University of Groningen

Switching Principal Component Analysis for Modeling Means and Covariance Changes Over Time

De Roover, Kim; Timmerman, Marieke E.; Van Diest, Ilse; Onghena, Patrick; Ceulemans, Eva

Published in:
Psychological Methods

DOI:
[10.1037/a0034525](https://doi.org/10.1037/a0034525)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Early version, also known as pre-print

Publication date:
2014

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

De Roover, K., Timmerman, M. E., Van Diest, I., Onghena, P., & Ceulemans, E. (2014). Switching Principal Component Analysis for Modeling Means and Covariance Changes Over Time. *Psychological Methods*, 19(1), 113-132. <https://doi.org/10.1037/a0034525>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Switching principal component analysis for modeling means and covariance changes over
time

Kim De Roover

KU Leuven

Marieke E. Timmerman

University of Groningen

Ilse Van Diest

KU Leuven

Patrick Onghena

KU Leuven

Eva Ceulemans

KU Leuven

Author Notes:

The research reported in this paper was partially supported by the fund for Scientific Research-Flanders (Belgium), Project No. G.0477.09 awarded to Eva Ceulemans, Marieke Timmerman and Patrick Onghena and by the Research Council of KU Leuven (GOA/2010/02). Correspondence concerning this paper should be addressed to Kim De Roover, KU Leuven Faculty of Psychology and Educational Sciences, Methodology of Educational Sciences Research Unit, Andreas Vesaliusstraat 2, 3000 Leuven, Belgium. E-mail: Kim.DeRoover@ppw.kuleuven.be.

The authors would like to thank Peter Kuppens, Francis Tuerlinckx, associate editor Daniel Bauer and four anonymous reviewers for comments on earlier versions of this paper.

Abstract

Many psychological theories predict that cognitions, affect, action tendencies, and other variables change across time in mean level as well as in covariance structure. Often, such changes are rather abrupt, because they are caused by sudden events. To capture such changes, one may repeatedly measure the variables under study for a single individual, and examine whether the resulting multivariate time series contains a number of phases with different means and covariance structures. The latter task is challenging, however. First, in many cases, it is unknown how many phases there are, and when new phases start. Second, often a rather large number of variables is involved, complicating the interpretation of the covariance pattern within each phase. To take up this challenge, we present Switching PCA. Switching PCA detects phases of consecutive observations or time points (in single subject data) with similar means and/or covariation structures, and performs a principal component analysis (PCA) per phase to yield insight into its covariance structure. An algorithm for fitting Switching PCA solutions as well as a model selection procedure are presented and evaluated in a simulation study. Finally, we analyze empirical data on cardio-respiratory recordings.

Keywords: principal component analysis, multivariate time series data, clustering, time contiguity, segmentation, dimensionality

Introduction

In the behavioral sciences, one often formulates hypotheses about the means of a set of variables and their covariance pattern. For instance, in emotion psychology, one investigates the strength and the covariance pattern of different emotion components, such as appraisals, action tendencies, facial expressions, physiological measures (e.g., skin conductance, blood pressure, heart rate, etc.) (Christie & Friedman, 2004; Mauss, Levenson, McCarter, Wilhelm, & Gross, 2005). Another example is found in the field of developmental psychology: when a set of Piagetian tasks (Piaget, 1972) are periodically administered to a child, the means and the covariances of the tasks will yield insight into which tasks require the same cognitive ability (Klausmeier & Sipple, 1982; Lawson & Nordland, 1976).

A complicating factor is that many theories predict that these means and covariance patterns will change across time due to external or internal events, yielding distinct phases of consecutive observations. For the emotion example, it is expected that some of the emotion components take on extreme values (i.e., response patterning; Sander, Grandjean, & Scherer, 2005) and covary stronger (i.e., synchronicity; Butler, 2011) when an emotion-eliciting event occurs, to enable a fast and efficient reaction of the organism (Mauss et al., 2005). In the developmental psychology example, changes in the means are predicted when the child reaches full function on certain tasks. Moreover, a number of theories on Piagetian tasks – for example, the parallel pattern of development by Flavell (1971) or the reciprocal-interaction model by Wohlwill (1973) – suggest that changes in the covariance structure might occur when the child reaches full function on one task (or subset of tasks) while other tasks requiring a similar cognitive ability have yet to be mastered (Klausmeier & Sipple, 1982). Many other examples can be found: Think of how personality-related (Mischel & Shoda, 1995) and attachment-related (Fraley, 2007) cognitions and affect come into play during

social encounters, or how we continuously coordinate the many components of our sensorimotor system (Repp, 2005).

Empirically studying such changes in means and covariance patterns is a challenging task. First, in many cases, it is unknown whether and when events (e.g., emotional trigger, mastery of specific conservation task) occur that mark the start of a new phase, implying that one does not know how many phases there are nor when they begin and end. Second, often a rather large number of variables is involved, complicating the interpretation of the covariance pattern within each phase. Therefore, a modeling technique is needed that on the one hand detects when a change in means and covariance structure takes place, revealing different phases, and on the other hand summarizes and yields insight into the covariance patterns.

Regarding the first aspect – detecting a change in means and covariance structure – a relevant modeling technique is switching regression (Hudson, 1966; Kiefer, 1978; Liu, Wu, & Zidek, 1997). Switching regression segments time series data of a single individual into a number of phases, based on changes in regression coefficients.

For the second aspect – summarizing the covariance of the variables in a specific phase – principal component analysis (PCA; Jolliffe, 2002; Meredith & Millsap, 1985; Pearson, 1901) as well as exploratory factor analysis (EFA; Lawley & Maxwell, 1962) can be used. In PCA, the variables are reduced to a few components that summarize the information in the data. EFA is a stochastic latent variable approach that finds common factors that explain the covariation of the variables under study. Although PCA and EFA are typically used to analyze subjects by variables data, considering time as fixed (the so-called R-

technique; Cattell, 1952), they can also yield insight into time points by variables data of a single subject (P-technique; Cattell, 1952; Jones & Nesselroade, 1990)¹.

