

Changing dynamics: Time-varying autoregressive models using generalized additive
modeling

Laura F. Bringmann

University of Leuven, Leuven, Belgium

Ellen L. Hamaker

University of Utrecht, Utrecht, the Netherlands

Daniel E. Vigo

Pontifical Catholic University of Argentina and National Research Council (CONICET),

Buenos Aires, Argentina

University of Leuven, Leuven, Belgium

André Aubert

University of Leuven, Leuven, Belgium

Denny Borsboom

University of Amsterdam, Amsterdam, the Netherlands

Francis Tuerlinckx

University of Leuven, Leuven, Belgium

Author Note

Address for correspondence: Department Quantitative Psychology and Individual Differences; University of Leuven; Tiensestraat 102 - box 3713; 3000 Leuven; Belgium; tel.+32163 26052; email: laura.bringmann@ppw.kuleuven.be. The research leading to the results reported in this paper was sponsored in part by Belgian Federal Science Policy within the framework of the Interuniversity Attraction Poles program (IAP/P7/06), as well as by grant GOA/15/003 from the KU Leuven, and grant G.0806.13 from the Fund of Scientific Research Flanders. We are thankful to Noémi Schuurman and Markus Eronen for their helpful comments on several versions of the manuscript.

Abstract

In psychology, the use of intensive longitudinal data has steeply increased during the past decade. As a result, studying temporal dependencies in such data with autoregressive modeling is becoming common practice. However, standard autoregressive models are often suboptimal as they assume that parameters are time-invariant. This is problematic if changing dynamics (e.g., changes in the temporal dependency of a process) govern the time series. Often a change in the process, such as emotional well-being during therapy, is the very reason why it is interesting and important to study psychological dynamics. As a result, there is a need for an easily applicable method for studying such non-stationary processes that result from changing dynamics. In this article we present such a tool: the semi-parametric TV-AR model. We show with a simulation study and an empirical application that the TV-AR model can approximate nonstationary processes well if there are at least 100 time points available and no unknown abrupt changes in the data. Notably, no prior knowledge of the processes that drive change in the dynamic structure is necessary. We conclude that the TV-AR model has significant potential for studying changing dynamics in psychology.

Keywords: Time series; Non-stationarity; Autoregressive models; Generalized Additive Models; Splines

Changing dynamics: Time-varying autoregressive models using generalized additive modeling

Humans are complex dynamic systems, whose emotions, cognitions, and behaviors fluctuate constantly over time (Nesselroade & Ram, 2004; L. P. Wang, Hamaker, & Bergeman, 2012). In order to study these within-person processes, and to determine how, why, and when individuals change over time, individuals need to be measured on a relatively large number of occasions (Bolger & Laurenceau, 2013; Ferrer & Nesselroade, 2003; Molenaar & Campbell, 2009; Nesselroade & Ram, 2004; Nesselroade & Molenaar, 2010), resulting in intensive longitudinal data that, if $N = 1$, are typically designated as time series (Walls & Schafer, 2006). Currently, a spectacular growth of studies gathering intensive longitudinal data is taking place (aan het Rot, Hogenelst, & Schoevers, 2012; Bolger, Davis, & Rafaeli, 2003; Mehl & Conner, 2012; Scollon, Prieto, & Diener, 2003). With this development, it has become possible to study dynamical processes of psychological phenomena in much greater detail than has hitherto been possible (Trull & Ebner-Priemer, 2013).

There are various aspects of within-person processes that one can choose to study in order to gather insights into psychological dynamics, of which *temporal dependence* is one particularly informative aspect (Boker, Molenaar, & Nesselroade, 2009; Hamaker, Ceulemans, Grasman, & Tuerlinckx, in press; McArdle, 2009). Temporal dependence concerns the degree to which current observations can be predicted by previous observations, for example, the degree to which an individual's emotional state at a given time point is predictive of her emotional state at subsequent time points (Jahng, Wood, & Trull, 2008; Kuppens, Allen, & Sheeber, 2010).

A popular approach to handling such temporal dependency is autoregressive (AR) modeling, a family of statistical models in which the structure of the time-dependency in the data is explicitly modeled through regression equations. Some autoregressive models are suited to study time dependence within a single individual (e.g., Hertzog &

Nesselroade, 2003; Molenaar, 1985; Rosmalen, Wenting, Roest, de Jonge, & Bos, 2012; Stroe-Kunold et al., 2012), whereas multilevel techniques can model time dependence within multiple individuals simultaneously (e.g., Bringmann et al., 2013; de Haan-Rietdijk, Gottman, Bergeman, & Hamaker, 2014; Song & Ferrer, 2012; Oravecz, Tuerlinckx, & Vandekerckhove, 2011). In addition, AR techniques can be applied in various frameworks, such as the Bayesian (e.g., Pole, West, & Harrison, 1994) and the structural equation modeling framework (SEM; e.g., Hamaker, Dolan, & Molenaar, 2003; McArdle, 2009; Voelkle, Oud, Davidov, & Schmidt, 2012).

A drawback of most AR models is that they are based on the assumption that the average value around which the process is fluctuating as well as the variance and the temporal dependency of the process are time-invariant. This is also known as the *stationarity assumption* (Chatfield, 2003). However, in the context of psychology this may not always be a realistic assumption. In fact, it could be argued that in many psychological time series studies a form of non-stationarity can be expected to be present (e.g., Bringmann, Lemmens, Huibers, Borsboom, & Tuerlinckx, 2014; Molenaar, De Gooijer, & Schmitz, 1992; Rosmalen et al., 2012; Tschacher & Ramseyer, 2009). Even more so, often the very reason why it is interesting and important to study dynamics of psychological processes lies in their non-stationary nature (Boker, Rotondo, Xu, & King, 2002; van de Leemput et al., 2014). For example, when an individual receives therapy, the aim is to accomplish change, such as a decrease in symptoms. Thus, instead of considering dynamics, such as temporal dependency, as static characteristics of an individual, it is more realistic to consider them as time-varying, which implies that standard AR models are unsuitable (Molenaar et al., 1992; Boker et al., 2002).

To overcome this limitation, time-varying AR (TV-AR) models have been developed (Dahlhaus, 1997). In these models, the parameters (the intercept and autoregressive parameter) of the AR model (most commonly an AR(1) model) are now allowed to vary over time, so the models can be applied to both stationary and non-stationary processes

(Chow, Zu, Shifren, & Zhang, 2011). Most time-varying AR models used in psychology and econometrics are based on the state-space modeling framework (Chow et al., 2011; Koop, 2012; Molenaar, 1987; Molenaar & Newell, 2003; Molenaar, Sinclair, Rovine, Ram, & Corneal, 2009; Mumtaz & Surico, 2009; Prado, 2010; Tarvainen, Hiltunen, Ranta-aho, & Karjalainen, 2004; Tarvainen, Georgiadis, Ranta-aho, & Karjalainen, 2006; West, Prado, & Krystal, 1999). The state-space framework is very general and encompasses a wide variety of models, such as dynamic linear models. Hence, the framework is very powerful due to its generality, but the downside is that it requires learning (state-space) notation with which most psychologists are unfamiliar. In addition, state-space models require the user to specify the way parameters of the time-varying model vary over time (Belsley & Kuh, 1973; Tarvainen et al., 2004; for a notable exception see Molenaar et al., 2009), but in practice the required theories about the nature of the change are often lacking (Tan, Shiyko, Li, Li, & Dierker, 2012), or must be handled via explicit incorporation of spline-based or other nonparametric functions into a (confirmatory) state-space framework (Tarvainen et al., 2006). Doing so may entail high computational demands when the dimension of the unknown change forms to be explored is high. Thus, there is a clear need for a time-varying AR method that functions without pre-specification and moreover is easy to apply for researchers in psychology.

As we will show in this paper, one solution is to implement TV-AR models based on *semi-parametric* statistical modeling using a well-studied elegant and easily applicable generalized additive modeling (GAM) framework (Hastie & Tibshirani, 1990; McKeown & Sneddon, 2014; Sullivan, Shadish, & Steiner, 2015; Wood, 2006). The crucial advantage of semi-parametric TV-AR models in general is that they are data-driven, and thus the shape of change need not be specified beforehand (Dahlhaus, 1997; Fan & Yao, 2003; Giraitis, Kapetanios, & Yates, 2014; Härdle, Lütkepohl, & Chen, 1997; Kitagawa & Gersch, 1985). Furthermore, no state-space notation is needed, since the TV-AR model is closely related to and can be specified and estimated within the familiar regression framework. Software

for applying the GAM framework is freely available in the *mgcv* package for the statistical software *R* (Wood, 2006). The package has well-functioning default settings, making it very user friendly.¹ By showing how the TV-AR model can be applied with existing and easy to use software, we hope to make the TV-AR method accessible for a broad audience of psychological researchers.

The structure of the paper is as follows. In the first section, a detailed explanation of the standard time-invariant AR is given. In the second section, we describe the general structure of the TV-AR model, and in the third section we explain in detail how the time-varying parameters are estimated, and also introduce the *mgcv* package in *R*, with which the TV-AR is estimated (McKeown & Sneddon, 2014; Wood, 2006). In the fourth section, we provide a simulation study and give guidelines on how to use the TV-AR model with the *mgcv* package. In the fifth section, we give an example from emotion dynamics research to illustrate the TV-AR method by applying it to two different subjects whose affect was measured over circa 500 days in the context of an isolation study, the MARS500 project (Basner et al., 2013; Tafforin, 2013; Vigo et al., 2013; Y. Wang et al., 2014). This section is followed by concluding remarks and the Appendix with a description of the *R*-code used throughout the article. Additional details of the simulation study can be found in the online supplemental material.

Standard time-invariant AR

In this section, the standard time-invariant autoregressive (AR) model is explained in more detail. Code for the equations and figures in this section can be found in the *R-code* in the Appendix under the heading *II. Standard time-invariant AR*.

Time series data consist of repeated measurements on one or more variable(s) taken from the same system (e.g., an individual, dyad, family, or organization). Typically, such data are statistically dependent, since all measures are taken from the same participant (e.g., answers on a questionnaire are likely to be related over time, Brandt & Williams,

2007; Velicer & Fava, 2003). This statistical dependence or autocorrelation that occurs in repeated measurement data is a central aspect that has to be accounted for when studying the underlying process. Furthermore, when this autocorrelation is not taken into account invalid estimates can occur.

In psychology, the standard model used to deal with this statistical dependency is a Gaussian discrete time AR model.² An AR model accounts for the statistical dependency by modeling it explicitly, or in other words, the time series is regressed on itself (Hamaker & Dolan, 2009). The most basic form is an AR model of lag order 1 or AR(1):

$$y_t = \beta_0 + \beta_1 y_{t-1} + \varepsilon_t. \quad (1)$$

This amounts to a linear regression model with an intercept β_0 , and the autoregressive coefficient β_1 , representing the degree and direction of the relation between a measurement at a previous (lagged) time point ($t - 1$) and current time point (t) of a single variable y (Velicer & Fava, 2003) and can be estimated with ordinary least squares (OLS). The part of observation y_t that cannot be explained by the previous observation y_{t-1} is referred to as the innovation ε_t (Chatfield, 2003). Other terms for the innovation are random shock, perturbation, or dynamic error.³ The innovations are assumed to be normally distributed with a mean of zero and variance σ_ε^2 (Hamilton, 1994).

The autoregressive coefficient β_1 can also be interpreted as the extent to which a current observation is predictable by the preceding observation (Hamaker & Dolan, 2009). A positive relationship indicates that high values of a variable (e.g., Positive Affect; PA) at one time point are likely to be followed by high values in the next time period (see left panel of Figure 1). In contrast, a negative relationship would predict the opposite, namely low values of the variable during the next time period (Chatfield, 2003; Velicer & Fava, 2003), which typically results in a jigsaw pattern (see right panel of Figure 1).

