thresholds in autoregressive models. *Journal of Time Series Analysis*, 7(3), 179–190.
- Chib, S. (2001). Markov chain monte carlo methods: Computation and inference. In *Handbook of econometrics* (Vol. 5, pp. 3569–3649). Elsevier.
- Cieslak, A., & Povala, P. (2015). Expected returns in treasury bonds. *The Review of Financial Studies*, 28(10), 2859–2901.
- Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association*, 74(366a), 427–431.
- Diebold, F. X., & Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2), 337–364.
- Duffee, G. R. (2011). *Forecasting with the term structure: The role of no-arbitrage restrictions* (Technical report, Working Paper). Johns Hopkins University.
- European Union. (2017). Commission delegated regulation (eu) 2017/653 of 8 march 2017 supplementing regulation (eu) no 1286/2014 of the european parliament and of the council on key information documents for packaged retail and insurance-based investment products (priips) by laying down regulatory technical standards with regard to the presentation, content, review and revision of key information documents and the conditions for fulfilling the requirement to provide such documents. *Official Journal of the European Union*, 60.
- Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis*, 15(2), 183–202.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and Hall/CRC.

- Gil-Alana, L. A. (2004). Long memory in the us interest rate. *International Review of Financial Analysis*, 13(3), 265–276.
- Goldfeld, S., & Quandt, R. (1973). The estimation of structural shifts by switching regressions. In *Annals of economic and social measurement* (Vol. 2, No. 4, pp. 475–485). NBER.
- Granger, C. W. J., & Andersen, A. P. (1978). *An introduction to bilinear time series models*. Vandenhoeck and Ruprecht: Göttingen.
- Haggan, V., & Ozaki, T. (1981). Modelling nonlinear random vibrations using an amplitude-dependent autoregressive time series model. *Biometrika*, 68(1), 189–196.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society*, 57, 357–384.
- Jansen, E. S., & Teräsvirta, T. (1996). Testing parameter constancy and super exogeneity in econometric equations. *Oxford Bulletin of Economics and Statistics*, 58(4), 735–763.
- Korn, R., & Wagner, A. (2019). *Praxishandbuch lebensversicherungsmathematik: Simulation und klassifikation von produktent*. VVW GmbH.
- Kozicki, S., & Tinsley, P. A. (2001). Shifting endpoints in the term structure of interest rates. *Journal of Monetary Economics*, 47(3), 613–652.
- Lanne, M., & Saikkonen, P. (2002). Threshold autoregressions for strongly autocorrelated time series. *Journal of Business & Economic Statistics*, 20(2), 282–289.
- Lim, K. S., & Tong, H. (1980). Threshold autoregressions, limit cycles, and data. *Journal of the Royal Statistical Society, B*, 42, 245–292.
- Litterman, R., & Scheinkman, J. (1991). Common factors affecting bond returns. *Journal of Fixed Income*, 1(1), 54–61.
- Millar, R. B., & Meyer, R. (2000). Non-linear state space modelling of fisheries biomass dynamics by using metropolis-hastings within-gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 49(3), 327–342.
- Morley, J. C. (2014). Nonlinear time series in macroeconomics. In R. Meyers (Ed.), *Encyclopedia of Complexity and System Science* (pp. 1–30). Berlin, Heidelberg: Springer.
- Quandt, R. E. (1958). The estimation of the parameters of a linear regression system obeying two separate regimes. *Journal of the American Statistical Association*, 53(284), 873–880.
- Rao, T. S., & Gabr, M. M. (2012). *An introduction to bispectral analysis and bilinear time series models* (Vol. 24). Springer Science & Business Media.
- Rose, A. K. (1988). Is the real interest rate stable? *The Journal of Finance*, 43(5), 1095–1112.
- Rudebusch, G. D., & Wu, T. (2008). A macro-finance model of the term structure, monetary policy and the economy. *The Economic Journal*, 118(530), 906–926.
- Schich, S. (1997). Schätzung der deutschen zinsstrukturkurve (Bundesbank Series 1 Discussion Paper). Volkswirtschaftliche Forschungsgruppe der Deutschen Bundesbank.
- Siklos, P. L., & Wohar, M. E. (1997). Convergence in interest rates and inflation rates across countries and over time. *Review of International Economics*, 5(1), 129–141.
- Stock, J. H., & Watson, M. W. (1988). Testing for common trends. *Journal of the American Statistical Association*, 83(404), 1097–1107.
- Teräsvirta, T. (1994). Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association*, 89(425), 208–218.
- Teräsvirta, T., Tjøstheim, D., Granger, C. W. J., et al. (2010). *Modelling nonlinear economic time series*. Oxford University Press Oxford.
- van den End, J. W., & van den End J. W. (2011). Statistical Evidence on the Mean Reversion of Interest Rates. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.1950596>
- Van Dijk, D., Koopman, S. J., Van der Wel, M., & Wright, J. H. (2014). Forecasting interest rates with shifting endpoints. *Journal of Applied Econometrics*, 29(5), 693–712.