Combining the key ideas behind switching regression and PCA, we present Switching PCA. Given multivariate time series data of a single subject, Switching PCA clusters the time points into (a priori unknown) phases of consecutive time points with similar means and/or covariation structures and performs a PCA per phase to yield insight into its covariance structure. We choose to adopt PCA rather than EFA because, unlike PCA, EFA is built on the a priori assumption that the observed variables measure a few latent variables (i.e., variables that cannot be measured directly and have a causal relationship to the observed variables). Indeed, the theoretical status of these latent variables is often problematic, in that some observed variables covary irrespective of the existence of some latent variable (Borsboom, Mellenbergh, & van Heerden, 2003). For instance, if we focus on physiological variables, which is often done in emotion research nowadays, many of them (e.g., inspiratory volume and respiration time) are expected to covary simply because they are biologically connected and not because a latent variable (e.g., ‘respiration’) is underlying the observed scores. Therefore, it is no surprise that for the analysis of physiological data component analysis approaches are more popular than EFA ones (e.g., Boiten, 1993; Friston, Frith, Liddle, & Frackowiak, 1993; Ramadan, Jacobs, Grigorov, & Kochlar, 2006). Moreover, PCA is

¹ PCA, EFA and confirmatory and/or mixture extensions thereof have also been applied to another kind of time series data, namely subjects by time points data (Besse & Ramsay, 1986; Dolan, Schmittman, Lubke, & Neale, 2005; Li, Duncan, Duncan, & Acock, 2001; McArdle & Epstein, 1987; Muthén et al., 2002; Nagin, 1999; Rao, 1958; Rice & Silverman, 1991; Tucker, 1958). These techniques, referred to as T-technique, functional PCA or growth curve models, fall outside the scope of this paper.

attractive from a computational point of view as it implies one algebraic operation only (i.e., eigenvalue decomposition).

The remainder of the paper is structured into five sections. The Model section presents the kind of data for which Switching PCA was developed and describes the Switching PCA model. Additionally, its relations with existing methods are discussed. The Data Analysis section presents the technical details of how to perform a Switching PCA analysis. Next, in the Application section, we analyze an empirical time series data set pertaining to physiological measures. Finally, the Discussion contains some directions for future research.

Model

Data Structure

Switching PCA is intended for the analysis of multivariate time series data from a single subject (i.e., an $N \times J$ data matrix $\mathbf{X}$, where N equals the number of observed time points and J equals the number of variables). The time points do not have to be equally spaced in time, but the spacing should be close enough to have sufficient information per expected phase to reliably study the corresponding means and covariances. Specifically, in the remainder of this paper, we assume that per phase one has at least as many time points as variables, although it is recommended to have many more. Which spacing is close enough to accomplish this, largely depends on the duration of the phenomena under study. For instance, as emotions may fluctuate rather rapidly across time, adequately capturing the covariance structure of its physiological correlates requires multiple time points a minute. On the other hand, for studying more slowly evolving processes like cognitive development, one measurement per week or month might suffice. When some of the variables are sampled more frequently than others, some averaging or imputation (e.g., linear interpolation; Aach & Church, 2001) may

be used to put all scores on the same time scale. With regard to preprocessing, arbitrary differences between the variables in measurement scale are usually eliminated in component analysis by scaling the data such that the variance of each variable equals one.

Model Specification

Looking for changes in mean level and in covariance structure over time, Switching PCA detects a limited number of unknown phases in the multivariate time series. To gain insight into these changes, the data within each phase are modeled by a separate PCA. To construct the Switching PCA model equation, we start from the model equation of a regular PCA model, which reconstructs the observed scores $\mathbf{x}_n$ ($1 \times J$) at the n -th time point as:

$$\mathbf{x}_n = \mathbf{m} + \mathbf{f}_n \mathbf{B}' + \mathbf{e}_n, \quad (1)$$

where $\mathbf{m}$ ($1 \times J$) contains the overall means of the variables, which are usually discarded by centering the data, $\mathbf{B}$ ($J \times Q$) is the loading matrix, which indicates the extent to which each variable is summarized and captured by the respective components, $\mathbf{f}_n$ ($1 \times Q$) contains the component scores of observation n , and $\mathbf{e}_n$ ($1 \times J$) denotes the vector of residuals for the n -th time point.

Adding the segmentation of the time series into phases, the Switching PCA model is given by

$$\mathbf{x}_n = \mathbf{m} + \sum_{c=1}^C u_{nc} \left(\bar{\mathbf{x}}^c + \mathbf{f}_n^c \mathbf{B}^{c'} \right) + \mathbf{e}_n, \quad (2)$$

where C is the total number of phases, u_{nc} is a binary score that indicates whether or not time point n belongs to phase c , $\bar{\mathbf{x}}^c$ ($1 \times J$) contains the phase means of phase c (in terms of deviations from the overall means $\mathbf{m}$), $\mathbf{B}^c$ ($J \times Q^c$) is the loading matrix of phase c , with Q^c

referring to the phase-specific number of components. The parameters $\mathbf{m}$, u_{nc} , $\bar{\mathbf{x}}^c$, $\mathbf{B}^c$, and $\mathbf{f}_n^c$ all have to be estimated in the analysis.

A time contiguity constraint is imposed on the u_{nc} scores, implying that each phase has to consist of consecutive time points². Note that this restriction means that all u_{nc} scores are fixed by determining the $C-1$ phase separations (i.e., the time points where the transitions from one phase to the next take place).

Regarding identification, the following remarks are in order: Without loss of fit, the variances of the component scores are fixed at one for each phase-specific component. This restriction implies that, in case the components of a particular phase are orthogonal, the loadings in this phase can be interpreted as covariances between the variables and the corresponding components. They are covariances rather than correlations, as they depend on both the correlations between the variables and components and the phase-specific standard deviations of the variables (which typically differ from one). To simplify the interpretation of the loadings, they can be rescaled into correlations instead of covariances³. Finally, like in

² We opted for at least J time points per phase to maximize the number of components that can be extracted. The maximum number of components that can be extracted in phase c equals the rank of the data in that phase, i.e., in practice the minimum of N^c , the number of time points in phase c , and J , the number of variables (Jolliffe, 2002).

³ In order to obtain correlations, the loadings can be divided by the phase-specific standard deviations of the variables. This implies that Equation 2 is rewritten as

$$\begin{aligned}\mathbf{x}_n &= \mathbf{m} + \sum_{c=1}^C u_{nc} \left(\bar{\mathbf{x}}^c + \mathbf{f}_n^c \mathbf{B}^{c'} (\mathbf{D}^c)^{-1} \mathbf{D}^c \right) + \mathbf{e}_n \\ &= \mathbf{m} + \sum_{c=1}^C u_{nc} \left(\bar{\mathbf{x}}^c + \mathbf{f}_n^c \tilde{\mathbf{B}}^{c'} \mathbf{D}^c \right) + \mathbf{e}_n\end{aligned},$$

PCA, the components of Switching PCA have rotational freedom per phase. Specifically, to enhance the interpretation of the component structures even further, the original loading matrices $\mathbf{B}^c$ or their rescaled versions $\tilde{\mathbf{B}}^c$ can be multiplied by an orthogonal or oblique rotation matrix without loss of fit, provided that the transformation is compensated for in the component scores.