An important assumption for an AR(1) model is stationarity. A distinction is made between strictly stationary and covariance-stationary (also known as weakly or

second-order stationary) processes. If a process is strictly stationary, the distribution of y_t and all joint distributions of y random variables are the same at all time points, and are thus time-invariant. Covariance-stationarity is a less strong assumption, as in this case only the first two moments of a distribution, the mean and the variance, and thus the parameters β_0 and β_1 , have to be time-invariant.⁴ Furthermore, stationarity also requires that the autoregressive coefficient must lie between -1 and 1 (boundaries not included). In this case, the mean μ and variance σ^2 of the process in Equation 1 can be expressed as

$$\mu = \frac{\beta_0}{1 - \beta_1} \quad (2)$$

$$\sigma^2 = \frac{\sigma_\varepsilon^2}{1 - \beta_1^2}, \quad (3)$$

showing that both are time-invariant (Chatfield, 2003; Hamilton, 1994).

Figure 1 shows two examples of a stationary process. Although the process fluctuates (changes) in both the left and right panel, the intercept, mean, autocorrelation and variance do not change over time. In an AR model, the intercept term β_0 only has a substantial interpretation if a score of 0 is a possible value in the sample.⁵ Therefore, we prefer to work with the mean μ , which can be interpreted as the value around which the process fluctuates.

Time-varying AR

Psychological data are often non-stationary, rendering a standard AR model inapplicable. In this section, we will therefore describe an alternative model, the TV-AR model, which can model non-stationarity. First, we will discuss non-stationarity, illustrated by two simulated examples with 150 time points (representing here the evolution of valence within an individual). Secondly, we will give a general overview of the TV-AR model. Information on statistical inference for the TV-AR model will be given in the next section. The code to make the figure in this section can be found in the *R-code* in the Appendix under the heading *III. Time-varying AR*.

There are several sources that can give rise to a non-stationary process in which the intercept, mean, autocorrelation and (or) variance change over time. In psychological research, the focus has mainly been on detecting a type of non-stationarity that is due to a (gradual) change in the mean of a process, which is visible as a trend in the data. Consider for example the left panel of Figure 2, in which a simulated process of hypothetical valence scores for an individual is shown. Here the autoregressive parameter does not change over time ($\beta_1 = 0.2$), but the intercept does, as represented by the dashed line, and therefore the mean also changes. Thus, a trend in the data is present.

To deal with a trend, common approaches in psychology have been *detrending* and *modeling the trend*. In the first method, data are made stationary by subtracting the values of a fitted trend from the individual data-points, thus removing the trend from the data (Hamaker & Dolan, 2009). A drawback associated with this way of dealing with non-stationarity is that it may remove important information from the data (Molenaar et al., 1992). In the second approach, stationarity is obtained through modeling the trend with, for example, linear growth curve modeling (Tschacher & Ramseyer, 2009). Both modeling the trend as well as detrending require specifying the functional form of the trend, which can be difficult, especially when convenient parametric forms are not applicable (Adolph, Robinson, Young, & Gill-Alvarez, 2008; Faraway, 2006; Tan et al., 2012). The TV-AR model that we will present has the advantage that it can detect trends in a data-driven way, and thus no pre-specifications are needed to account for a trend in the data.

Detrending or modeling the trend makes the process trend-stationary. However, when detrending, often only the trend due to a changing intercept is removed, and what is overlooked is that non-stationarity and trends can also occur due to changes in the autocorrelation.⁶ For example, Figure 2 (right panel) shows a process that is non-stationary due to a change in the autocorrelation. The autoregressive function changes linearly over time, from a high value ($\beta_1 = 0.65$) to a lower one ($\beta_1 = 0.2$). At first, the

data are characterized by a high autocorrelation, which disappears towards the end of the time series. This is evident in the figure: First there are large oscillations (a signature of a high autocorrelation), which then become smaller towards the end of the time series (indicating low autocorrelation). Removing or modeling a trend as described above will not deal with this source of non-stationarity, leaving the process covertly non-stationary. This is an important reason why TV-AR models, which can detect and model both changes in the intercept and autocorrelation simultaneously, are important.

Another reason why TV-AR models are useful is that they can test for non-stationarity. There are several tests to check for stationarity, such as the Dickey Fuller test (which can be used to test whether a unit root is present in the time series; Dickey & Fuller, 1979), and the KPSS test (which can be used to test whether the mean is stable over time, or whether it follows a linear trend; Kwiatkowski, Phillips, Schmidt, & Shin, 1992). However, there is no specific test that checks for non-stationarity due to changing autoregression or a changing mean that follows a different trajectory than a linear trend. With the TV-AR model, we present a method that can test the time invariance of the autoregressive parameter, and simultaneously check whether a trend is due to a time-varying intercept and/or a time-varying autoregressive parameter (see Figure 2). Moreover, this method allows for instantly modeling such non-stationarity.

The defining feature of a TV-AR model is that the coefficients of the model are allowed to vary over time, following an unspecified function of time (Dahlhaus, 1997; Giraitis et al., 2014). To this end, we specify

$$y_t = \beta_{0,t} + \beta_{1,t}y_{t-1} + \varepsilon_t \quad (4)$$

where the intercept $\beta_{0,t}$ and the autoregressive $\beta_{1,t}$ coefficients are now functions that can change over time.⁷ The innovations still form a white noise process so that the values of ε_t are independently and identically distributed, which implies that their variance is constant over time.

An important assumption of the TV-AR model is that, even though the functional

form of $\beta_{0,t}$ and $\beta_{1,t}$ can be any function, change in the parameter values is restricted to be gradual, that is, there should be no sudden transitions. This assumption implies that the TV-AR model, as defined here, is not appropriate for time series with abrupt changes or sudden jumps. Thus, researchers should decide whether or not continuous change in parameters is plausible on the basis of the substantive knowledge of the problem at hand. If sudden, qualitative transitions are expected (e.g., as would be the case in some areas of cognitive development or in mental disorders with a sudden onset) then the current methodology would not be advisable. However, if the point at which an abrupt change takes place is known, one can model the change with a TV-AR model. One could specify, for example, a TV-AR model before and after an intervention. Additionally, although a TV-AR model is designed for handling non-stationary processes, the process is still required to be *locally stationary*, meaning that $-1 < \beta_{1,t} < 1$, for all t (Dahlhaus, 1997).

Assuming that the change is restricted to be gradual and the process is locally stationary, the model implied mean is (Giraitis et al., 2014):⁸

$$\mu_t \approx \frac{\beta_{0,t}}{1 - \beta_{1,t}}. \quad (5)$$

Similarly, due to the fact that the autoregressive coefficient is allowed to vary over time, the variance of the time series is now also time-varying, that is,

$$\sigma_t^2 \approx \frac{\sigma_\varepsilon^2}{1 - \beta_{1,t}^2}. \quad (6)$$

Note that since μ_t can vary over time, in the literature μ_t is often interpreted as the *attractor* (also known as baseline or equilibrium) rather than the mean of the process (Giraitis et al., 2014; Hamaker, 2012; Oravecz et al., 2011). As is the case in a time-invariant AR model, the intercept and the changing mean (attractor or trend) are distinct features of a process. The intercept typically does not have a direct psychological interpretation, whereas the attractor represents the underlying trend in the time series (see Figure 2).

Inference of the TV-AR model: Splines and generalized additive models

In this section, we discuss how to estimate the time-varying parameters in the TV-AR model using the generalized additive model (GAM) framework. GAM models are expanded general linear models (GLMs), such that one or more terms are replaced with a non-parametric (smooth) function (Keele, 2008; Wood, 2006). This makes GAM models semi-parametric models, since predictor variables (i.e., in our case y_{t-1}) can either be modeled as in standard regression (e.g., β_1) or in a non-parametric way (e.g., $\beta_{1,t}$). We focus in this section on the nonparametric representation. Code for the figures can be found in the *R-code* in the Appendix under the heading *IV. Inference of the TV-AR model*.

The non-parametric smooth functions used here are based on regression splines. Regression splines are piecewise polynomial functions that are joined (smoothly) at breakpoints called knots (Hastie & Tibshirani, 1990). In order to clarify the concept further, we will give a simulated example (based on Wood, 2006). Specifically, data are simulated for $n = 20$ time points according to a sine wave: $y_t = \sin\left(\frac{2\pi t}{20}\right) + \epsilon_t$, where $\epsilon_t \sim N(0, 0.3^2)$. We denote the time points in the data as t_i with $i = 1, \dots, 20$. The data are represented as the small black dots in the first and last panel of Figure 3. To fit these data, we start with a simplified TV-AR model

$$y_t = \beta_{0,t} + \varepsilon_t \quad (7)$$

with only a time-varying intercept and no autoregressive parameter.

The goal is to find the function $\beta_{0,t}$ that tracks the general relation between y and t (which for this example is the sine wave underlying the data) as well as possible. In order to find the optimal smooth function estimating $\beta_{0,t}$, the following penalized least squares loss function is minimized:

$$\sum_{i=1}^n [y_i - \beta_{0,t_i}]^2 + \lambda \int_{-\infty}^{+\infty} [\beta_{0,t}'']^2 dt. \quad (8)$$

In the first part of Equation 8 one can recognize the ordinary least squares minimization $\sum_{i=1}^n [y_i - \beta_{0,t_i}]^2$, which measures the distance between the function and data points. The

last part is the roughness penalty $\lambda \int_{-\infty}^{+\infty} [\beta''_{0,t}]^2 dt$. This is an integrated squared second derivative that defines wiggleness, since the second derivative is a measure of curvature of the function whereas the integral sums up this measure along the entire domain of the function (Keele, 2008). Note that the square is needed to treat negative and positive curvature identically. The λ is a tuning parameter that controls the smoothness of the function. Small values of λ practically eliminate the penalty, thereby not penalizing for wiggleness and opening the possibility for wiggly functions. Large values of λ give a lot of weight to the penalty, thereby penalizing for wiggleness and restricting the possibility for wiggly functions. Minimizing the whole function leads to an optimal trade-off between goodness of fit and smoothness.⁹

The solution to the problem in Equation 8, denoted $\hat{\beta}_{0,t}$, can be expressed as a finite weighted sum of a set of predefined functions, known as basis functions. This can be written as follows:

$$\hat{\beta}_{0,t} = \hat{\alpha}_1 R_1(t) + \hat{\alpha}_2 R_2(t) + \hat{\alpha}_3 R_3(t) + \cdots + \hat{\alpha}_K R_K(t), \quad (9)$$

where we have expressed the solution in terms of K basis functions $R_1(t), \dots, R_K(t)$ and t represents the predictor variable (time, in our case). The basis functions can be evaluated at every time t_i in the data and therefore the values $R_1(t_i), \dots, R_K(t_i)$ can be collected in a $n \times K$ design matrix X so that the optimal regression weights can be determined by linear regression methods (see below).

Various options exist for choosing the smoothing basis, that is, the set of basis functions R_1 to R_K . Commonly used smoothing bases are *cubic regression splines* and *thin plate regression splines* (the latter being the default setting in the package *mgcv*), which represent alternative strategies with different properties of how the basis functions are chosen (Wood, 2006). Cubic regression splines are segmented cubic polynomials joined at the knots, and are constrained to be continuous at the knot points as well as to have a continuous first and second derivative (Fitzmaurice, Davidian, Verbeke, & Molenberghs, 2008). With cubic regression splines the locations of knots have to be chosen, the default

setting in the *mgcv* package being that the knot points are automatically placed (equally spaced) over the entire range of data.

In contrast, the thin plate regression splines approach automatically starts with one knot per observation and then uses an eigen-decomposition to find the basis coefficients that maximally account for the variances in the data. Thus, thin plate regression splines circumvent the choice of knot locations, reducing subjectivity brought into the model fit (Wood, 2006). Furthermore, unlike cubic regression splines, thin plate regression splines can handle smoothing in high-dimensional problems (e.g., when multiple independent variables occur). However, in one-dimensional problems, such as the one considered here, cubic and thin plate regression splines will lead to very similar solutions.

For our example, we have chosen a thin plate regression spline smoothing basis with $K = 6$ basis functions. The six basis functions are plotted in the panels 2-7 of Figure 3. The first two basis function are defined as $R_1(t) = 1$ and $R_2(t) = t$. Here one can recognize the constant and the first predictor variable of a standard linear regression model. The other four basis functions ($R_3 - R_6$) have a more complicated shape (for examples of such functions, see Gu, 2002; Keele, 2008; Wood, 2006). Additionally, in thin plate regression every basis function that is added is wigglier than the previous basis function. For example, basis function R_6 is wigglier than R_5 . Note that in contrast to cubic splines, where the basis functions depend on the knot location, in thin plate splines a basis function cannot be associated with a knot location. Furthermore, the basis functions are evaluated at every value of t (also with the cubic spline smoothing basis). This is important to point out, as regression splines are defined as segmented polynomials that are joined at the knot points, so evaluations of the basis functions may prima facie seem to be restricted to particular segments.