AUTHOR BIOGRAPHIES

Christoph Berninger is research associate and PhD candidate at the Department of Statistics, LMU Munich, Germany. He received his master's degree in Financial and Insurance Mathematics from the LMU Munich in 2016. His research interests mainly include time series analysis, financial econometrics, interest rate modelling, and machine learning.

Almond Stöcker started his doctorate at LMU Munich and currently continues his work at HU Berlin, School of Business and Economics. With a broad interest in statistical methodology and application, he currently focusses on extending conventional regression to modelling non-standard response properties. This includes distributional regression, mixed models and especially models for functional, shape, and compositional data.

David Rügamer is currently interim professor for Data Science at the Department of Statistics, LMU Munich, Germany. He also leads two research focus groups at the chair of Statistical Learning and Data Science on Boosting and Probabilistic Machine Learning. His research addresses the combination of statistics and machine and deep learning and is motivated by applications from various fields such as biomechanics, economics, medicine, neuroscience, or psychology.

How to cite this article: Berninger, C., Stöcker, A., & Rügamer, D. (2022). A Bayesian time-varying autoregressive model for improved short-term and long-term prediction. *Journal of Forecasting*, 41(1), 181–200. <https://doi.org/10.1002/for.2802>

APPENDIX A: FULL CONDITIONAL DISTRIBUTIONS

A.1 | The full conditional distribution of $\tilde{\alpha}$

The prior distribution of $\tilde{\alpha}$ is a centered Gaussian process with a specific covariance structure $\tilde{\Sigma}$, i.e.,

$$\tilde{\alpha} = (\alpha_1, \dots, \alpha_T, \dots, \alpha_{T+h}) \sim \mathcal{N}_T(\mathbf{0}, \tilde{\Sigma})$$

The following derivations will be independent of the specific choice of $\tilde{\Sigma}$. By defining

$$\Delta_j = x_{j+1} - \beta x_j$$

as well as $\mathbf{\Delta} = (\Delta_0, \dots, \Delta_{T-1})^\top$ and the fact that

$$\phi(x_t | \alpha_t + \beta x_{t-1}, \sigma^2) = \phi(\alpha_t | \Delta_{t-1}, \sigma^2)$$

allows a straightforward derivation of the full conditional of $\tilde{\alpha}$:

$$\begin{aligned} p(\tilde{\alpha} | \beta, \sigma^2, \mathbf{x}) &\propto p(\mathbf{x} | \tilde{\alpha}, \beta, \sigma^2) p(\tilde{\alpha} | \beta, \sigma^2) \\ &\propto p(\mathbf{x} | \alpha, \beta, \sigma^2) p(\tilde{\alpha}) \\ &\propto \exp\left(-\frac{1}{2\sigma^2} (\alpha - \mathbf{\Delta})^\top (\alpha - \mathbf{\Delta})\right) \\ &\quad \exp\left(-\frac{1}{2} \tilde{\alpha}^\top \tilde{\Sigma}^{-1} \tilde{\alpha}\right) \\ &\propto \exp\left(-\frac{1}{2} (\tilde{\alpha}^\top \tilde{\Sigma}_{post}^{-1} \tilde{\alpha} - 2\tilde{\alpha}^\top \tilde{\Sigma}_{post}^{-1} \underbrace{\tilde{\Sigma}_{post} \tilde{\Delta}_0 \frac{1}{\sigma^2}}_{=\tilde{\mu}_{post}})\right) \end{aligned}$$

with $\tilde{\Sigma}_{post}^{-1} = \tilde{\Sigma}^{-1} + \frac{1}{\sigma^2} \begin{pmatrix} \mathbf{I}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}$ and $\tilde{\Delta}_0 = (\Delta^\top, \mathbf{0})^\top$.