Finally, we want to emphasize that data containing strong phase changes, can always be fitted well by a regular PCA, provided that a sufficiently high number of components is used. However, such an analysis will hardly shed light on these phase changes and will yield components that mix up two sources of (co)variance: i.e., between-phase differences in means and within-phase (co)variances (see Illustrative application).

Related Methods

First of all, some extensions of PCA and EFA exist that allow for differences in component or factor structure by clustering the rows of a data matrix and simultaneously fitting a PCA or EFA model to each cluster (i.e., for PCA: subspace K-means, Timmerman, Ceulemans, De Roover, & Van Leeuwen, 2013; for EFA: mixtures of factor analyzers, McLachlan & Peel, 2000; Yung, 1997). Yet, those techniques are unsuitable for the research questions at hand, because the resulting clusters may contain any combination of consecutive and non-consecutive time points and therefore cannot be interpreted as phases.

In the remainder of this section, we discuss two types of techniques that can be used to segment time series data of a single subject into different phases, based on means and/or

where $\mathbf{D}^c$ is a $J \times J$ diagonal matrix that contains the standard deviations of the variables in phase c , and $\tilde{\mathbf{B}}^c$ the rescaled or ‘standardized’ loadings.

covariances: change-point detection and hidden markov modeling (HMM). Change-point detection methods perform statistical tests to detect if and when a mean or covariance shift takes place in the time series (e.g., Aue, Hörmann, Horváth, & Reimherr, 2009; Lavielle, & Teyssière, 2006; Sullivan & Woodall, 2000). For instance, some of these methods test, for each time point n , whether the mean and/or covariance matrix up until this time point (i.e., for time points 1 to n) is different from that of the time points after n (time points $n+1$ to N). If a change is detected, the time series is split up in two parts according to the most likely location of the change-point and the test is repeated for both parts etc., until no further change-points are found. An important difference with Switching PCA is thus that it looks for multiple change-points in a sequential manner. Another major difference is that these methods do not explicitly fit a model to the data and as such provide no insight into which covariances change.

A HMM models the data as a sequence of latent states, according to the scores on the variables (Baum & Petrie, 1966; Raftery, 1985). Different variants have been proposed, including latent transition analysis (LTA; Graham, Collins, Wugalter, Chung, & Hansen, 1991; Velicer, Martin, & Collins, 1996) and the regime-switching state-space model (Kim & Nelson, 1999; Lodewyckx, Tuerlinckx, Kuppens, Allen, & Sheeber, 2011). While Switching PCA is a deterministic model, HMMs are stochastic. Focusing on the research questions presented in the Introduction – how to model changes in mean level and covariance structure across time – the key differences pertain to the use of dimension reduction, and the nature of the states, as we will now discuss subsequently.

Typically, a HMM is intended for modeling a limited number of variables (say 5) only and, thus, performs no dimension reduction per state. Moreover, the HMM state sequence is, in most cases, based on the means only. Specifically, every state is characterized by a vector of means and a covariance matrix, while the state-specific covariance matrices are restricted

to confine the number of parameters. Typical restrictions of the state-specific covariance matrices are diagonality (i.e., local independence) or equality across (subgroups of) states (see, e.g., Gales, 1999). Another option that is particularly interesting for the present paper is imposing an exploratory (Rosti & Gales, 2004) or confirmatory (Boker, Neale, & Rausch, 2004; Schmittmann, Dolan, van der Maas, & Neale, 2005) factor model on each state or subgroup of states. An important difference between the ‘factor analyzed HMM’ of Rosti and Gales (2004) and Switching PCA is that the former is developed for making the HMM estimation possible and not for obtaining a state sequence that models shifts in means as well as in the underlying factor model. Moreover, a state-specific common factor model builds on the assumption of underlying latent variables we meant to avoid (see Introduction). Additionally, as HMM is mostly applied to only a few variables, often for multiple subjects simultaneously, and is mainly used to model differences in means and not in covariances, it is unclear how many time points would be needed per state for the associated estimation procedures to converge, especially in case the covariance matrices are nondiagonal and different across the states of a single subject.

The second important difference between Switching PCA and HMM pertains to the nature of the states. Both approaches assume a limited number of states, as the Switching PCA phases are modeled separately and therefore may be considered distinct states. In a HMM, one models the probability of going from one state to the other at consecutive time points. The transition probabilities (also called the ‘Markov dependency’) are typically time-invariant⁴ and typically not restricted. The latter implies that states generally reoccur across time (i.e., people can go back and forth between the latent states). In contrast, in a Switching

⁴ In HMM, time-dependent transition probabilities can be covered when modeled as a function of a known covariate (Diebold, Lee, & Weinbach, 1994; Filardo, 1998).

PCA, the sequence of states is severely constrained: the states are consecutively ordered, implying that one cannot go back to a previous state, and one cannot skip a certain state. If one were to define the transition probabilities of a Switching PCA model, only the transition probabilities towards the next state are non-zero, while all other transition probabilities are zero. The non-zero transition probabilities are time dependent, but their time dependency is not explicitly modeled. The time dependency occurs because of the time contiguity constraint and the restriction on the minimal number of time points within a phase: That is, if time points $n-1$ and $n+1$ clearly belong to a particular phase, time point n will be assigned to that phase as well.

Data Analysis

In this section, we present the Switching PCA objective function, an algorithm for performing a Switching PCA analysis, and a stepwise procedure for selecting the most appropriate number of phases and components. This section is rather technical and may be skipped by readers who are mainly interested in the key idea of the method and its usefulness for substantive research.

Objective Function

When conducting component analysis, one typically looks for the model estimates that, given a prespecified number of components, minimize the sum of the squared residuals

$$SSE = \sum_{n=1}^N \|\mathbf{e}_n\|^2, \quad (3)$$

where $\| \|^2$ indicates the computation of the sum-of-squares of the residuals of time point n .