After choosing the smoothing basis and the number of basis functions, estimating the time-varying function $\beta_{0,t}$ simply boils down to the estimation of the weights (denoted as α_i above) of the linear combination in a penalized regression sense (see below). In Figure 3,

the final panel shows the weighted basis functions as well as the sine wave that is the final smooth function (i.e., $\hat{\beta}_{0,t}$, the thick dashed line).

Using a regression spline based method to estimate a smooth function raises the question of how many basis functions are needed to get a good fit. The usual approach is to place more basis functions than can reasonably be expected to be necessary, after which the function's smoothness is controlled by the roughness or wiggleness penalty as described earlier ($\lambda \int_{-\infty}^{+\infty} [\beta''_{0,t}]^2 dt$; see Wood, 2006). An attractive feature of spline regression methods is that the penalized loss function eventually boils down to a relatively simple penalized regression problem (see Wood, 2006). Thus, one can choose a reasonably large number of basis functions (so that in principle even very wiggly functions can be handled by the model), but then too wiggly components of the basis functions that are unnecessary are downplayed based on the value of the penalization tuning parameter λ . For instance, in our example the wiggliest basis function R_6 (panel 7 in Figure 3) is clearly penalized, as it appears as an almost flat horizontal line in the last panel of Figure 3.

Of course, the next question is then: What is a good value for the penalty parameter λ ? If the value of λ is too small, the estimated function is not smooth enough, but if λ is set too high, the function may oversmooth the data. Commonly, the optimal value of λ is determined using the *generalized cross validation* method (GCV; Golub, Heath, & Wahba, 1979). The idea of (ordinary) cross validation is that first a model, in this case a regression spline with a certain value of λ , is fitted on part of the data, for example leaving one datum out. In a second step, it is measured how well the estimated model can predict the other part of the data, for example the datum that was left out. However, with splines this process is computationally intensive and sensitive to transformations of the data (Wood, 2006). Therefore, the generalized cross validation score is used instead, which follows the same principle, but is invariant to transformations (Keele, 2008). The lowest GCV score indicates the optimal λ value and thus optimal smoothness of the estimated smooth function.

All of the above steps are implemented in the *mgcv* package in *R* (Wood, 2006). Using this software, one only has to define sufficiently many basis functions. The default for all splines is 10 basis functions. For the current example, detecting the relation between y and t , the command in *R* would be `gam(y~s(t,bs='tp',k=6))`, where the function `s` indicates the use of a smooth function for its argument (the predictor `t` in this case), `bs` indicates which smoothing basis is used (thin plate in this case), and `k` indicates the number of basis functions (see also the *R-code* in the Appendix). In addition to the GCV score and the estimated smooth function, the *mgcv* package also provides 1) p -values, 2) a measure of nonlinearity (edf and ref.df), 3) 95% confidence intervals (CIs) and 4) model fit indices, all of which we elaborate on below.

1. The p -values for the smooth function result from a test of the null hypothesis that the smooth time-varying function is actually zero over the whole time range (Wood, 2013).
2. As non-parametric smooth functions (such as $\beta_{0,t}$) are difficult to represent in a formulaic way, a graphical representation is usually needed to get insight into the form of the function (see for instance Figure 3; Faraway, 2006). However, besides a plot of the smooth function, the *mgcv* package also provides a measure of nonlinearity in the form of the *effective degrees of freedom* (edf). Basically, the edf refers to the number of parameters needed to represent the smooth functions. At first sight, one may think that this is equal to the number of basis functions, but because of the penalization that is not the case. The reason why the penalization decreases the effective degrees of freedom is that the parameters are not free to vary because of the penalizations (Wood, 2006). The higher the edf, the more wiggly the estimated smooth function is, and an edf of 1 indicates a linear effect (Shadish, Zuur, & Sullivan, 2014). Furthermore, the edf also gives an indication of how much penalization took place and thus may serve as a diagnostic: The closer the edf is to the number of basis functions, the lower the penalization. Usually, an edf close to the

number of basis functions means that additional basis functions should be added to capture the shape of the function. The `ref.df` is the reference degree of freedom used for hypothesis testing (Wood, 2013).

3. The 95% confidence intervals (CIs) around the smooth curve reflect the uncertainty of the smooth function. As the confidence intervals are obtained through a Bayesian approach, they are strictly speaking credible intervals, or Bayesian confidence intervals as referred to by Wood (see Wood, 2006).
4. Finally, model selection criteria can be retrieved with the package (such as BIC and AIC), where the lowest fit indices indicate the best model fit. When using the BIC and AIC for penalized models, note that the degrees of freedom are determined by the `edf` number and not by the number of parameters (see for more information Hastie & Tibshirani, 1990).

We have assumed a simple model with only a time-varying intercept to explain the fundamentals of splines. For the more realistic general TV-AR model, the time-varying autoregressive function is estimated in a similar way (see for further information Wood, 2006).

Guidelines regarding the TV-AR model: a simulation study

To evaluate how the TV-AR model performs under different circumstances using the default settings, we carried out a simulation study. In addition, we investigated the robustness of our method against violations of the assumption of gradual change, by considering also functions that change non-gradually. We will give here a general overview of the simulation conditions. In the supplementary material the simulation setup is described in detail. In addition, there is *R-code* in the Appendix exemplifying some of the simulation results under the heading *V. Guidelines TV-AR model*.

In the simulation study, we varied three factors: the generating function, low or high

values for the model parameters, and the sample size. First, we had 5 generating functions for the intercept $\beta_{0,t}$ and the autoregressive parameter $\beta_{1,t}$: 1) both are invariant over time, 2) both increase linearly over time, 3) both follow a cosine function over time, 4) both follow a random walk and 5) both follow a stepwise function (see also Figure 4). Note that the random walk and the stepwise function are non-gradually changing functions. Strictly, the TV-AR model is thus not expected to recover these functions. Instead, we consider these functions to investigate the robustness of TV-AR in non-gradual conditions. The second factor we varied was the maximum absolute values of the parameters (low or high maximum value). The third factor was sample size (30, 60, 100, 200, 400, 1000).

Estimation was executed using five models: A) a TV-AR model using the default settings (a thin plate regression spline basis using 10 basis functions); B) a TV-AR model with only a time-varying intercept and a time-invariant autoregressive parameter using the default settings; C) a TV-AR model with only a time-varying autoregressive parameter using the default settings; D) a standard time-invariant AR model; and E) a thin plate regression spline basis using 30 basis functions.

We evaluated the estimates of all models with mean squared errors (MSE) and coverage probabilities (CP) (see the Appendix for a detailed explanation of these measures). Furthermore, we analyzed how well the BIC, AIC and GCV could distinguish between time-varying and time-invariant processes. Last, we looked at the significance of the parameters and the effective degrees of freedom (edf) if applicable.

Results and guidelines

The results show that the time-varying AR model was able to estimate all gradually changing generating functions (invariant, linear, cosine) very well using the default settings of the *mgcv* package in *R* (i.e., using 10 basis functions and thin plate regression splines; see Figure 4 and 5). Around 200 time points were needed for detecting a small change, such as a small linear increase over time, but large changes could already be detected with

60 time points.

In general, none of the model selection methods (BIC, AIC and GCV) performed well in selecting the correct model out of models A, B, C and D (e.g., with 100 time points in the high condition of the linear increase, the BIC selects the correct model (model A) in only 60% of the cases). However, the BIC does relatively well in distinguishing between the time-invariant model D and the time-varying models (the three variants A, B and C combined). For example, with 100 time points in the high condition of the linear increase, the BIC selects the correct class (invariant versus time-varying) in circa 97% of the cases.¹⁰

As the BIC cannot be used for selecting the exact time-varying model (model A, B or C), additional criteria are needed. One possibility is to fit a TV-AR model and check the significance of the parameters (intercept and autoregressive parameter). If the intercept is significant, one can be confident that the intercept is time-varying, especially with at least circa 100 time points. This is because the TV-AR model automatically includes an (standard time-invariant) intercept, and significance implies that another, time-varying, intercept is needed. In contrast, in the case of the autoregressive parameter, significance entails that the parameter is valuable for the model, and thus should be kept, but it does not give information about whether it is a time-varying parameter or not. Additionally, a high edf is an indication that the parameter is time-varying, but note that the edf cannot be used to discriminate between time-invariant parameters and linearly increasing time-varying parameters, as they will often both have an edf of circa 2.

Even when the assumption of gradual change was violated, the TV-AR model was still able to estimate the general pattern of change (i.e., the trend-like fluctuations in the random walk), but not abrupt changes (such as in the stepwise function) or fast changes (i.e., the small-magnitude fluctuations in the random walk process). An exception was the condition with 1000 time points of the stepwise function, where the large jump could be detected quite well (see Figure 5). To get satisfying estimations in these cases, more time points are needed, and the amount of basis functions should be large enough. In general, it

is advisable to always check whether you have enough basis functions. A good indication that you do not have enough basis functions and should increase their number is that the effective degrees of freedom (edf) come close to the number of basis functions (Wood, 2006). The simulation study showed that the average coverage probabilities of especially the non-gradually changing functions are clearly improved by increasing the number of basis functions (in this case from 10 to 30 basis functions; see Table 1). This lines up well with the advice given in general to have a high enough number of basis functions to allow for enough wiggleness in the estimated function (Wood, 2006).

An empirical example

We applied the TV-AR model to data of two individuals who took part in a long isolation study, the MARS500 project, in which psychological and physiological data have been collected to study the effects of living in an enclosed environment for the duration of a real potential mission to Mars (i.e, 520 days; for more information see <http://www.esa.int/Mars500>). We focus here on emotional inertia, which is studied in the context of affective research. Emotional inertia is defined as the temporal dependency of individual emotions, or the self-predictability of emotions, and is typically modeled with an AR model (Kuppens et al., 2010; Suls, Green, & Hillis, 1998). However, a study by Koval and Kuppens (2012) showed that emotional inertia is not a trait-like characteristic, but is itself prone to change, causing the data to be non-stationary (see also de Haan-Rietdijk et al., 2014; Koval et al., in press). They showed, among other things, that individuals who anticipated a social stressor had a significant decrease in their emotional inertia, which means that to model the process of inertia correctly, the autoregressive parameter should be allowed to vary over time. In the MARS500 example, being isolated can be seen as a social stressor. Furthermore, it is plausible that the longer one is isolated, the more social stress there is. To study if and how inertia changed due to social isolation,

we analyzed time series data from two persons involved in the MARS500 study using the TV-AR model.