This is the kernel of a multivariate Gaussian distribution with covariance $\tilde{\Sigma}_{post}$ and mean vector $\tilde{\mu}_{post}$, i.e

$$\tilde{\alpha} | \beta, \sigma^2, \mathbf{x} \sim \mathcal{N}(\tilde{\mu}_{post}, \tilde{\Sigma}_{post}).$$

A.2 | The full conditional distribution of ρ

If an AR-covariance structure is assumed for $\tilde{\alpha}$ the latent α -process can be written in the following form

$$\alpha_t = \rho \alpha_{t-1} + \eta_t,$$

where ρ determines the correlation between two successive time steps and η_t is a Gaussian white noise process, i.e., $\eta_t \sim \mathcal{N}(0, \tau^2)$.

The full conditional distribution of ρ can be therefore derived as follows:

$$p(\rho | \tau^2, \alpha) \propto \mathcal{L}(\rho, \tau^2) \cdot p(\rho) = \prod_{j=0}^{T-1} \phi(\alpha_{T-j} | \rho \alpha_{T-j-1}, \tau^2) \cdot p(\rho). \quad (\text{A.1})$$

The likelihood $\mathcal{L}(\cdot)$ in the above equation can be reformulated as

$$\mathcal{L} \propto \exp\left(-\frac{1}{2\tau^2} \left\{ -2\rho \left[\sum_{j=0}^{T-1} \alpha_{T-j} \alpha_{T-j-1} \right] + \rho^2 \left[\sum_{j=0}^{T-1} \alpha_{T-j-1}^2 \right] \right\}\right).$$

The calculation is similar to the one in Appendix 2.5. Defining the two terms in square brackets as η and χ , respectively, we get

$$\mathcal{L} \propto \exp\left(-\frac{1}{2\tau^2} \{-2\rho\eta + \rho^2\chi\}\right).$$

Plugging this into (15) and using a normal prior with parameters μ_ρ, σ_ρ^2 for ρ , we have

$$\begin{aligned} p(\rho | \tau^2, \alpha) &\propto \exp\left(-\frac{1}{2} \left\{ \frac{\rho^2 \chi}{\tau^2} - 2\frac{\rho \eta}{\tau^2} \right\}\right) \\ &\quad \exp\left(-\frac{1}{2} \left\{ \frac{\rho^2}{\sigma_\rho^2} - 2\frac{\rho \mu_\rho}{\sigma_\rho} \right\}\right) \\ &\propto \exp\left(-\frac{1}{2} \left\{ \rho^2 \cdot \left(\frac{\chi}{\tau^2} + \sigma_\rho^{-2} \right) - 2\rho \left(\frac{\eta}{\tau^2} + \frac{\mu_\rho}{\sigma_\rho} \right) \right\}\right) \end{aligned}$$

and thus $\rho | \tau^2, \alpha \sim \mathcal{N}(\mu_{\rho, post}, \sigma_{\rho, post}^2)$ with

$$\sigma_{\rho, post}^2 = \left(\frac{\chi}{\tau^2} + \sigma_\rho^{-2} \right)^{-1}$$

and

$$\mu_{\rho, post} = \left(\frac{\eta}{\tau^2} + \frac{\mu_\rho}{\sigma_\rho} \right) \sigma_{\rho, post}^2.$$

If a truncated normal prior is used, the truncation is transferred to the full conditional distribution.

A.3 | The full conditional distribution of β

Analogously to (10) and (11), we have

$$\begin{aligned} p(\beta|\alpha, \sigma^2, \mathbf{x}) &\propto \mathcal{L}(\beta, \alpha, \sigma^2) \cdot p(\beta) \\ &= \prod_{j=0}^{T-1} \phi(x_{T-j}|\alpha_{T-j} + \beta x_{T-j-1}, \sigma^2) \cdot p(\beta). \end{aligned} \quad (\text{A.2})$$

By defining $\check{d}_{t-j} := x_{t-j} - \alpha_{t-j}$ and as

$$\phi(x_{t-j}|\alpha_{t-j} + \beta x_{t-j-1}, \sigma^2) = \phi(\beta x_{t-j-1}|\check{d}_{t-j}, \sigma^2)$$

the likelihood $\mathcal{L}(\cdot)$ in the above equation can be reformulated as

$$\mathcal{L} \propto \exp\left(-\frac{1}{2\sigma^2} \left\{ -2\beta \left[\sum_{j=0}^{T-1} \check{d}_{T-j} x_{T-j-1} \right] + \beta^2 \left[\sum_{j=0}^{T-1} x_{T-j-1}^2 \right] \right\}\right).$$