However, in our case, where the number of components Q^c varies over the phases, this may lead to the majority of the time points being erroneously assigned to phase(s) with a higher Q^c -value, since using more components yields more modeling freedom (i.e., more component scores per time point), which often reduces the SSE -value (see De Roover, Ceulemans, Timmerman, Nezlek, & Onghena, 2013). Specifically, this phenomenon may occur when the phases have a relatively similar component structure. Therefore, building on the work of De Roover et al. (2013), we propose to penalize the Switching PCA SSE for the number of component scores to be estimated, yielding the following objective function:

$$L = NJ \log(SSE) + 2 \sum_{c=1}^C N^c Q^c, \quad (4)$$

where N^c is the number of time points in the c -th phase. A technical derivation of this objective function is given in Appendix A. The key idea is that a time point is only added to a phase with a higher Q^c -value when the resulting increase in fit outweighs the increase in the number of component scores to be estimated. Herewith, we make the assumption that the residuals are independently, identically and normally distributed (see Appendix A for more details).

Often, it is interesting to know how much of the variance in a given data set is accounted for by a certain Switching PCA model. The percentage of variance accounted for (VAF%) can be obtained from the SSE as follows:

$$\text{VAF\%} = \frac{SST - SSE}{SST} \times 100 \quad (5)$$

where $SST = \sum_{n=1}^N \|\mathbf{x}_n - \mathbf{m}\|^2$.

Algorithm

To minimize the objective function L , we developed an alternating least squares (ALS) algorithm, which has been implemented in Matlab R2011b and can be obtained freely from the first author. Also, stand-alone software for applying Switching PCA will be developed in the future (similar to the component analysis packages provided by De Roover, Ceulemans & Timmerman, 2012, and Wilderjans, Ceulemans, Kiers & Meers, 2009). Conceptually speaking, the algorithm consists of the following steps (technical details of the algorithm can be found in Appendix B):

- 1) Randomly segment the time series into C phases: The algorithm selects at random $C-1$ phase transition points among the N time points, provided that the resulting phases contain at least J time points. This initializes the u_{nc} scores in Equation 2.
- 2) Estimate the phase-specific means and PCA models of the initial phases.
- 3) Find the best possible transition points: The algorithm alternately optimizes the timing of the transitions between the subsequent phases. Specifically, the algorithm searches for the optimal position for each phase transition, while keeping the other phase transitions fixed (see Figure 1) and subject to the restriction that each phase contains at least J time points. With each relocation of a phase transition, the phase-specific means and PCA models are updated. This procedure (i.e., step 3) is repeated until the phase transitions no longer change.

To reduce the probability of retaining a solution that is a local minimum, the algorithm uses a multistart procedure, i.e., with different random segmentations of the time series (step 1). To improve the efficiency of this multistart procedure, it first generates a large number (e.g., 125) of random segmentations of the time series which are evaluated in terms of their fit to the data. Next, the 20% best fitting random segmentations (e.g., 25 out of 125) are retained and

used as initial segmentation (i.e., in step 1). The best solution out of the different runs is retained as the final solution. The user has to specify the desired number of random starts.

[Insert Figure 1 about here]

Model Selection

When performing Switching PCA, one has to specify the number of phases and the numbers of phase-specific components. Yet, the most appropriate number of phases C , C^{best} , is often unknown, as well as the most appropriate number of components Q^c for each phase, $Q^{c,\text{best}}$. Therefore, one usually fits a number of Switching PCA models with different C - and Q^c -values and retains the model with the best balance of fit and number of phases and components. The number of models under consideration rapidly grows very large, however. For example, if C^{max} and Q^{max} equal six (both values should be larger than can be reasonably expected for the data set in question), 923 different models are obtained (see Appendix C). Therefore, we propose a more time-efficient, stepwise model selection strategy (similar to the one proposed in De Roover et al., 2013). Specifically, this procedure, that is integrated in the Switching PCA software, consists of the following steps (technical details can be found in Appendix C):

- 1) **Initial screening:** Select the best model from the subset of models that impose an equal number of components Q^c in all phases (only these models have to be fitted), using the automated scree test proposed by Ceulemans and Kiers (2006; see also Wilderjans, Ceulemans & Meers, 2013). An attractive feature of this test is that it also allows for a ranking of promising models (Bulteel, Wilderjans, Tuerlinckx & Ceulemans, 2013), which is useful as the scree test is, as all tests, susceptible to sampling fluctuations.

Therefore, it often makes sense to also consider the second best solution and make the final model selection decision based on interpretability and stability.

- 2) Re-evaluate the best number of components $Q^{c,best}$ for each phase separately: Using the phase segmentation of the time series that was found in step 1, compute PCA solutions with one up to $Q^{\max}$ components per phase and determine the best number of components $Q^{c,best}$ for each phase separately using the above mentioned automated scree test.
- 3) Estimate the Switching PCA model with the selected numbers of phases and components: Perform a Switching PCA analysis with 25 random starts and a 26th ‘rational’ start using the phase segmentation obtained in step 1, and retain the best solution.
- 4) Check convergence: Repeat step 2 using the phase segmentation that results from step 3, to evaluate whether changes in the transitions points affect the selected Q^c s. If this is the case, repeat step 3 and 4.

This procedure will be illustrated in the Illustrative Application section.

Simulation Studies

Problem

To evaluate the performance of the Switching PCA algorithm and model selection procedure, we performed two simulation studies. In both studies a large number of data sets are generated according to a known Switching PCA model, manipulating several characteristics which were selected on the basis of previous research (without claiming to be exhaustive): number of time points and variables, number of phases and within-phase components, similarity of the component loadings of subsequent phases, phase lengths,

between-phase variance (i.e., the amount of variance accounted for by the phase-specific means, which is computed as $\frac{1}{N} \sum_{c=1}^C N^c (\bar{x}_j^c)^2$ for variable j , where $\bar{x}_j^c$ is the deviation of the phase-specific mean from the overall mean m_j). Subsequently, these data are perturbed with different amounts of noise. In Simulation Study 1, we assessed the sensitivity of the algorithm to local minima as well as the goodness-of-recovery of the phases, the phase means, and the phase-specific loading matrices. These aspects are evaluated under the assumption that the correct numbers of phases C and phase-specific components Q^c are known. Moreover, we also assess whether the proposed stepwise model selection procedure succeeds in indicating the correct C - and Q^c -values. In Simulation Study 2, we examined whether and how model estimation and selection is affected by the presence of a first-order autocorrelation and/or cross-lagged correlation among the components of each phase.

Simulation Study 1

Design and procedure.