Method

Data description. The MARS500 study consisted of six healthy male participants (average age was 34 years), who all signed a written informed consent before participating in this experiment. In accordance with the Declaration of Helsinki, the protocol was approved by The Ethics Committee of the University Hospital Gasthuisberg of Leuven (Belgium) and the ESA Medical Board before the research was conducted. We focus here on the dynamics of the variable ‘valence’ of two participants. Each morning, the participants indicated on a 21×21 grid how they were feeling at that moment. The horizontal axis of the grid referred to valence and the vertical axis to arousal. Only the valence score (on 21-point scale) will be analyzed here. A high score indicates experience of highly positive feelings, and a low score experience of highly negative feelings.¹¹ There was 29% and 18% missingness in the data of participant 1 and 2 respectively (see Figure 6 for the raw data).¹²

Analyses. We consider the following four models:

In *model 1*, both the intercept and the autoregressive parameter are allowed to vary over time. The time-varying autoregressive parameter implies that the temporal dependency or emotional inertia (i.e., how self-predictable the emotion is) changes over time. Since the mean (or the attractor of the process) is a function of the intercept and the autoregressive parameter, it most likely also changes over time in this model:¹³

$$Valence_t = \beta_{0,t} + \beta_{1,t}Valence_{t-1} + \varepsilon_t. \quad (10)$$

In *model 2*, the intercept is allowed to fluctuate over time, but the autoregressive parameter is fixed over time, meaning that the temporal dependency (or emotional inertia) is time-invariant. Due to the changing intercept, the person’s attractor also changes over

time:

$$Valence_t = \beta_{0,t} + \beta_1 Valence_{t-1} + \varepsilon_t. \quad (11)$$

In *model 3*, the intercept is fixed over time, while the autoregressive parameter is allowed to vary over time. As indicated in the description of *model 1*, a time-varying autoregressive parameter means that the temporal dependency (or emotional inertia) of the process changes over time. However, fixing the intercept implies that the attractor changes over time, but this is fully accounted for by changes in the temporal dependency (i.e., the autoregressive parameter):

$$Valence_t = \beta_0 + \beta_{1,t} Valence_{t-1} + \varepsilon_t. \quad (12)$$

Finally, *model 4* is the standard AR(1) model, in which both the intercept and the autoregressive parameter are time-invariant; as a result the mean (i.e., a time-invariant attractor) is also fixed over time. This means that the temporal dependency (or emotional inertia) is completely constant over time, that is, both the temporal dependency (or emotional inertia) and the attractor value of the process remain the same over time:

$$Valence_t = \beta_0 + \beta_1 Valence_{t-1} + \varepsilon_t. \quad (13)$$

Following the guidelines presented in the previous section, we first checked if the process was time-varying or not. For this purpose, we used the BIC: If the BIC selects model 1, 2 or 3 the process is probably changing over time, and otherwise (i.e., if model 4 is selected) the process is probably time-invariant. In the latter case, a standard AR model should be used; otherwise a TV-AR model is appropriate. Secondly, to check which parameters are time-varying, we considered whether the smooth parameters were significantly different from zero and thus were needed in the model. As noted before, a significant intercept indicates that this parameter is time-varying, whereas a significant autoregressive parameter does not entail that it is time-varying. Therefore, in a third step, when the autoregressive parameter was significant we checked if the edf was higher than 1.

Additionally, we checked whether the residuals (estimated innovations $\hat{\varepsilon}_t$) indicated autocorrelation over time, satisfied the equal variance assumption and were normally distributed.

The analyses reported here were based on the default settings, that is, a thin plate regression spline basis with 10 basis functions (i.e., $K = 10$). We also ran all of the analyses with a cubic regression spline basis and thin plate regression splines with 30 basis functions (i.e., $K = 30$), but all results were highly similar and led to the same conclusions.

Results

As can be seen in Figure 6 (left panel), in the data of participant 1, a clear trend is apparent, whereas the data for participant 2 do not contain any clear time trend (Figure 6 right panel). For both participants the assumptions held for the selected models: The residuals for both participants did not indicate any autocorrelation over time, did not violate the equal variance assumption and were normally distributed.

For participant 1, the BIC indicated that the underlying process was varying over time and thus non-stationary (*model 2* was selected as the best model, although the differences between model 1 and 2 were fairly small, see Table 2). Consequently, fitting the TV-AR model showed that the function of the intercept was significantly different from zero ($F = 3.42$, $p = 0.0046$, $edf = 4.50$, $ref.df = 5.20$), while the function of the autoregressive parameter was not ($F = 0.87$, $p = 0.51$, $edf = 5.01$, $ref.df = 5.62$). Thus, only a time-varying intercept was needed in the TV-AR model. Based on visually inspecting Figure 7, the function of the intercept process (upper panel) is clearly varying over time, whereas the CIs of the function of the autoregressive parameter (middle panel) always include zero (the zero is represented by the dashed gray line) and the function does not clearly go up or down at any point in time. Taking all of these considerations into account, *model 2*, with a time-varying intercept and a time-invariant autoregressive parameter of zero, seems to be the best fitting model.

For participant 2, the BIC indicated that *model 3* had the best model fit and thus a TV-AR model was estimated. In line with this result, *model 1* (Equation 10) implied that the function of the autoregressive parameter was significant and should be kept in the model ($F = 8.32$, $p < 0.0001$, $edf = 5.17$, $ref.df = 6.15$), while the function of the intercept was not significant and thus time-invariant ($F = 0.15$, $p = 0.70$, $edf = 1.00$, $ref.df = 1.00$). Although significance does not imply that the autoregressive parameter is time-varying, the edf was clearly higher than 1. In addition, visual inspection of Figure 8 also clearly indicates that the autoregressive function (middle panel) of participant 2 changes over time. Thus, *model 3*, with a time-invariant intercept and a time-varying autoregressive parameter, seems to be the best model.

In sum, in the data for participant 1, no inertia or autocorrelation of valence in the data is apparent, but rather it is the intercept that changes (see Figure 7 panel 3). In this specific case, the attractor is equal to the intercept as the autoregressive parameter equals zero. Participant 1 simply feels less happy as the isolation experiment proceeds, as represented by the changing intercept and attractor. This is not necessarily in contradiction with the results found by Koval and Kuppens (2012) as we do not know how much emotional inertia participant 1 had before the isolation experiment. It is possible, for example, that this participant had some level of emotional inertia before going into isolation, but as soon as the experiment started, his emotional inertia decreased to zero, which would be in line with the previous findings of Koval and Kuppens (2012). In contrast, participant 2 starts the isolation experiment relatively happy and with a high spill-over of valence (high inertia), but already after a few days, his inertia decreases until it gets to zero around 100 days, and also his valence becomes more negative (see the attractor in the last panel of Figure 8). Towards the end of the experiment, there is again a light increase in his feeling of happiness and his inertia. This result is in line with research of Koval and colleagues, which suggests that as stress increases (the longer one is isolated) inertia decreases, and thus affect becomes less predictable (Koval & Kuppens, 2012).

Note that if one had ignored this non-stationarity in the data, a standard autoregressive model (thus, *model 4*) would have led to inaccurate conclusions about these two participants. For participant 1, ignoring non-stationarity would have led to inferring a highly significant autoregressive coefficient ($\beta_1 = 0.85$, $t(325) = 27.43$, $p < 0.0001$), that is, an extremely high inertia or a high predictability of his valence. For participant 2, ignoring non-stationarity would have led to the conclusion that there was a positive inertia ($\beta_1 = 0.20$, $t(420) = 4.29$, $p < 0.0001$), and the fact that his inertia was actually varying over time would have gone unnoticed.

In general, even though inertia is already well known to vary in strength greatly across individuals, it is still often studied as a trait of an individual. With the TV-AR model we can study inertia throughout the whole study period, creating an inertia value for every single time point. In future studies, it would be fruitful to take into account that inertia can change over time, even from day to day or faster, and of course, also in other contexts than social stress.

Furthermore, these two applications show how important it is in general to use a TV-AR model, as different conclusions would have been drawn with a standard AR model. In addition, with the TV-AR model trends as well as (time-varying) autoregressive parameters can be detected in one step: Even though the first example above (participant 1) involves a trend-stationary process and pre-specifying the exact (non-linear as the edf of 4.50 indicates) trend would have led to the same conclusions, this would have been much more difficult than with the TV-AR model. Psychological data can be non-stationary for various reasons, and the TV-AR model offers a simple exploratory tool for detecting such changing dynamic processes.

Discussion

In this paper, we have introduced a new way to study changing dynamics: the semi-parametric TV-AR model. This model fills a gap in the literature, because most

standard autoregressive models do not take into account non-stationarity, even though many psychological processes are likely to be non-stationary. Therefore, there is a need for an easily applicable method for studying such non-stationarity or changing dynamics. The semi-parametric TV-AR model presented in this article is exactly such a tool.

As shown by the simulations and application in this paper, the TV-AR model can estimate non-stationary processes well and has significant potential for studying changing dynamics in psychology. For example, the TV-AR model can help to detect and specify different kinds of non-stationarity in the data. Currently, it is common practice to focus on the trend that is apparent in the data, and to transform the time series so that it becomes trend stationary. However, even if the trend could be perfectly specified, which is often difficult, non-stationarity may not be fully accounted for, since the autocorrelation structure of the data can also change over time. Furthermore, a changing autocorrelation is not easy to detect visually, nor is there a test to detect such non-stationarity. With the semi-parametric TV-AR model, all such problems can be dealt with in one single step: Trends in the data and changes in the autoregressive process can be detected at once, and even more importantly, no pre-specifications are necessary, as has been shown in the real data application.

It is therefore clear that the semi-parametric TV-AR model is important in the case of non-stationary data. However, its potential range of application is much broader. As little is known about how and when psychological dynamics change, we would recommend to always run a TV-AR model next to a standard AR model as part of regular analysis if enough time points (circa 100) are available. In this way, the model can be used as a diagnostic tool for probing whether there is non-stationarity in the time series, and for detecting and specifying changing dynamics, such as the trend. For example, if the time series turns out to have a trend that is linear instead of non-parametric, a simpler parametric model can be specified based on the TV-AR analyses.

We have considered the simplest form of a TV-AR model, and will now elaborate on

some of the extensions that are possible. We studied temporal dependency with a lag order 1 TV-AR model, but one can imagine that the temporal dependency is not only apparent between the two closest occasions, but also between occasions further apart, in which case a TV-AR model with lag order 2 or larger is necessary. Such extra lags can be easily added into a TV-AR model in the same manner as they are added into standard AR models through the inclusion of more lagged predictors.

Another sensible extension involves generalization of the model to multivariate data. The TV-AR model is currently only applicable to the univariate case, while it is often more realistic that a variable is not only predicted by itself, but also by other variables, which evokes the need to analyze psychological dynamics as a multivariate system. Such an extension would lead to a time varying vector AR (TV-VAR) model, and comes with new challenges, as both auto-correlations and cross-correlations would have to be modeled in this case. Yet another natural, but even more challenging, extension would be a TV-AR multilevel extension based on current multilevel (V)AR models (Bringmann et al., 2013; de Haan-Rietdijk et al., 2014; Jongerling, Laurenceau, & Hamaker, 2015). To the best of our knowledge, this is currently not possible, as the *mgcv* software cannot be used to estimate a flexible smooth function for the population (i.e., the population average) and to allow for flexible interindividual variation for that smooth function. An additional extension could be time-varying error variance, so that also the time-varying variance of a process could be fully accounted for. However, with current software, only the intercept and the autoregressive parameter (and not the error variance) can be modeled as time-varying parameters. Further research should also consider the combination of gradual and abrupt changes, so that when the point of an abrupt change is known, it could be easily adjusted in the TV-AR model.

Even though the TV-AR model is easily applicable, the number of time points needed is a potential limitation. While 100 time points per participant would be preferable, currently most longitudinal studies in psychology gather around 60 time points or less

(aan het Rot et al., 2012). Another limitation of the TV-AR is the assumption of gradual change. Although we have shown in the simulation study that with many time points and a large abrupt change the TV-AR model is quite robust and still gives an indication of the sudden jump, other models are probably more suitable for studying sudden change. Such models include the threshold autoregressive model (TAR) (e.g., Hamaker, 2009; Hamaker, Grasman, & Kamphuis, 2010), its multilevel extension, multilevel TAR (de Haan-Rietdijk et al., 2014), or the regime-switching state-space model (cf. Hamaker & Grasman, 2012; Kim & Nelson, 1999).

Furthermore, as the semi-parametric TV-AR model is an exploratory tool, the standard errors of the time-varying parameters are likely to be less satisfactory compared to confirmatory, raw-data maximum likelihood approaches, such as the state-space approach. Additionally, estimating a TV-AR model in a state-space modeling framework has the advantage that measurement error can be taken into account, which is not possible with the semi-parametric TV-AR model (Schuurman, Houtveen, & Hamaker, 2015). Thus, future research should aim at comparing the exploratory semi-parametric TV-AR model with confirmatory approaches.

In sum, the semi-parametric TV-AR model presented here is an easy to use tool for detecting and modeling non-stationarity. Many extensions are possible, and future research is needed to uncover all the possibilities and limitations of this innovative framework. By introducing the model and explaining its application in standard software, we hope to have made it available to a broad range of psychologists studying human dynamics.