You can find a more detailed calculation in Appendix 2.5. Defining the two terms in square brackets as η and χ , respectively, we get

$$\mathcal{L} \propto \exp\left(-\frac{1}{2\sigma^2} \{-2\beta\eta + \beta^2\chi\}\right).$$

Plugging this into (16) and using a normal prior with parameters $\mu_\beta, \sigma_\beta^2$ for β , we have

$$\begin{aligned} p(\beta|\alpha, \sigma^2, \mathbf{x}) &\propto \exp\left(-\frac{1}{2} \left\{ \frac{\beta^2\chi}{\sigma^2} - 2\frac{\beta\eta}{\sigma^2} \right\}\right) \\ &\exp\left(-\frac{1}{2} \left\{ \frac{\beta^2}{\sigma_\beta^2} - 2\frac{\beta\mu_\beta}{\sigma_\beta} \right\}\right) \\ &\propto \exp\left(-\frac{1}{2} \left\{ \beta^2 \cdot \left(\frac{\chi}{\sigma^2} + \sigma_\beta^{-2} \right) - 2\beta \left(\frac{\eta}{\sigma^2} + \frac{\mu_\beta}{\sigma_\beta} \right) \right\}\right) \end{aligned}$$

and thus $\beta|\mathbf{x}, \alpha, \sigma^2 \sim \mathcal{N}(\mu_{\beta, \text{post}}, \sigma_{\beta, \text{post}}^2)$ with

$$\sigma_{\beta, \text{post}}^2 = \left(\frac{\chi}{\sigma^2} + \sigma_\beta^{-2} \right)^{-1}$$

and

$$\mu_{\beta, \text{post}} = \left(\frac{\eta}{\sigma^2} + \frac{\mu_\beta}{\sigma_\beta} \right) \sigma_{\beta, \text{post}}^2.$$

If a truncated normal prior is used, the truncation is transferred to the full conditional distribution.

A.4 | The full conditional distribution of σ^2

In this section, we derive the full conditional distribution of σ^2 . As before,

$$\begin{aligned} p(\sigma^2|\alpha, \beta, \mathbf{x}) &\propto \prod_{j=0}^{T-1} \phi(x_{T-j}|\alpha_{T-j} + \beta x_{T-j-1}, \sigma^2) \\ &\cdot p(\sigma^2) \cdot p(\beta|\sigma^2), \end{aligned}$$

which is equal to

$$\begin{aligned} &(\sigma^2)^{-\frac{t}{2}} \exp\left(-\frac{1}{2\sigma^2} \sum_{j=0}^{T-1} \epsilon_{T-j}^2\right) \cdot p(\sigma^2) \cdot p(\beta|\sigma^2) \\ &=: (\sigma^2)^{-\frac{t}{2}} \exp\left(-\frac{1}{2\sigma^2} \kappa\right) \cdot p(\sigma^2) \cdot p(\beta|\sigma^2). \end{aligned}$$

By using an inverse gamma distribution with shape and scale parameters a, b , or short $\mathcal{IG}(a, b)$, for the prior of σ^2 we get

$$\begin{aligned} &(\sigma^2)^{-(t/2)} \exp\left(-\frac{1}{2\sigma^2} \kappa\right) \cdot (\sigma^2)^{-(a+1)} \exp(-b/\sigma^2) \\ &\cdot (\sigma^2)^{-\frac{1}{2}} \exp\left(-\frac{1}{2\sigma^2 \sigma_\beta^2} (\beta - \mu_\beta)^2\right) \\ &= \frac{1}{\Phi\left(\frac{1-\mu_\beta}{\sigma\sigma_\beta}\right) - \Phi\left(\frac{-1-\mu_\beta}{\sigma\sigma_\beta}\right)} \\ &(\sigma^2)^{-(t/2)} \exp\left(-\frac{1}{2\sigma^2} \kappa\right) \cdot (\sigma^2)^{-(a+\frac{3}{2})} \\ &\exp\left(-\frac{b + \frac{(\beta - \mu_\beta)^2}{2\sigma_\beta^2}}{\sigma^2}\right) \frac{1}{\Phi\left(\frac{1-\mu_\beta}{\sigma\sigma_\beta}\right) - \Phi\left(\frac{-1-\mu_\beta}{\sigma\sigma_\beta}\right)} \approx \\ &(\sigma^2)^{-(\frac{t+1}{2}+a+1)} \exp\left(-\frac{\frac{\kappa}{2} + b + \frac{(\beta - \mu_\beta)^2}{2\sigma_\beta^2}}{\sigma^2}\right). \end{aligned}$$