Eight factors were crossed in a complete factorial design:

1. *number of time points* N at 3 levels: 180, 240, 300;
2. *number of variables* J at 2 levels: 12, 24;
3. *number of phases* C at 3 levels: 2, 3, 4;
4. *phase lengths* at 3 levels: J time points are assigned to each phase and the remaining ones are distributed (respecting the time contiguity restriction) as follows (see Milligan, Soon, & Sokol, 1983): equal (equally distributed over all phases), unequal with minority (10% added to a randomly selected phase, remaining time points

- equally divided over the other phases), unequal with majority (60% added to a randomly chosen phase, remaining time points equally divided over the other phases);
5. *numbers of within-phase components Q^c* at 4 levels: [2 2 (2) (2)], [4 4 (4) (4)], [2 1 (2) (1)], [4 2 (4) (2)], where the elements of the vector indicate the Q^c -values and the length of the vector depends on the number of phases C , e.g., '[4 2 (4) (2)]' indicates that Q^1 equals 4, Q^2 equals 2, and, in case of a third and fourth cluster (factor 3), Q^3 equals 4, and Q^4 equals 2;
 6. *similarity of component loadings of subsequent phases* at 2 levels: random phase loadings with low similarity, simple structure phase loadings with medium similarity;
 7. *ratio b of the between-phase variance to the total structural variance* at 3 levels: .00, .25, .50;
 8. *error variance e* at 2 levels: .20, .40.

For each cell of the factorial design, 20 data matrices $\mathbf{X}$ were generated. Specifically, the number of consecutive time points within each phase was determined according to factors 1, 3 and 4. The $N^c \times Q^c$ component score matrices $\mathbf{F}^c$ (containing the component scores $\mathbf{f}_n^c$ of all N^c time points in phase c) were generated by randomly sampling entries from a standard normal distribution. The random J (factor 2) $\times Q^c$ (factor 5) phase loading matrices $\mathbf{B}^c$ (level 1 of factor 6) were created by uniformly sampling the loadings between -1 and 1 and rescaling the rows such that their sum-of-squares equals 1. In the simple structure conditions (level 2 of factor 6), binary loadings were used such that (1) all variables have a zero loading on all but one component; (2) in case of an identical Q^c for all phases, the sets of variables that constitute the respective components differ from those of the preceding phase with regard to one variable only (see example in Figure 2A); and (3) in case of varying Q^c , the sets of variables that constitute the components are assembled in such a way that the components of

the phases with lower Q^c can be obtained by a pairwise merge of two components of the preceding phase (see example in Figure 2B). These simple structure loading matrices were selected to achieve moderately similar phase-specific structures⁵. To avoid the presence of one very dominant component and to make sure that all components are sufficiently salient (see, e.g., Ceulemans, Timmerman, & Kiers, 2011; Timmerman & Kiers, 2000), we imposed the following two restrictions: (a) each component should have an eigenvalue larger than $.02*J$ and thus explain more than 2% of the structural within-phase variance; (b) none of the components should have an eigenvalue larger than $.60*J$ and thus explain more than 60% of the structural within-phase variance. These loading matrices are rowwise rescaled such that the sum-of-squares of each row equals $(1-b)*(1-e)$ (factors 7 and 8). Next, to add between-phase variance to the data, phase means are randomly sampled from a uniform distribution and rescaled such that, in the final data matrix $\mathbf{X}$, the expected proportion of variance accounted for by the phase means $\bar{\mathbf{x}}^c$ is $(b)*(1-e)$ over all variables. When b equals zero, the phases have identical means and differ only in terms of their covariance structures. The error matrix $\mathbf{E}$ is randomly sampled from a standardnormal distribution and rescaled such that the

⁵ To quantify the degree of similarity among the phase loading matrices, a mean RV-coefficient “ RV_{mean} ” was calculated for each data set. The RV-coefficient (Robert & Escoufier, 1976) is a rotation-independent correlation between two matrices, which allows for the number of columns to differ between the matrices and which takes values between 0 and 1. For each data set, the RV-coefficient is computed for each pair of subsequent true phase loading matrices and then averaged over the pairs to obtain RV_{mean} . On average, RV_{mean} amounts to .14 ($SD = .09$) and .70 ($SD = .08$) for the random and simple structure loadings respectively, which indicates that the random loadings matrices are very different among subsequent phases while the simple structure loading matrices are moderately similar as intended by the manipulation.

expected error variance e equals .20 or .40 (factor 8). Finally, the $\mathbf{X}$ matrices are constructed using Equation 2 and rescaled to have a variance of one per variable.

[Insert Figure 2 about here]

In total, 51,840 simulated data matrices were generated, as we created 20 replicates per cell of the design. Each data matrix $\mathbf{X}$ is analyzed with the Switching PCA algorithm, using 25 random starts (i.e., 20% best out of 125 evaluated random starts). Using the correct number of phases and components, the average computation time for one data set amounted to 13 seconds when conducted on an Intel® Core™ i7-3770K processor of a personal computer, with a clock frequency of 3.4 to 3.9 GHz and a RAM speed of 1600 MHz.

To reduce the computational burden, the evaluation of the model selection procedure is confined to the first five replications in each cell. To each of these 12,960 data sets, we applied the stepwise model selection procedure, setting the $C^{\max}$ - and $Q^{\max}$ -values to six. On average, this took about 9 minutes of computation time per data set on the same processor.

Results.

Model estimation.

Sensitivity to local minima.

The Switching PCA algorithm may end in a local minimum in case the number of phases is larger than two, which implies that another solution exists that has an even smaller value on the objective function. To evaluate the sensitivity of the algorithm to local minima, we compared the loss function value (Equation 4) of the retained solution to that of the solution that is obtained when starting the algorithm from the correct segmentation. If the

former loss function value is larger than the latter, the multistart solution is a local minimum for sure. This way, we established that a local minimum was found for 2,221 out of the 51,840 data sets (i.e., 4.28%). The majority of the local minima (2,152) are obtained in the conditions with four underlying phases and with unequal phase lengths with majority, probably because it is more difficult to recover the correct segmentation when more phase separations need to be estimated, especially when the differences in phase length are more extreme. In those more difficult cases, using more random starts (e.g., 50 or 100) may be a solution.

Goodness-of-recovery of the structural phases.

To examine the recovery of the underlying phases, the *Adjusted Rand Index* (*ARI*, Hubert & Arabie, 1985) is calculated between the true and estimated segmentations. The *ARI* equals one if both segmentations are identical and zero when their agreement is at chance level. The overall mean *ARI* is .98 ($SD = 0.08$), which indicates an excellent recovery of the phases. The vast majority (i.e., 10,073) of the 11,035 data sets for which the phase recovery is not perfect (i.e., $ARI < 1.00$) belong to the conditions with four phases and/or with 0% between-phase structural variance. The mean *ARI* values (and 95% confidence intervals) are presented in Figure 3A as a function of the corresponding manipulated factors.