References

- aan het Rot, M., Hogenelst, K., & Schoevers, R. A. (2012). Mood disorders in everyday life: A systematic review of experience sampling and ecological momentary assessment studies. *Clinical Psychology Review*, *32*(6), 510–523.
- Adolph, K. E., Robinson, S. R., Young, J. W., & Gill-Alvarez, F. (2008). What is the shape of developmental change? *Psychological Review*, *115*(3), 527–543.
- Basner, M., Dinges, D. F., Mollicone, D., Ecker, A., Jones, C. W., Hyder, E. C., . . . Sutton, J. (2013). Mars 520-d mission simulation reveals protracted crew hypokinesia and alterations of sleep duration and timing. *Proceedings of the National Academy of Sciences*, *110*(7), 2635–2640.
- Belsley, D. A., & Kuh, E. (1973). Time-varying parameter structures: An overview. *Annals of Economic and Social Measurement*, *2*(4), 375–379.
- Bisconti, T. L., Bergeman, C., & Boker, S. M. (2004). Emotional well-being in recently bereaved widows: A dynamical systems approach. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, *59*(4), 158–167.
- Boker, S. M., Molenaar, P. C. M., & Nesselroade, J. R. (2009). Issues in intraindividual variability: Individual differences in equilibria and dynamics over multiple time scales. *Psychology and Aging*, *24*(4), 858–862.
- Boker, S. M., Rotondo, J. L., Xu, M., & King, K. (2002). Windowed cross-correlation and peak picking for the analysis of variability in the association between behavioral time series. *Psychological Methods*, *7*(3), 338–355.
- Bolger, N., Davis, A., & Rafaeli, E. (2003). Diary methods: Capturing life as it is lived. *Annual Review of Psychology*, *54*(1), 579–616.
- Bolger, N., & Laurenceau, J.-P. (2013). *Intensive longitudinal methods: An introduction to diary and experience sampling research*. New York, NY: Guilford Press.
- Brandt, P. T., & Williams, J. T. (2007). *Multiple time series models. series: Quantitative applications in the social sciences*. Thousand Oaks, CA: Sage Publications Inc.

- Bringmann, L. F., Lemmens, L. H., Huibers, M. J., Borsboom, D., & Tuerlinckx, F. (2014). Revealing the dynamic network structure of the Beck Depression Inventory-II. *Psychological Medicine*, 1–11.
- Bringmann, L. F., Vissers, N., Wichers, M., Geschwind, N., Kuppens, P., Peeters, F., ... Tuerlinckx, F. (2013). A network approach to psychopathology: New insights into clinical longitudinal data. *PloS One*, 8(4), e60188.
- Chatfield, C. (2003). *The analysis of time series: An introduction* (Fifth ed.). Boca Raton, FL: Chapman and Hall/CRC.
- Chow, S.-M., Zu, J., Shifren, K., & Zhang, G. (2011). Dynamic factor analysis models with time-varying parameters. *Multivariate Behavioral Research*, 46(2), 303–339.
- Dahlhaus, R. (1997). Fitting time series models to nonstationary processes. *The Annals of Statistics*, 25(1), 1–37.
- Deboeck, P. R. (2013). Dynamical systems and models of continuous time. In T. D. Little (Ed.), *The Oxford handbook of quantitative methods in psychology: Vol. 2: Statistical analysis* (pp. 411–431). Oxford, England: Oxford University Press.
- de Haan-Rietdijk, S., Gottman, J. M., Bergeman, C. S., & Hamaker, E. L. (2014). Get over it! A multilevel threshold autoregressive model for state-dependent affect regulation. *Psychometrika*, 1–25.
- 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.
- Fan, J., & Yao, Q. (2003). *Nonlinear time series: Nonparametric and parametric methods*. New York, NY: Springer.
- Faraway, J. J. (2006). *Extending the linear model with R: Generalized linear, mixed effects and nonparametric regression models*. Boca Raton, FL: Chapman and Hall/CRC.
- Ferrer, E., & Nesselroade, J. R. (2003). Modeling affective processes in dyadic relations via dynamic factor analysis. *Emotion*, 3(4), 344–360.

- Fitzmaurice, G., Davidian, M., Verbeke, G., & Molenberghs, G. (2008). *Longitudinal data analysis*. Boca Raton, FL: Chapman and Hall/CRC.
- Giraitis, L., Kapetanios, G., & Yates, T. (2014). Inference on stochastic time-varying coefficient models. *Journal of Econometrics*, 179(1), 46–65.
- Golub, G. H., Heath, M., & Wahba, G. (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2), 215–223.
- Gu, G. (2002). *Smoothing spline anova models*. New York: Springer.
- Hamaker, E. L. (2009). Using information criteria to determine the number of regimes in threshold autoregressive models. *Journal of Mathematical Psychology*, 53(6), 518–529.
- Hamaker, E. L. (2012). Why researchers should think within-person: A paradigmatic rationale. In M. R. Mehl & T. S. Conner (Eds.), *Handbook of research methods for studying daily life* (pp. 43–61). New York, NY: Guilford.
- Hamaker, E. L., Ceulemans, E., Grasman, R. P. P. P., & Tuerlinckx, F. (in press). Modeling affect dynamics: State-of-the-art and future challenges. *Emotion Review*.
- Hamaker, E. L., & Dolan, C. V. (2009). Idiographic data analysis: Quantitative methods from simple to advanced. In J. Valsiner, P. C. M. Molenaar, M. Lyra, & N. Chaudhary (Eds.), *Dynamic process methodology in the social and developmental sciences* (pp. 191–216). New York, NY: Springer-Verlag.
- Hamaker, E. L., Dolan, C. V., & Molenaar, P. C. M. (2003). Arma-based sem when the number of time points T exceeds the number of cases N : Raw data maximum likelihood. *Structural Equation Modeling*, 10(3), 352–379.
- Hamaker, E. L., & Grasman, R. (2012). Regime switching state-space models applied to psychological processes: Handling missing data and making inferences. *Psychometrika*, 77(2), 400–422.
- Hamaker, E. L., Grasman, R. P. P. P., & Kamphuis, J. H. (2010). Regime-switching models to study psychological processes. In P. C. M. Molenaar & K. Newell (Eds.),

- Individual pathways of change* (pp. 155–168). Washington, DC: American Psychological Association.
- Hamilton, J. D. (1994). *Time series analysis*. Princeton, NJ: Princeton university press.
- Härdle, W., Lütkepohl, H., & Chen, R. (1997). A review of nonparametric time series analysis. *International Statistical Review*, 65(1), 49–72.
- Hastie, T. J., & Tibshirani, R. J. (1990). *Generalized additive models*. Boca Raton, FL: Chapman and Hall/CRC.
- Hertzog, C., & Nesselroade, J. R. (2003). Assessing psychological change in adulthood: An overview of methodological issues. *Psychology and aging*, 18(4), 639–657.
- Jahng, S., Wood, P. K., & Trull, T. J. (2008). Analysis of affective instability in ecological momentary assessment: Indices using successive difference and group comparison via multilevel modeling. *Psychological Methods*, 13(4), 354–375.
- Jongerling, J., Laurenceau, J.-P., & Hamaker, E. L. (2015). A multilevel AR (1) model: Allowing for inter-individual differences in trait-scores, inertia, and innovation variance. *Multivariate Behavioral Research*, 50(3), 334–349.
- Keele, L. J. (2008). *Semiparametric regression for the social sciences*. Chichester, England: John Wiley & Sons.
- Kim, C.-J., & Nelson, C. R. (1999). *State-space models with regime switching: classical and gibbs-sampling approaches with applications*. Cambridge: MIT press.
- Kitagawa, G., & Gersch, W. (1985). A smoothness priors time-varying ar coefficient modeling of nonstationary covariance time series. *IEEE Transactions on Automatic Control*, 30(1), 48–56.
- Koop, G. (2012). Using VARs and TVP-VARs with many macroeconomic variables. *Central European Journal of Economic Modelling and Econometrics*, 4(3), 143–167.
- Koval, P., Brose, A., Pe, M. L., Houben, M., Erbas, Y. I., Champagne, D., & Kuppens, P. (in press). Emotional inertia and external events: The roles of exposure, reactivity, and recovery. *Emotion*.

- Koval, P., & Kuppens, P. (2012). Changing emotion dynamics: Individual differences in the effect of anticipatory social stress on emotional inertia. *Emotion, 12*(2), 256–267.
- Kuppens, P., Allen, N. B., & Sheeber, L. B. (2010). Emotional inertia and psychological maladjustment. *Psychological Science, 21*, 984–991.
- Kwiatkowski, D., Phillips, P. C., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics, 54*(1), 159–178.
- McArdle, J. J. (2009). Latent variable modeling of differences and changes with longitudinal data. *Annual Review of Psychology, 60*, 577–605.
- McKeown, G. J., & Sneddon, I. (2014). Modeling continuous self-report measures of perceived emotion using generalized additive mixed models. *Psychological Methods, 19*(1), 155–174.
- Mehl, M. R., & Conner, T. S. E. (2012). *Handbook of research methods for studying daily life*. New York, NY: Guilford Press.
- Molenaar, P. C. M. (1985). A dynamic factor model for the analysis of multivariate time series. *Psychometrika, 50*(2), 181–202.
- Molenaar, P. C. M. (1987). Dynamic assessment and adaptive optimization of the psychotherapeutic process. *Behavioral Assessment, 9*, 389–416.
- Molenaar, P. C. M., & Campbell, C. G. (2009). The new person-specific paradigm in psychology. *Current Directions in Psychological Science, 18*(2), 112–117.
- Molenaar, P. C. M., De Gooijer, J. G., & Schmitz, B. (1992). Dynamic factor analysis of nonstationary multivariate time series. *Psychometrika, 57*(3), 333–349.
- Molenaar, P. C. M., & Newell, K. M. (2003). Direct fit of a theoretical model of phase transition in oscillatory finger motions. *British Journal of Mathematical and Statistical Psychology, 56*(2), 199–214.
- Molenaar, P. C. M., Sinclair, K. O., Rovine, M. J., Ram, N., & Corneal, S. E. (2009). Analyzing developmental processes on an individual level using nonstationary time

- series modeling. *Developmental psychology*, 45(1), 260.
- Mumtaz, H., & Surico, P. (2009). Time-varying yield curve dynamics and monetary policy. *Journal of Applied Econometrics*, 24(6), 895–913.
- Nesselroade, J. R., & Molenaar, P. (2010). Emphasizing intraindividual variability in the study of development over the life span. In R. M. Lerner & W. F. Overton (Eds.), *The handbook of life-span development* (pp. 30–54). Hoboken, NJ: Wiley.
- Nesselroade, J. R., & Ram, N. (2004). Studying intraindividual variability: What we have learned that will help us understand lives in context. *Research in Human Development*, 1(1-2), 9–29.
- Oravecz, Z., Tuerlinckx, F., & Vandekerckhove, J. (2011). A hierarchical latent stochastic differential equation model for affective dynamics. *Psychological Methods*, 16(4), 468.
- Pole, A., West, M., & Harrison, J. (1994). *Applied Bayesian forecasting and time series analysis*. Boca Raton, FL: Chapman and Hall/CRC.
- Prado, R. (2010). Characterization of latent structure in brain signals. In S.-M. Chow, E. Ferrer, & F. Hsieh (Eds.), *Statistical methods for modeling human dynamics: An interdisciplinary dialogue* (pp. 123–153). New York, NY: Routledge, Taylor and Francis.
- Rosmalen, J. G., Wenting, A. M., Roest, A. M., de Jonge, P., & Bos, E. H. (2012). Revealing causal heterogeneity using time series analysis of ambulatory assessments: Application to the association between depression and physical activity after myocardial infarction. *Psychosomatic Medicine*, 74(4), 377–386.
- Schuurman, N. K., Houtveen, J. H., & Hamaker, E. L. (2015). Incorporating measurement error in n= 1 psychological autoregressive modeling. *Frontiers in psychology*, 6.
- Scollon, C. N., Prieto, C.-K., & Diener, E. (2003). Experience sampling: promises and pitfalls, strength and weaknesses. *Journal of Happiness Studies*, 4, 5–34.
- Shadish, W. R., Zuur, A. F., & Sullivan, K. J. (2014). Using generalized additive (mixed) models to analyze single case designs. *Journal of School Psychology*, 52(2), 149–178.