In the last step, we omitted the last term, which results from the truncation, as in our application the truncation is not very restrictive such that this term is

close to 1. Thus the full conditional distribution is approximately also an inverse gamma distribution with parameters $\tilde{a} = \frac{t+1}{2} + a$ and $\tilde{b} = \frac{\kappa}{2} + b + \frac{(\beta - \mu_\beta)^2}{2} \sigma_\beta^2$, i.e.,

$$\sigma^2 | \alpha, \beta, \mathbf{x} \sim \mathcal{IG}(\tilde{a}, \tilde{b}).$$

A.5 | Rewriting the likelihood of the parameters

By defining $\check{d}_{t-j} := x_{t-j} - \alpha_{t-j}$, the likelihood of the parameters can be reformulated as follows:

$$\begin{aligned} \mathcal{L}(\beta, \alpha, \sigma^2) &= \prod_{j=0}^{T-1} \phi(x_{T-j} | \alpha_{T-j} + \beta x_{T-j-1}, \sigma^2) \\ &= \prod_{j=0}^{T-1} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x_{T-j} - \alpha_{T-j} - \beta x_{T-j-1})^2}{2\sigma^2}\right) \\ &= \prod_{j=0}^{T-1} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\check{d}_{T-j} - \beta x_{T-j-1})^2}{2\sigma^2}\right) \\ &= \prod_{j=0}^{T-1} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\check{d}_{T-j}^2 - 2\beta x_{T-j-1} \check{d}_{T-j} + \beta^2 x_{T-j-1}^2)}{2\sigma^2}\right) \\ &\propto \prod_{j=0}^{T-1} \exp\left(-\frac{1}{2\sigma^2} \left\{ -2\beta x_{T-j-1} \check{d}_{T-j} + \beta^2 x_{T-j-1}^2 \right\}\right) \\ &= \exp\left(-\frac{1}{2\sigma^2} \left\{ -2\beta \sum_{j=0}^{T-1} \check{d}_{T-j} x_{T-j-1} + \beta^2 \sum_{j=0}^{T-1} x_{T-j-1}^2 \right\}\right). \end{aligned}$$

APPENDIX B: METROPOLIS-HASTINGS WITHIN GIBBS SAMPLER ROUTINE

Starting with an initial sample $(\alpha^{(0)}, \beta^{(0)}, (\sigma^2)^{(0)}, \rho^{(0)}, (\tau^2)^{(0)})$, where

$$(\tau^2)^{(0)} = f(\beta^{(0)}, (\sigma^2)^{(0)}, \rho^{(0)}, \text{Var}(x_t))$$

as specified in (8), we first draw a sample of $\tilde{\alpha}$ values from its full conditional distribution. We proceed with the Metropolis-Hastings algorithm step by drawing from the conditional distributions of ρ , σ^2 and β as derived in Section 2.3. Furthermore, τ^2 is set according to (6) such that a prior specified long run variance is met. We calculate the density value of the proposal distribution q specified in (12), i.e.,

$$\begin{aligned} q(\rho^{(n+1)}, \beta^{(n+1)}, (\sigma^2)^{(n+1)} | (\tau^2)^{(n)}, \beta^{(n)}, (\sigma^2)^{(n)}, \alpha, \mathbf{x}) &= \\ q(\rho^{(n+1)} | (\tau^2)^{(n)}, \alpha, \mathbf{x}) q((\sigma^2)^{(n+1)} | & \\ \beta^{(n)}, \alpha, \mathbf{x}) q(\beta^{(n+1)} | (\sigma^2)^{(n)}, \alpha, \mathbf{x}). & \end{aligned}$$