[Insert Figure 3 about here]

To investigate whether the recovery of the phases is related to how much they overlap in terms of means and covariances, we calculated the Morisita measure for overlap between two multivariate normal distributions (a value of 1 indicates perfect overlap; Lu, Smith, & Good, 1989) for each pair of subsequent phases (i.e., comparing the simulated data of subsequent phases) and averaged these values per data set. We expect most overlap when

phases hardly differ in terms of their means (i.e., 0% between-phase structural variance) and have similar component structures (i.e., simple structure loadings with medium similarity). The obtained mean Morisita values varied between .00 and .64, with an average of .05 ($SD = .10$). Although this seems to suggest that, in general, the overlap is pretty low, this conclusion is incorrect in the sense that the Morisita overlap measure is based on the complete $J \times J$ covariance matrices of subsequent phases and therefore is negatively correlated with the number of variables. For example, the average Morisita value is equal to .30 and .08 for 12 and 24 variables, respectively, in case of 0% between-phase structural variance and simple structure loadings. In other words, the Morisita values should be interpreted relatively (i.e., compared to other data sets with the same number of variables) rather than absolutely. The Spearman rank correlation between the ARI and the mean Morisita value is $-.46$ ($p < .001$) and $-.43$ ($p < .001$) for the data sets with 12 and 24 variables respectively, suggesting that the recovery of the phase separations is indeed related to the amount of phase overlap.

Goodness-of-recovery of the phase means.

The recovery of the phase means is evaluated using the mean absolute difference between the true and estimated phase means:

$$MAD_{means} = \frac{\sum_{c=1}^C \sum_{j=1}^J |\bar{x}_j^{cT} - \bar{x}_j^{cM}|}{\sum_{c=1}^C \sum_{j=1}^J |\bar{x}_j^{cT}|} \quad (6)$$

where $\bar{x}_j^{cT}$ and $\bar{x}_j^{cM}$ denote the true and estimated mean, respectively, for variable j within phase c . The true means of a particular phase are computed as the means of the error perturbed data within that phase, according to the true phase separations. MAD_{means} is zero when the phase means are perfectly recovered, while higher values indicate worse recovery.

Thus, the overall average MAD_{means} of 0.07 ($SD = .23$) indicates a very good recovery of the phase means. The phase means are recovered perfectly for 40,805 or 79% of the simulated data sets. As we expect the recovery of the phase means to depend strongly on the recovery of the phase segmentation, we computed the Pearson correlation among the ARI and the MAD_{means} over all simulated data sets. This correlation amounts to $-.89$, implying that the phase means are recovered worse when the estimated phase transitions diverge more from the true ones. Therefore, it is no surprise that MAD_{means} values larger than zero mainly (i.e., 10,073 of the 11,035 cases) occur in the conditions with four phases and/or with 0% between-phase structural variance.

Goodness-of-recovery of the phase-specific loading matrices.

The recovery of the phase-specific loading matrices was assessed by computing congruence coefficients φ (Tucker, 1951) between the true and (rescaled) estimated component loadings and averaging these coefficients across components and phases:

$$GOPL = \frac{\sum_{c=1}^C \sum_{q=1}^{Q^c} \varphi(\mathbf{B}_q^{cT}, \tilde{\mathbf{B}}_q^{cM})}{\sum_{c=1}^C Q^c} \quad (7)$$

with $\mathbf{B}_q^{cT}$ and $\tilde{\mathbf{B}}_q^{cM}$ indicating the true and estimated q -th component for phase c , respectively.

Moreover, to deal with the rotational freedom of the Switching PCA model, each $\tilde{\mathbf{B}}_q^{cM}$ matrix was orthogonally procrustes rotated towards the true loading matrix of the corresponding phase. The goodness-of-phase-loading-recovery or $GOPL$ statistic takes values between 0 (no recovery at all) and 1 (perfect recovery).

On average, the mean *GOPL* amounts to .98 ($SD = .05$), which indicates an excellent recovery of the loadings. Indeed, according to Lorenzo-Seva and ten Berge (2006), congruence coefficients larger than .95 indicate that components are essentially equal and this is the case for no less than 48,075 (93%) of the 51,840 simulated data sets. The majority (i.e., 3,157) of the 3,765 cases with a *GOPL* smaller than .95, are situated in the conditions with 50% between-phase structural variance and/or 40% error variance. This implies that in these conditions the within-phase structural variance is easily masked by the error variance. The mean *GOPL* values (and 95% confidence intervals) in function of the corresponding manipulated factors are given in Figure 3B.

Model selection.

The stepwise model selection procedure selects the correct Switching PCA model – i.e., the correct number of phases C as well as the correct number of components Q^c for each phase c – for 69% of the cases (8,983 out of 12,960). For 2,212 (56%) of the 3,977 errors the number of phases of the selected model is incorrect⁶. The majority of these mistakes against the number of phases (i.e., 2,200 out of 2,212) happen when the data contain more than two phases, the phases have a more similar component structure (simple structure loadings) and/or the phases differ only in covariances and not in means (0% between-phase structural variance). If we consider the ranking of the different models that results from the model

⁶ The correctness of the retained number of phases is somewhat correlated to the mean Morisita value (Lu, Smith, & Good, 1989) for overlap between subsequent phases, i.e., the Spearman correlation amounts to $-.32$ ($p < .001$), both for the data sets with 12 variables and with 24 variables.

selection procedure, it appears that for 1,185 of the data sets for which an incorrect number of phases is indicated as best by the scree test, the second best number of phases is the correct one.

For 1,765 data sets, the selected model has a correct number of phases but an incorrect number of components within at least one phase. Most (1,763) of these mistakes are made when four components are underlying one or more of the phases, when only 12 variables are available, and/or when the data contain 40% error variance. For 699 of these data sets, the second best Q^c was the correct one for the phase(s) with an incorrect best Q^c .

Simulation Study 2

Design and procedure.

In Simulation Study 2, we used the same eight factors and levels as in the first study and included an additional factor pertaining to *serial dependencies*. This factor was manipulated at 3 levels: $[\rho_{ac}, \rho_{cc}]$ equal to $[\.70, .00]$, $[\.00, .70]$, $[\.70, .70]$. ρ_{ac} denotes the expected value of the lag-one autocorrelation among the component scores of each component within each phase. ρ_{cc} indicates the expected correlation among the lag-one scores on component 1 and the scores on component 2 (lag-one cross-correlation), within each phase with two or more components.