- Song, H., & Ferrer, E. (2012). Bayesian estimation of random coefficient dynamic factor models. *Multivariate Behavioral Research*, 47(1), 26–60.
- Stroe-Kunold, E., Wesche, D., Friederich, H.-C., Herzog, W., Zastrow, A., & Wild, B. (2012). Temporal relationships of emotional avoidance in a patient with anorexia nervosa—A time series analysis. *The International Journal of Psychiatry in Medicine*, 44(1), 53–62.
- Sullivan, K. J., Shadish, W. R., & Steiner, P. M. (2015). An introduction to modeling longitudinal data with generalized additive models: Applications to single-case designs. *Psychological methods*, 20(1), 26–42.
- Suls, J., Green, P., & Hillis, S. (1998). Emotional reactivity to everyday problems, affective inertia, and neuroticism. *Personality and Social Psychology Bulletin*, 24(2), 127–136.
- Tafforin, C. (2013). The mars-500 crew in daily life activities: An ethological study. *Acta Astronautica*, 91, 69–76.
- Tan, X., Shiyko, M. P., Li, R., Li, Y., & Dierker, L. (2012). A time-varying effect model for intensive longitudinal data. *Psychological Methods*, 17(1), 61–77.
- Tarvainen, M. P., Georgiadis, S. D., Ranta-aho, P. O., & Karjalainen, P. A. (2006). Time-varying analysis of heart rate variability signals with a Kalman smoother algorithm. *Physiological Measurement*, 27(3), 225–239.
- Tarvainen, M. P., Hiltunen, J. K., Ranta-aho, P. O., & Karjalainen, P. A. (2004). Estimation of nonstationary EEG with Kalman smoother approach: An application to event-related synchronization. *IEEE Transactions on Biomedical Engineering*, 51(3), 516–524.
- Trull, T. J., & Ebner-Priemer, U. (2013). Ambulatory assessment. *Annual Review of Clinical Psychology*, 9, 151–176.
- Tschacher, W., & Ramseyer, F. (2009). Modeling psychotherapy process by time-series panel analysis. *Psychotherapy Research*, 19(4-5), 469–481.
- van de Leemput, I. A., Wichers, M., Cramer, A. O., Borsboom, D., Tuerlinckx, F.,

- Kuppens, P., ... Scheffer, M. (2014). Critical slowing down as early warning for the onset and termination of depression. *Proceedings of the National Academy of Sciences*, 111(1), 87–92.
- Velicer, W. F., & Fava, J. L. (2003). Time series analysis. In J. A. Schinka & W. F. Velicer (Eds.), *Handbook of psychology: Research methods in psychology, vol. 2* (pp. 581–606). Hoboken, NJ: Wiley.
- Vigo, D. E., Tuerlinckx, F., Ogrinz, B., Wan, L., Simonelli, G., Bersenev, E., ... Aubert, A. E. (2013). Circadian rhythm of autonomic cardiovascular control during Mars500 simulated mission to Mars. *Aviation, Space, and Environmental Medicine*, 84(10), 1023–1028.
- Voelkle, M. C., & Oud, J. H. (2013). Continuous time modelling with individually varying time intervals for oscillating and non-oscillating processes. *British Journal of Mathematical and Statistical Psychology*, 66(1), 103–126.
- Voelkle, M. C., Oud, J. H., Davidov, E., & Schmidt, P. (2012). An SEM approach to continuous time modeling of panel data: Relating authoritarianism and anomia. *Psychological Methods*, 17(2), 176–192.
- Walls, T. A., & Schafer, J. L. (2006). *Models for intensive longitudinal data*. Oxford, England: Oxford University Press.
- Wang, L. P., Hamaker, E. L., & Bergeman, C. S. (2012). Investigating inter-individual differences in short-term intra-individual variability. *Psychological Methods*, 17(4), 567–581.
- Wang, Y., Jing, X., Lv, K., Wu, B., Bai, Y., Luo, Y., ... Li, Y. (2014). During the long way to Mars: Effects of 520 days of confinement (Mars500) on the assessment of affective stimuli and stage alteration in mood and plasma hormone levels. *PloS One*, 9(4), e87087.
- West, M., Prado, R., & Krystal, A. D. (1999). Evaluation and comparison of EEG traces: Latent structure in nonstationary time series. *Journal of the American Statistical*

Association, 94(446), 375–387.

Wood, S. N. (2006). *Generalized additive models: An introduction with R*. Boca Raton, FL: Chapman and Hall/CRC.

Wood, S. N. (2013). On p-values for smooth components of an extended generalized additive model. *Biometrika*, 100(1), 221–228.

Footnotes

¹Note that a time-varying effect model that also allows fitting a semi-parametric TV-AR model has recently been developed in SAS (Tan et al., 2012). However, it is less general and has fewer options for fitting a TV-AR model (e.g., at the moment it is only suitable for normally distributed time-varying models).

²In discrete time AR models the measurements of the process are assumed to be equally spaced, meaning that the distance between the measurements is the same through the whole study. If time points were not equally spaced, the autoregressive coefficient would have a different meaning across occasions. This is in contrast to continuous time AR models, where the intervals between time points do not have to be equal (see for more information: Bisconti, Bergeman, & Boker, 2004; Deboeck, 2013; Oravecz et al., 2011; Voelkle & Oud, 2013; Voelkle et al., 2012).

³The term dynamic error is used to pit this error against the well-known measurement error. The difference between the two error terms is that while measurement error is occasion-specific, affecting the scores only at a single occasion, dynamic error tends to affect subsequent occasions as well due to the underlying temporal dependency in the process (Schuurman et al., 2015). In the current study we restrict our focus to processes without measurement error.

⁴As we study normally distributed processes here, it is interesting to note that in this case covariance-stationarity implies strict stationarity, since a normal distribution is completely defined by its first two moments (Chatfield, 2003, p. 36).

⁵The intercept β_0 is the expected score when the observation at the previous occasion was zero (i.e., $y_{t-1} = 0$). When the scale that is used does not include the score zero, the intercept is typically not interesting.

⁶Note that a trend can be also caused by a unit root process, such as a random walk. In this case, the process has to be differenced in order to become stationary (see, for example, Hamilton, 1994).

⁷Note that in Giraitis et al. (2014) $\beta_{1,t}$ is specified as $\beta_{1,t-1}$. Here we use the standard notation used in Dahlhaus (1997).

⁸To derive a model-implied mean of the TV-AR, we can write

$$\begin{aligned}
 \mu_t &= E[\beta_{0,t} + \beta_{1,t}y_{t-1} + \varepsilon_t] \\
 &= E[\beta_{0,t}] + E[\beta_{1,t}y_{t-1}] + E[\varepsilon_t] \\
 &= \beta_{0,t} + \beta_{1,t}\mu_{t-1} \\
 &\approx \beta_{0,t} + \beta_{1,t}\mu_t
 \end{aligned} \tag{14}$$

where the latter approximation results from the fact that, in contrast to a standard AR model where we have $E[y_t] = E[y_{t-1}] = \mu$, the expectations of y_t and y_{t-1} are not exactly equal for a TV-AR model. However, since the parameters $\beta_{0,t}$ and $\beta_{1,t}$ are only allowed to change gradually, we can assume that μ_{t-1} is reasonably well approximated by μ_t , so that we have Equation 5. The derivation of the time-varying variance is similar to the derivation of the time-varying mean.

⁹Note that the least squares criterion can be used here because we assume continuous normally distributed data. In the more general case, the least squares criterion is replaced by minus the likelihood.

¹⁰Note that the AIC and GCV were not as accurate as the BIC. For example, with 100 time points in the high condition of the linear increase, the AIC and GCV selected the correct class (invariant versus time-varying) in only 73% and 76% of the cases respectively.

¹¹Although the measurement was done on a daily basis, on some days there were multiple measures, which was due to extra physiological tests that required additional measurements of valence and arousal. In these cases, we only used the first measure of the day.

¹²Note that the TV-AR model can also be used with missing data, although the more missingness the less power one has to detect the underlying process. Additionally, one has to assume that the missingness is (completely) at random.

¹³Of course it is possible, though unlikely, that the changes in the autoregressive parameter are exactly countered by the changes in the intercept (see Equation 5). In this case, the attractor would be time-invariant, while the temporal dependency would fluctuate over time.

Table 1

Coverage Probabilities (CP) of the autoregressive function in % using thin plate regression splines. Here the average CP of every simulation condition is given. Low and high stand for low and high value conditions for the maximum absolute values of the time-varying parameters. Note that the last line in the table uses the same settings as the previous line, except now 30 instead of 10 basis functions (K) are used.

N	True underlying function									
	<i>Invariant</i>		<i>Linear</i>		<i>Cosine</i>		<i>Random</i>		<i>Step</i>	
	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>
30	86	67	89	83	89	83	92	87	89	78
60	92	84	93	91	91	84	94	88	91	83
100	93	90	93	91	92	85	93	86	92	83
200	95	92	95	94	89	92	92	84	90	79
400	95	93	95	94	87	94	91	81	86	80
1000	95	95	95	95	89	96	86	78	82	82
1000 $K = 30$	95	94	94	95	91	96	87	83	84	87

Table 2

Model selection for participants 1 and 2 using the BIC indices. Lowest fit indices are in bold.

Model	BIC Participant 1	BIC Participant 2
Model 1	688	1,896
Model 2	684	1,894
Model 3	696	1,890
Model 4	868	1,899




Figure 1. Simulated time series with a positive (left) and a negative (right) autocorrelation for a valence process of a single individual. The valence process ranged from 0 to 10, with 0 indicating feeling very unhappy and 10 indicating very happy. The process was simulated for 150 time points with an intercept (β_0) of 3 (left) and 6 (right; see dashed line in both graphs) and an autoregressive coefficient (β_1) of 0.5 (left) and -0.5 (right), meaning that there was a positive (left) or negative (right) dependency in the data. Notice that here the intercept as such has no further meaning and is different from the mean. In the left graph, the mean (μ ; shown by the solid black line) is $3/(1 - 0.5) = 6$, indicating that on average this individual felt quite happy. In the right graph, the mean is $6/(1 + 0.5) = 4$, indicating that on average this individual felt slightly unhappy.




Figure 2. Simulated data of a valence process (with 0 indicating feeling very unhappy and 10 indicating very happy) with time-varying parameters. In the left panel, the autoregressive coefficient is time-invariant ($\beta_1 = 0.2$), while the intercept is time-varying ($\beta_{0,t}$; ranging from 3 to 5); in the right panel, the autoregressive coefficient is time-varying ($\beta_{1,t}$; gradually changing from 0.65 to 0.2), while the intercept is time-invariant ($\beta_0 = 2$). The attractor in the left panel (μ_t ; shown by the solid black line) changes from 4 to 7, indicating that this individual felt a bit unhappy at first, but at the end of the time series felt happy, whereas the attractor in the right panel changes from circa 6 to 2.5, indicating that this individual felt happy at first, but at the end felt unhappy.




Figure 3. The six basis functions for the curve $\beta_{0,t}$ using a cubic regression spline basis. Just as in standard regression, all basis functions $R_i(t)$ are weighed by multiplying them with their corresponding α_i coefficients. The contribution of each basis function to the solution is estimated using penalized regression and the $\hat{\beta}_{0,t}$ (the thick black dashed line in the bottom right panel) is a weighted sum.




Figure 4. Graphical representations of the generating functions of the autoregressive parameter for the low condition. The different true underlying functions $\beta_{1,t}$ are represented as thick black solid lines and the estimated $\hat{\beta}_{1,t}$ as grey solid lines, the grey dashed lines being the 95% CIs. The estimations are based on the median of the MSE values of the 1000 replications.




Figure 5. Graphical representations of the generating functions of the autoregressive parameter for the high condition. The different true underlying functions $\beta_{1,t}$ are represented as thick black solid lines and the estimated $\hat{\beta}_{1,t}$ as grey solid lines, the grey dashed lines being the 95% CIs. The estimations are based on the median of the MSE values of the 1000 replications.




Figure 6. The raw data of the variable valence for participant 1 (left) and participant 2 (right).




Figure 7. Estimation results for the TV-AR model for participant 1. Every panel represents a different parameter of the TV-AR model: the upper panel the intercept, the middle the autoregressive and the lowest the attractor. Note that the attractor process is plotted over the actual valence scores (represented in grey).