We further calculate the density value of the proposal distribution for the parameters of the previous step conditional on the new drawn parameter, i.e.,

$$\begin{aligned} q(\rho^{(n)}, \beta^{(n)}, (\sigma^2)^{(n)} | (\tau^2)^{(n+1)}, \beta^{(n+1)}, (\sigma^2)^{(n+1)}, \alpha, \mathbf{x}) &= \\ q(\rho^{(n)} | (\tau^2)^{(n+1)}, \alpha, \mathbf{x}) \cdot q((\sigma^2)^{(n)} | \beta^{(n+1)}, \alpha, \mathbf{x}) \cdot & \\ q(\beta^{(n)} | (\sigma^2)^{(n+1)}, \alpha, \mathbf{x}). & \end{aligned}$$

The true conditional posterior density is given by

$$\begin{aligned} p(\rho, \beta, \sigma^2 | \alpha, \mathbf{x}) &\propto p(\mathbf{x} | \beta, \sigma^2, \alpha) \\ p(\alpha | \rho, f(\rho, \beta, \sigma^2)) p(\rho) p(\beta | \sigma^2) p(\sigma^2). & \end{aligned}$$

The acceptance probability is calculated by

$$P_{\text{accept.}} = \min\left(1, \frac{p(\rho^{(n+1)}, \beta^{(n+1)}, (\sigma^2)^{(n+1)} | \alpha^{(n+1)}, \mathbf{x})}{p(\rho^{(n)}, \beta^{(n)}, (\sigma^2)^{(n)} | \alpha^{(n+1)}, \mathbf{x})} \cdot \frac{q(\rho^{(n)}, \beta^{(n)}, (\sigma^2)^{(n)} | \rho^{(n+1)}, \beta^{(n+1)}, (\sigma^2)^{(n+1)}, \alpha^{(n+1)}, \mathbf{x})}{q(\rho^{(n+1)}, \beta^{(n+1)}, (\sigma^2)^{(n+1)} | \rho^{(n)}, \beta^{(n)}, (\sigma^2)^{(n)}, \alpha^{(n+1)}, \mathbf{x})}\right).$$

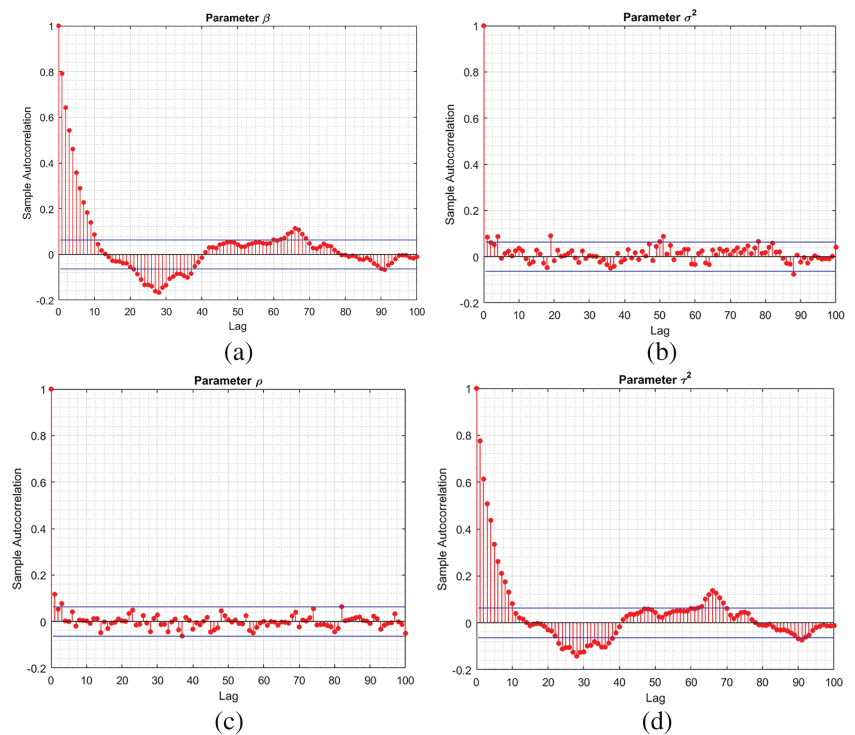
A new drawn sample is accepted if a uniform distributed random variable is smaller than the acceptance probability. Otherwise the sample from the previous step is taken. After a burn-in period the parameter set $(\tilde{\alpha}^{(m)}, \beta^{(m)}, (\sigma^2)^{(m)})$ is approximately distributed according to the joint posterior distribution $p(\tilde{\alpha}, \beta, \sigma^2 | \mathbf{x})$.