For each cell of this design, five data matrices $\mathbf{X}$ were generated, yielding 38,880 data matrices. The component score matrices $\mathbf{F}^c$ are sampled from a multivariate normal distribution, with the means equal to zero, the variances equal to one, and the covariance between the first and second components equal to ρ_{cc} , and zero covariances among all other components in conditions with more than two components. Subsequently, to impose the lag-

one cross-correlation, the scores on component 1 are shifted such that the n -th score becomes the $(n-1)$ th score (note that the first score becomes the last score). Next, to each column of $\mathbf{F}^c$ a first-order recursive filter is applied to induce the lag-one autocorrelations (Hamilton, 1994). Finally, the resulting component scores are rescaled to an expected variance of one by multiplying them by $\sqrt{1 - \rho_{ac}^2}$. Note that the filtering procedure does not affect the expected lag-one cross-correlation among components 1 and 2. All other matrices were generated as described for Simulation study 1. The resulting data sets were analyzed in the same way.

Results.

With respect to model estimation, for only 1,712 (4.40%) of the data sets, the algorithm ended in a local minimum for sure. The average *ARI* equals .97 ($SD = 0.08$), the average *MAD_{means}* 0.07 ($SD = 0.24$) and the average *GOPL* .97 ($SD = 0.05$). The conditions in which the model estimation performance is somewhat lower are the same as in Simulation Study 1. Specifically, 8,477 of the 9,369 data sets for which *ARI* is lower than 1 (and *MAD_{means}* larger than zero) have four underlying phases and/or 0% between-phase structural variance, and 3,415 of the 3,880 data sets with a *GOPL* smaller than .95 are situated in the conditions with 50% between-phase structural variance and/or 40% error variance. We conclude that the performance of the Switching PCA algorithm is not affected by lag-one auto- and cross-correlations among the components.

With respect to model selection, the correct Switching PCA model is selected for 57% of the simulated data sets, which is about 12% less than in Simulation study 1. When looking at the three levels of serial dependencies separately, we find percentages of 55%, 70% and 46%, respectively. This implies that model selection is mainly affected by the lag-one autocorrelations of .70 in levels 1 and 3 of factor 9 (especially in case of the simple structure

loadings). The number of phases is correctly assessed for 86% of these two types of data sets (i.e., for levels 1 and 3 of factor 9 combined) in case of random loadings and for 67% of these data sets in case of simple structure loadings. Thus, we conclude that the selection of the number of phases C is more challenging when rather similar covariance structures per phase are combined with the autocorrelations. Taking the second best number of phases into account may mend the C -selection, as this is the correct one for 9% of the random loading cases and 18% of the simple structure cases with an autocorrelation of .70. Also, the selection of the number of components per phase (Q^c) has become more difficult, as indicated by the fact that, for levels 1 and 3 of the serial dependencies factor, the number of components is selected incorrectly in at least one phase for 66% of the data sets for which the number of phases is selected correctly. For 14% of these data sets, the correct model complexity may be obtained by also considering the second best number of components for each phase.

Conclusion

Based on the above results, we conclude that in the considered range of simulated conditions (which of course are not exhaustive) and given the correct numbers of phases and components, the Switching PCA algorithm only seldomly yields a local minimum for sure when using a multistart procedure with 25 random starts. However, we advice to use more random starts (say, 50 or 100) when more phases are estimated, which may be unequal in length. In future research, it may be worthwhile to investigate the occurrence of local minima when model complexity is over- or underestimated.

Regarding parameter estimates, the algorithm seems to recover the phase separations, the phase means, and the within-phase loading matrices very well. The recovery of the phase separations may be somewhat influenced by the restriction that each phase contains at least J

time points (see Footnote 2), however. This restriction may, on the one hand, be too stringent in the sense that shorter phases cannot be captured and, on the other hand, it ensures the identifiability of the principal components but still is quite minimalistic to ensure a solid estimation of the within-phase model. Therefore, it would be useful to evaluate the influence of the phase length restriction in future research.

Regarding model selection, Simulation Study 1 indicated that the stepwise selection procedure selects the correct model in the majority of the cases despite the difficulty of the task, i.e., selecting the correct number of phases as well as the correct number of components for each phase. Nonetheless, the model selection procedure is far from flawless since selection mistakes are made for 31% of the simulated data sets. Therefore, when analysing real data, it is important to realize that the procedure is just a heuristic and to carefully consider the rank orderings that result from it, to select a few promising candidate solutions for further inspection. Indeed, for many of the model selection mistakes, the second best number of phases or components was the correct one in the simulations. Finally, Simulation Study 2 reveals that the model selection performance is somewhat hampered by the presence of autocorrelations.

Illustrative Application

To illustrate our new method, we collected and analysed cardio-respiratory recordings from a young, healthy man who underwent three, experimentally induced states: resting baseline (11 min.), inhalation of 7.5% CO₂-enriched air (5 min.), and recovery (6 min.) This procedure is regularly used in research on panic or respiratory symptom perception (e.g., Pappens, de Peuter, Vansteenwegen, Van den Bergh, & Van Diest, 2012; Fannes et al., 2008), and was approved by the local ethical committee. The participant wore a face mask that was

connected to a non-rebreathing valve allowing for the separation of inspiratory and expiratory air. A wide tube (60 cm long, with a diameter of 5 cm) was connected to the inspiratory side of the non-rebreathing valve and led through the wall to the experimenter room. At the experimenter side, a three-way valve was mounted on this inspiratory tube that allowed the experimenter to switch between regular roomair and the CO₂ mixture.

After providing his informed consent, the participant breathed roomair during a resting baseline. After 11 minutes, the valve was switched to the CO₂ mixture for 5 minutes. Finally, it was switched back to roomair for another 6 minutes. The participant was unaware of both switches. Respiration and heart rate were measured continuously by means of respiratory inductive plethysmography (RIP) and electrocardiography (ECG), using the LifeShirt System[®] (Vivometrics Inc., Ventura, CA). Two inductance belts at the level of the ribcage and the abdomen, sewn into a LifeShirt garment, were connected to a digital processing unit, including a data storage card. Vivologic software (Vivometrics Inc., Ventura, CA) was used to extract nine physiological variables (see Table 1) for each breath. The resulting data set contains 134 time points for the baseline, 66 time points for the CO₂ inhalation, and 90 time points for the recovery. As the scale differences of the physiological variables are not interesting from a substantive point of view, we rescaled each variable to a variance of one.