Figure 8. Estimation results for the TV-AR model for participant 2. Every panel represents a different parameter of the TV-AR model: the upper panel the intercept, the middle the autoregressive and the lowest the attractor. Note that the attractor process is plotted over the actual valence scores (represented in grey).

```
##%#####%##
##%#####%##
##### Appendix #####
##%#####%##
##%#####%##
```

```
#This file is part of the paper #Bringmann et al. 2015
#"Changing dynamics: TV-AR models using generalized additive modeling"
#in Psychological Methods.
```

```
#There are five main parts in this file:
```

```
#I. General code: This section contains the necessary packages and functions
#in order to run the code in the other sections.
```

```
#II. Standard time invariant AR: In this section we will provide the code
#to run an AR model.
```

```
#III. Time-varying AR: In this section we will provide the code
#to run a TV-AR model.
```

```
#IV. Inference of the TV-AR model: This section contains the code that results in
#Figure 3 of the paper.
```

```
#V. Guidelines TV-AR model: In this section some of the simulation results are
#exemplified.
```

```
##%#####%##
#####I. General code#####
##%#####%##
#First load the necessary packages:
library("mgcv")# for the gam function, version 1.8-3
library("hydroGOF")# for calculating the mean squared error (MSE), version 0.3-8
library("mvtnorm")# for the CI_mean function, version 1.0-2
```

```
###There are 5 auxiliary functions which are used throughout this code to simulate data###
```

```
##1. The five generating functions (genfun) options##
```

```
#1A. Time invariant
```

```
invariant<-function(N,MaxAbsValue) #N is the sample size, and MaxAbsValue is the maximum
absolute value of the function.
{genfun=rep(NA,(N)) #Creating the function first with NAs (missing values).
  genfun=rep(MaxAbsValue,(N)) #Here the actual invariant function is created.
  return(genfun)
}
```

```
#1B. Linear
```

```
linear<-function(N,MaxAbsValue,MinAbsValue) #N is the sample size, and MaxAbsValue is the
maximum absolute value of the function.
  #You can overwrite the MinAbsValue, which has been set to zero as a default.
{genfun=rep(NA,(N)) #Creating the function first with NAs (missing values).
  genfun=seq(MinAbsValue,MaxAbsValue,length.out=(N)) #Here the actual linear function is
  created.

  return(genfun)
}
```

#1C. Sine

```
sine<-function(N,MaxAbsValue) #N is the sample size, and MaxAbsValue is the maximum absolute
value of the function.
{genfun=rep(NA,(N)) #Creating the function first with NAs (missing values).
  tt=1:(N) #Defining a time parameter in order to create the sine function.
  genfun=MaxAbsValue*sin(2*pi*tt/(N)) #Here the actual sine function is created.
  return(genfun)
}
```

#1D. Random walk

```
#The code for the random walk is based on Giraitis et al. (2014).
random.walk<-function(N,MaxAbsValue) #N is the sample size, and MaxAbsValue is the maximum
absolute value of the function.
{genfun=rep(NA,(N)) # Here several parameters have to be created (using again NAs).
  a=rep(NA,(N))
  v=rep(NA,(N))
  a[1]=0
  v[1]=0
  phi=0
  for (t in 2:(N)){
    v[t]=phi*v[t-1]+rnorm(1)
    a[t]=v[t]+a[t-1] # This parameter is created to make sure that the random walk is bounded.
  }
  rho=MaxAbsValue
  absmax=max(abs(a))
  genfun=rho*a/absmax #Here the actual random walk function is created.

  return(genfun)
}
```

#1E. Stepwise

```
stepwise<-function(N,MaxAbsValue) #N is the sample size, and MaxAbsValue is the maximum
absolute value of the function.
{genfun=rep(NA,(N)) #Creating the function first with NAs (missing values).
  genfun=c(rep(0,((N)/2)),rep(MaxAbsValue,((N)/2))) #Here the actual stepwise function is
  created.
  return(genfun)
}
```

##2. The function to choose one of the above generating functions##

#This function is built into the next function (to create the data).

```
choose.coef<-function(FUN1,N,MaxAbsValue,MinAbsValue=0){
  FUNchoose=c(invariant=invariant,linear=linear,sine=sine,random.walk=random.walk, stepwise=
  stepwise)
  if(FUN1==2 | FUN1=="linear"){FUNchoose[[FUN1]](N,MaxAbsValue,MinAbsValue)}
  else{FUNchoose[[FUN1]](N,MaxAbsValue)}
}
```

##3. The function to create the data y##

```
creat.y<-function(cho=c("invariant","invariant"),MaxAbsValue=c(1,0.2),MinAbsValue=c(0,0),N=
100,sd_innovation=1) #If you do not specify anything, these are the default settings.
  #With cho you choose one of the five parameter generating functions.
  #With MaxAbsvalue you choose the low or high condition (the maximum absolute value of the
  function).
  #N is the sample size
  #sd_innovation is the amount of error
```

```

{
  if(cho[1]=="linear"){beta_0=choose.coef(cho[1],N=N,MaxAbsValue[1],MinAbsValue[1])} #
  creating the intercept function
  else{beta_0=choose.coef(cho[1],N=N,MaxAbsValue[1])}

  if(cho[2]=="linear"){beta_1=choose.coef(cho[2],N=N,MaxAbsValue[2],MinAbsValue[2])} #
  creating the autoregressive function
  else{beta_1=choose.coef(cho[2],N=N,MaxAbsValue[2])}

  muu=beta_0/(1-beta_1) #deriving the attractor function
  y=rep(NA,(N)) #creating y
  y[1]=rnorm(1,mean=beta_0[1]/(1-beta_1[1]),sd=sd_innovation/(1-(beta_1[1])^2)) #creating the
  first value of y

  for (t in 2:(N)){ #filling in the rest of the values of y
    y[t]=beta_0[t]+beta_1[t]*y[t-1]+rnorm(1,sd=sd_innovation)
  }

  return(list(y=y,beta_0=beta_0,muu=muu,beta_1=beta_1))
}

##4. This is the function for creating the CI for the attractor##
# Since the mean (attractor) is estimated indirectly, we have to estimate the credible
intervals (CI) ourselves.
CI_mean<-function(mod,tt,N){
  # mod is the TV-AR model we will estimate
  # tt is a time variable (this will also be created separately in the actual simulation)
  # N is the sample size
  newd=data.frame(tt=tt,yL=rep(1,N))
  Xp=predict(mod,newd,type="lpmatrix",seWithMean = TRUE)
  k=dim(Xp)[2]/2 #Since we have two smooth functions (beta_0,t and beta_1,t) to estimate,
  the basis functions per smooth function have to be divided by 2
  Xp.beta_0=Xp[,1:k]# basis functions
  Xp.beta_1=Xp[, (k+1):(2*k)]
  Numbrep=10000
  modr<-rmvnorm(Numbrep,coef(mod),mod$Vp+diag(2*k)*10^(-14))#There were really small
  eigenvalues, but they were negative, so we added a very small number to delete such effects.
  res <- rep(0,Numbrep)
  mu.sm=matrix(NA,N,Numbrep)
  for (i in 1:Numbrep){
    beta_0.sm <- Xp.beta_0 %*% modr[i,1:k]
    beta_1.sm <- Xp.beta_1 %*% modr[i,(k+1):(2*k)]
    mu.sm[,i] <- beta_0.sm/(1-beta_1.sm)}
  mu.uncert=apply(mu.sm,1,quantile,c(.025,.5,.975))
  return(list(mu.uncert=mu.uncert))
}

##5. This is the function for lagging the data##
lagmatrix <- function(x,max.lag) embed(c(rep(NA,max.lag), x), max.lag+1)

#####
##### II. Standard time invariant AR #####
#####
#Before you start, please load all libraries and auxiliary functions
#in section "I. General code".
#Equation 1:

```

```

#In order to get a better understanding of equation 1, we will simulate an AR(1) model with
the help of equation 1.
#We first create the simulation function to see how it is done.
AR1<-function(NT,beta_0,beta_1,sd_innovation){
y<-rep(NA,NT) #We create y. NT is the number of time points. NA stands for missing values.
#y is thus first created as a vector of missing values with length NT.

#As every value needs to be regressed on its previous time point, we need to create the
first value of y
#or the first observation. This is done by drawing from a stationary marginal distribution
#(meaning it is not conditioned on the previous time point):
y[1]<-rnorm(1,mean=beta_0[1]/(1-beta_1[1]),sd=sd_innovation/(1-(beta_1[1])^2))

for (t in 2:NT){ #In this for-loop the time points 2 to NT are created.
  #beta_0 is the intercept, beta_1 the autoregressive coefficient.
  #rnorm(1,mean=0,sd=sd_innovation) is the innovation or error drawn from a normal
distribution with mean zero and standard deviation (sd)
  #sd_innovation (standard deviation of the innovation; 1 in this case, thus variance and sd
are the same here).
  y[t]<-beta_0+beta_1*y[t-1]+rnorm(1,mean=0,sd=sd_innovation)
}

return(y)
}

#Using the simulation function to create the figures
par(mfrow=c(1,2))
set.seed(1235)#seed is set so that the exact same figure is made.
y1<-AR1(NT=150,beta_0=3,beta_1=0.5,sd_innovation=1)
y2<-AR1(NT=150,beta_0=6,beta_1=-0.5,sd_innovation=1)

#Plot the AR(1) model with a positive autocorrelation#
plot.ts(y1,col="grey50",ylab="Valence",ylim=c(0,10))
beta_0=3;beta_1=0.5;NT=150;sd_innovation=1
lines(rep(beta_0/(1-beta_1),NT), col="black",lty=1)#This is the equation 2: the mean of the
process.
lines(rep(beta_0,NT), col="black",lty=2) # The intercept
#The lines show that the intercept and mean represent something different.

#Plot the AR(1) model with a negative autocorrelation#
plot.ts(y2,col="grey50",ylab="Valence",ylim=c(0,10))
beta_0=6;beta_1=-0.5;NT=150
lines(rep(beta_0/(1-beta_1),NT), col="black",lty=1)#This is the equation 2: the mean of the
process.
lines(rep(beta_0,NT), col="black",lty=2) #the intercept
#The lines show that the intercept and mean represent something different.

#Estimating an AR(1) model using OLS#
y1L=lagmatrix(y1,1)[,2] #We previously made y1, now we lag the data.

lagmatrix(y1,1)#Compare the original y1 (first column) with the lagged y1L (second column).
#Now we first estimate the AR1 model with the use of an OLS estimation, a simple linear model:
lm(y1~y1L) #beta_1 is y1L
#Then we compare it with the correlation coefficient
cor(y1,y1L,use="complete")#Lag.1 is the correlation coefficient and in this case thus the
autocorrelation.
#We see that they are almost exactly the same.