APPENDIX C: DIAGNOSTICS OF THE MCMC ALGORITHM

In this section, we present diagnostics of the Metropolis-Hastings within Gibbs sampler (MHwGS) routine used for the dataset in our application.

To investigate the distribution of the 10-year interest rate at a forecasting horizon of 40 years, we use the prior distributions as specified in Section 3.3.1 and draw 10.000 samples by applying the MHwGS routine as described in Appendix A1. The chosen proposal distribution results in an acceptance rate of 30.92%. After a burn-in phase of 100 samples, we apply thinning to reduce autocorrelation and take every 10th parameter set to simulated paths of the first principal component. The autocorrelation functions for selected parameters are visualized in Figure C1.

FIGURE C1 Autocorrelation function of the parameters, which stay constant over time



APPENDIX D: BACKTEST RESULTS

TABLE D1 Results of the out-of-sample 1-month ahead forecasting

Maturity	Mean	Std. Dev.	RMSE
<i>The BTVC-AR(1)-Factor model</i>			
1 year	-0.0268	0.2566	0.0659
3 years	-0.0469	0.2289	0.0541
5 years	-0.0681	0.2402	0.0617
10 years	-0.0640	0.2346	0.0586
<i>The Gauss2++ model</i>			
1 year	-0.0808	0.2361	0.0618
3 years	-0.1037	0.2252	0.0610
5 years	-0.1203	0.2139	0.0598
10 years	-0.1429	0.2130	0.0654
<i>The dynamic Nelson-Siegel model</i>			
1 year	-0.0290	0.2615	0.0685
3 years	-0.0462	0.2311	0.0550
5 years	-0.0653	0.2410	0.0617
10 years	-0.0589	0.2340	0.0577

TABLE D2 Results of the out-of-sample 3-month ahead forecasting

Maturity	Mean	Std. Dev.	RMSE
<i>The BTVC-AR(1)-Factor model</i>			
1 year	-0.1264	0.5064	0.2697
3 years	-0.1505	0.4640	0.2358
5 years	-0.1725	0.4354	0.2174
10 years	-0.1625	0.3875	0.1751
<i>The Gauss2++ model</i>			
1 year	-0.2057	0.5329	0.3236
3 years	-0.2707	0.4702	0.2923
5 years	-0.3098	0.4208	0.2714
10 years	-0.3435	0.3875	0.2667
<i>The dynamic Nelson-Siegel model</i>			
1 year	-0.1327	0.5152	0.2803
3 years	-0.1482	0.4665	0.2374
5 years	-0.1643	0.4343	0.2137
10 years	-0.1478	0.3827	0.1668

TABLE D3 Results of the out-of-sample 6-month ahead forecasting

Maturity	Mean	Std. Dev.	RMSE
<i>The BTVC-AR(1)-Factor model</i>			
1 year	-0.2809	0.7683	0.6631
3 years	-0.3093	0.6941	0.5725
5 years	-0.3311	0.6330	0.5062
10 years	-0.3110	0.5462	0.3920
<i>The Gauss2++ model</i>			
1 year	-0.4094	0.8105	0.8184
3 years	-0.5402	0.6768	0.7457
5 years	-0.6098	0.6090	0.7393
10 years	-0.6545	0.5824	0.7693
<i>The dynamic Nelson-Siegel model</i>			
1 year	-0.2900	0.7857	0.6951
3 years	-0.3022	0.7045	0.5825
5 years	-0.3130	0.6380	0.5008
10 years	-0.2812	0.5446	0.3727

TABLE D4 Results of the out-of-sample 12-month ahead forecasting

Maturity	Mean	Std. Dev.	RMSE
<i>The BTVC-AR(1)-Factor model</i>			
1 year	-0.5956	0.9591	1.2652
3 years	-0.6264	0.7861	1.0041
5 years	-0.6526	0.6834	0.8881
10 years	-0.6275	0.5986	0.7484
<i>The Gauss2++ model</i>			
1 year	-0.9047	1.0709	1.9546
3 years	-1.1531	0.7939	1.9541
5 years	-1.2745	0.7255	2.1458
10 years	-1.3345	0.8060	2.4246
<i>The dynamic Nelson-Siegel model</i>			
1 year	-0.6004	0.9961	1.3424
3 years	-0.6024	0.8218	1.0316
5 years	-0.6098	0.7096	0.8702
10 years	-0.5657	0.6024	0.6793