```

```

#In addition we see that the true value of 0.5 is slightly underestimated.
#We will get a better estimation with more time points.
#Moreover, in an AR model with more time points, beta_1 and the correlation coefficient will
become identical.

#####
##### III. Time-varying AR #####
#####
#Before you start, please load all libraries and auxiliary functions
#in section "I. General code".
set.seed(2242)
par(mfrow=c(1,2))

#Note that the function for simulating a time-varying process (called "creat.y") is very
similar to the AR1 function,
#except now beta_0[t] and beta_1[t] are allowed to vary over time (t).
#for (t in 2:(N)){ #filling in the rest of the values of y
#y[t]=beta_0[t]+beta_1[t]*y[t-1]+rnorm(1,sd=sd_innovation)
#}

#Trend in the data due to a change in intercept
data<-creat.y(cho=c("linear","invariant"),MaxAbsValue=c(5,0.2),MinAbsValue=c(3,0),N=150,
sd_innovation=1)
plot.ts(data$y,col="grey50",ylab="Valence",ylim=c(0,10))
lines(data$beta_0, col="black",lty=2)
lines(data$mu, col="black",lty=1)

#Trend in the data due to a change in the autocorrelation
data<-creat.y(cho=c("invariant","linear"),MaxAbsValue=c(2,0.2),MinAbsValue=c(0,0.65),N=150,
sd_innovation=1)
plot.ts(data$y,col="grey50",ylab="Valence",ylim=c(0,10))
lines(data$beta_0, col="black",lty=2)
lines(data$mu, col="black",lty=1)

#####
##### IV. Inference of the TV-AR model #####
#####
#Before you start, please load all libraries and auxiliary functions
#in section "I. General code".
#### This section is based on chapter 4 of the book:
#Wood, S. N. (2006). Generalized additive models: An introduction with R. Boca Raton, FL:
Chapman and Hall/CRC.

set.seed("3022")# To get the exact same figure, we set a seed.
nT=20# The number of time points
P=20 # This is necessary for creating data that follows a sine wave
t=seq(1:nT) # The sequence of time points
y=sin(2*pi*t/P)+rnorm(nT)*.3 # Creating the data
par(mfrow=c(2,4))
plot(t,y,pch=20,xlab=expression(italic("t")),ylab=expression(italic("y"))) #Plotting the data

k=6# This is the number of basis functions
mod=gam(y~s(t,bs="tp",k=k))# Using the gam function to estimate the beta_0,t function with a
thin plate spline basis

cc=matrix(mod$coefficients)# the estimated alpha coefficients
tt=seq(1,nT,.01)# values for prediction

```

```
newd=data.frame(t=tt)# transform to a data.frame
Xp=predict(mod,newd,type="lpmatrix") # matrix containing the values of the linear predictor

k=dim(Xp)[2]# shows that the prediction matrix Xp has indeed 6 basis functions
for (i in c(1,6,5,2,3,4)){plot(tt,Xp[,i],type="l",ylim=c(-1,1.2),xlab=expression(italic("t")),
,ylab=expression(italic("y")))}
      if ((i>1) & (i<k)){ abline(v=as.numeric(mod$smooth[[1]]$xp[i]),col=
"grey")}}
}
plot(tt,Xp[,1]*cc[1],type="l",ylim=c(-1,1.2),ylab="Weighted basis functions and their sum",
xlab=expression(italic("t"))))
for (i in 1:k){
  lines(tt,Xp[,i]*cc[i],col="grey42")# Plot weighed basis functions
lines(tt,Xp%*%cc,col="black",lwd=4,lty=2) #and the their sum (i.e., the smooth function of
beta_0,t)
points(t,y,pch=20)
```

```
#####
##### V. Guidelines TV-AR model #####
#####
#Before you start, please load all libraries and auxiliary functions
#in section "I. General code".
#We will now illustrate some of the simulation results with 2 examples.
```

```
#####
##### Example A: A time invariant function #####
#####
set.seed(2345)
N=100 #the number of time points in this simulated example
#We will first create the data and true values with the function creat.y
#In this example, we will create data that is time-invariant, and the true values are 1.5
for the intercept
#and 0.5 for the autoregressive parameter.
Simulation1<-creat.y(cho=c("invariant","invariant"),MaxAbsValue=c(1.5,0.5),N=N)
valence_variable<-Simulation1$y #Here we have the data, for example, valence measured daily
```

```
#If you have imported data into R or simulated data (as we do here) you first have to lag
your data.
#We will only use 1 lag, for example, if the variable (e.g., valence) was measured daily we
want to know
#the predictive value this variable has on the next day.
valence_variable_lag1<-lagmatrix(as.vector(valence_variable),1)[,2]#The first argument is
the data and the "k" argument is the number of lags.
plot.ts(valence_variable,xlab="Day",ylab="Valence")#We see here the raw data.
```

```
#Now we will estimate four models to see whether the process is time-varying or
time-invariant.
#The BIC will indicate whether we need the TV-AR model or the AR model.
nrT=length(valence_variable)# This defines the length (or the total number time points).
tt=1:nrT# tt is the time variable for the days the variable was measured.
```

```
#In the following 4 models we see:
#1. A time-varying model where both the intercept (s(tt)) and the autoregressive parameter
s(tt,by=valence_variable_lag1) are time-varying,
#2. A model where only the autoregressive parameter s(tt,by=valence_variable_lag1) is
time-varying and
#the intercept is deleted completely as a time-invariant intercept is already in the model
```

by default

#3. A model where only the intercept is time-varying and the autoregressive parameter
valence_variable_lag1 is time-invariant

#4. A standard AR model with both parameters being time-invariant

```
model_tvar_1=gam(valence_variable ~ s(tt,k=10,bs="tp")+s(tt,by=valence_variable_lag1,k=10,bs="tp"))
```

```
model_tvar_2=gam(valence_variable ~ s(tt,by=valence_variable_lag1,k=10,bs="tp"))
```

```
model_tvar_3=gam(valence_variable ~ s(tt,k=10,bs="tp")+valence_variable_lag1)
```

```
model_tvar_4=gam(valence_variable ~ valence_variable_lag1)
```

#Note that "k" is the number of basis functions and "bs" is the spline basis.

#We have made the default settings explicit, and thus we use 10 basis functions and the thin plate regression splines here.

#You can adjust the number of basis functions and the basis itself separately for the two smooth functions.

#Now we let the BIC select the best fitting model (which is the one having the lowest BIC)

```
which.min(c(BIC(model_tvar_1),BIC(model_tvar_2),BIC(model_tvar_3),BIC(model_tvar_4)))
```

#The BIC indicates that model 4 is the best fitting model and thus a standard time-invariant model is needed.

#A standard time-invariant AR model can be fitted with the "gam()" function as we did above (model4), but also with the "lm()" function or "ar()" or "arima()" functions.

#For example, using "ar()":

```
ar(valence_variable)#this gives us the true autoregressive parameter 0.5
```

#Also note that one can use a full TV-AR model 1. In this case, one also sees that the intercept and the autoregressive parameter are time-invariant.

```
summary(model_tvar_1)
```

#Approximate significance of smooth terms:

```
#  edf Ref.df      F p-value
```

```
#s(tt)                1      1  0.066  0.798
```

```
#s(tt):valence_variable_lag1  2      2 16.347 6.7e-07 ***
```

#In the summary one sees, first of all, that the intercept (s(tt)) is not significant, indicating that it is time-invariant.

#One also sees that the autoregressive function is significant, but this does not tell us whether it is time-invariant or not.

#Thus we have to look at the edf, which is 2, and could indicate either a straight time-invariant line or a linear increase.

#Therefore, we will plot both the intercept and the autoregressive function.

```
plot.gam(model_tvar_1,select=1,rug=FALSE,shift=coef(model_tvar_2)[1],seWithMean = TRUE,ylab="Intercept",xlab="Day",ylim=c(-3,6)) #This gives you the intercept beta0, which is in this case time-invariant (we can easily fit a horizontal line here).
```

```
plot.gam(model_tvar_1,select=2,rug=FALSE,ylim=c(-1,1),seWithMean = TRUE,ylab="Autoregressive parameter ",xlab="Day")
```

#They both look time-invariant.

#Indeed, as we simulated the data, we can now see that the model captured the true underlying (time-invariant)

#process perfectly.

```
plot.gam(model_tvar_1,select=1,rug=FALSE,shift=coef(model_tvar_2)[1],seWithMean = TRUE,ylab="Intercept",xlab="Day",ylim=c(-3,6)) #This gives you the intercept beta0, which is in this case time-invariant (we can easily fit a horizontal line here).
```

```
lines(Simulation1$beta_0,col="red") # Since the data was simulated, we know the true value (the red line), and we see it is well captured by the TV-AR model.
```

```
plot.gam(model_tvar_1,select=2,rug=FALSE,ylim=c(-1,1),seWithMean = TRUE,ylab="Autoregressive parameter ",xlab="Day")
```

```
lines(Simulation1$beta_1,col="red") # Since the data was simulated, we know the true value (the red line), and we see it is well captured by the TV-AR model.
```

```

#####
#% Example B: A time-varying linear function %
#####
set.seed(2345)
N=100 #the number of time points in this simulated example
#We will first create the data and true values with the function creat.y
#In this example, we will create data that is time-varying and the functions
#increase from 0 to 1.5 for the intercept and 0.5 for the autoregressive parameter.
Simulation1<-creat.y(cho=c("linear","linear"),MaxAbsValue=c(1.5,0.5),N=N)
valence_variable<-Simulation1$y #Here we have the data, for example, valence measured daily

#If you have imported data into R or simulated data (as we do here), you first have to lag
your data.
valence_variable_lag1<-lagmatrix(as.vector(valence_variable),1)[,2]#The first argument is
the data and the "k" argument is the number of lags.
plot.ts(valence_variable,xlab="Day",ylab="Valence")#We see here the raw data.

#Now we will estimate four models to see if the process is time-varying or time-invariant.
#The BIC will indicate whether we need the TV-AR model or the AR model.
nrT=length(valence_variable)# This defines the length (or the total number time points).
tt=1:nrT# tt is the time variable for the days the variable was measured.

#In the following 4 models we see:
#1. A time-varying model where both the intercept (s(tt)) and the autoregressive parameter
s(tt,by=valence_variable_lag1) are time-varying,
#2. A model where only the autoregressive parameter s(tt,by=valence_variable_lag1) is
time-varying and
#the intercept is deleted completely as a time-invariant intercept is already in the model
by default
#3. A model where only the intercept is time-varying and the autoregressive parameter
valence_variable_lag1 is time-invariant
#4. A standard AR model with both parameters being time-invariant
model_tvar_1=gam(valence_variable ~ s(tt,k=10,bs="tp")+s(tt,by=valence_variable_lag1,k=10,bs=
"tp"))
model_tvar_2=gam(valence_variable ~ s(tt,by=valence_variable_lag1,k=10,bs="tp"))
model_tvar_3=gam(valence_variable ~ s(tt,k=10,bs="tp")+valence_variable_lag1)
model_tvar_4=gam(valence_variable ~ valence_variable_lag1)
#Note that "k" is the number of basis functions and "bs" is the spline basis.
#We have made the default settings explicit, and thus we use 10 basis functions and the thin
plate regression splines here.
#You can adjust the number of basis functions and the basis itself separately for the two
smooth functions.

#Now we let the BIC select the best fitting model (which is the one having the lowest BIC)
which.min(c(BIC(model_tvar_1),BIC(model_tvar_2),BIC(model_tvar_3),BIC(model_tvar_4)))
#The BIC indicates that model 1 is the best fitting model and thus a time-varying model is
needed.
#In this case, it selected a model in which both the intercept and autoregressive parameter
vary over time.
#Let us take a look at the summary:
summary(model_tvar_1)
#Approximate significance of smooth terms:
#
#           edf Ref.df F      p-value
#s(tt)          1      1 8.537 0.00433 **
#s(tt):valence_variable_lag1  2      2 7.367 0.00105 **
#In the summary, one sees first of all that the intercept (s(tt)) is significant, indicating
that it is time-varying.

```

```

#Interestingly, the edf is only 1, indicating a straight line. However, both the BIC as well
as the significant intercept
#indicate that it is time-varying.
#One also sees that the autoregressive function is significant, but this does not tell us
whether it is time-invariant.
#Thus, we have to look at the edf, which is 2, and could indicate either a straight
time-invariant line or a linear increase.
#Therefore, we will plot both the intercept and autoregressive function.
plot.gam(model_tvar_1,select=1,rug=FALSE,shift=coef(model_tvar_2)[1],seWithMean = TRUE,ylab=
"Intercept",xlab="Day",ylim=c(-3,6)) #This gives you the intercept beta0, which is in this
case time-invariant (we can easily fit a horizontal line here).
plot.gam(model_tvar_1,select=2,rug=FALSE,ylim=c(-1,1),seWithMean = TRUE,ylab="Autoregressive
parameter ",xlab="Day")
#They both look time-varying. This shows that the edf is not an ideal indicator of a
time-varying process.

```

```

#Indeed, as we simulated the data, we can now see that the model captured the true
underlying (time-varying)
#process perfectly.
plot.gam(model_tvar_1,select=1,rug=FALSE,shift=coef(model_tvar_2)[1],seWithMean = TRUE,ylab=
"Intercept",xlab="Day",ylim=c(-3,6)) #This gives you the intercept beta0, which is in this
case time-invariant (we can easily fit a horizontal line here).
lines(Simulation1$beta_0,col="red") # Since the data was simulated, we know the true value
(the red line), and we see it is well captured by the TV-AR model.
plot.gam(model_tvar_1,select=2,rug=FALSE,ylim=c(-1,1),seWithMean = TRUE,ylab="Autoregressive
parameter ",xlab="Day")
lines(Simulation1$beta_1,col="red") # Since the data was simulated, we know the true value
(the red line), and we see it is well captured by the TV-AR model.

```