[Insert Table 1 about here]

As indirect cardiorespiratory recordings are very prone to measurement error, we visually inspected the recorded signals for movement artefacts. We found that in some breaths a high inspiratory volume co-occurred with a low expiratory volume, or vice versa, which may indicate an artefact caused by movements, coughs or hiccups. Therefore, we removed 51 time points for which expiratory volume differed more than one standard deviation from the inspiratory volume, leaving 115 baseline time points, 40 CO₂ inhalation time points and 84

recovery time points (for which each variable was again rescaled to a variance of one). These data are plotted in Figure 4.

[Insert Figure 4 about here]

To start off the model selection procedure, Switching PCA analyses were performed with one to six phases and zero (i.e., only modeling the means) to four components per phase. In Figure 5 the VAF% values of the different Switching PCA solutions are plotted against Q for each value of C . When comparing the scree lines for the different numbers of phases, we see that going from one to two phases and from two to three phases gives a large increase in the VAF% (especially for the lower Q -values), while adding more phases leads to substantially smaller improvements in fit. Therefore, three phases seem indicated for this data set. Note that the gap between the scree lines for two and three phases becomes considerably smaller for the larger Q -values, because in those cases less phases are needed to reach a good fit (i.e., interdependence in model selection; Bauer & Curran, 2004; Molenaar & von Eye, 1994). For instance, the fit of the regular four-component PCA model (i.e., with one phase) is almost as good as that of the model with three phases with three components each. As mentioned in the Model Specification section, a drawback of the former model is that its components capture both the between-phase differences in means and the different within-phase covariance structures, making it difficult to gain insight into both types of differences. This is illustrated by Figure 6 which plots the scores on the first three principal components of the regular PCA solution: The first component mainly captures an increase in some physiological parameters during the CO₂ inhalation as well as a decrease afterwards, whereas the second component unravels a mix of more subtle manipulation-related differences in level as well as within-phase covariances, and the third component seems to address within-phase covariances only. Therefore, we prefer to select a model with three phases.

[Insert Figures 5 and 6 about here]

The next step of the model selection procedure focuses on the scree line for three phases. The increase in fit seems to level off after three components. These observations are supported by the scree ratios in Table 2, which also point to three phases and three components per phase (see Appendix C for more information on how to compute and use the scree ratio's).

[Insert Table 2 about here]

Subsequently, the best number of components is re-evaluated for each phase separately. The VAF% of the phase-specific PCA models are plotted against the number of components Q^c in Figure 7. In this figure, we see an elbow at three, one and two components for phases one to three respectively. The phase-specific scree ratios in Table 3 confirm that three, one and two components seem optimal, as the corresponding ratios are the highest. To examine whether the modeling assumptions regarding the residuals are tenable, we inspected boxplots of the residuals per phase, which are displayed in Figure 8 and reveal some but relatively mild heteroscedasticity.

[Insert Figure 7, Table 3, and Figure 8 about here]

When we inspect the estimated phases of the selected model (with a VAF% of about 84%), the first phase includes 127 time points, i.e., the complete baseline part and the first 12 time points of the inhalation part, which suggests that the CO₂ inhalation had a delayed effect on the level and the covariance of the physiological parameters (see Figure 4). The second phase contains the remaining time points of the CO₂ inhalation and the first eight recovery time points. The third phase comprises the 76 remaining ones. Thus, it seems warranted to label the three phases 'baseline', 'CO₂ inhalation' and 'recovery', respectively.

Given the estimated phases, we first examine the differences in the means $\bar{\mathbf{x}}^c$ ($c = 1, 2, 3$) in Table 4. Inspiratory (ViVol) and expiratory volumes (VeVol), minute ventilation (Vent) and the peak inspiratory acceleration (PiaAB) are obviously highest and respiratory timing parameters are lowest (more rapid breathing) in the CO₂ inhalation phase (see Figure 4). This is exactly what can be expected since inhaling CO₂-enriched air leads to an increase in arterial CO₂ pressure (pCO₂), due to the excess level of CO₂ in the blood (hypercapnic state), and thereby reflexively increases the neuronal respiratory drive (as reflected in PiaAB; Cotton, Sheiban, & Engel, 1988; DiFiore et al., 1997) and concomitantly also minute ventilation (Pappens et al., 2012; Woods, Charney, Loke, & Goodman, 1986). Compared to the baseline phase, heart rate is increased and the cardiac RR interval is decreased during the CO₂ inhalation phase. This can be explained by the increased 'work of breathing' (i.e., the respiratory muscles must work harder) during the CO₂ inhalation phase, as well as by the sympathetic activation associated with a hypercapnic state (Pappens et al., 2012). In the recovery phase, heart rate was still increased and respiratory timing parameters were still decreased, suggesting a slower or incomplete recovery in these parameters. This contrasts with the respiratory volumes and minute ventilation, which returned to their more natural levels (baseline).

[Insert Table 4 about here]

The between-phase and within-phase variances of the different variables in Table 5, show that especially the respiratory volume and ventilation measures and the peak inspiratory acceleration vary strongly between phases, and consequently less within the phases, as can be expected from the above described effects of CO₂ inhalation. Moreover, the within-phase variances indicate that the respiratory timing parameters show the highest variability during the baseline phase and are relatively more stable within the CO₂ inhalation and recovery phases. This is consistent with the presence of an imperative metabolic drive to breathe during

the CO₂ phase, which is gradually decreasing again in the recovery phase, leaving less room for behavioral and random influences on breathing that typically add to variability. However, for many of the other variables, the variance pattern is reversed in that variance is highest during the inhalation phase, due to the gradual changes in level.

[Insert Tables 5 and 6 about here]

Next, we inspect the phase loading matrices given in Table 6, which are standardized (see Footnote 3) and normalized varimax rotated to improve the interpretability. For the baseline phase, the first component can be labeled ‘ventilation by inspiration time’, given the highly positive loading of ventilation (Vent) and the highly negative ones of inspiration and total time (Ti, Tt). Thus, this component suggests that ventilation varies mostly because of variations in the speed of the inspiration. Similarly, the second component for the baseline phase can be called ‘volume and expiration time’, with higher respiratory volume scores (ViVol and VeVol) being associated with a longer expiration time (Te) (and consequently with a longer total time Tt). This co-occurrence of a higher respiratory volume and a longer expiration time may be explained by the fact that the participant sighed a few times during the baseline, as a sigh is characterized by a quick, deep inspiration followed by a slower expiration. The third baseline component is a heart rate component, indicating that a higher heart rate (HR) is – logically – related to a shorter RR interval.

The loadings on the single component of the CO₂ inhalation phase indicate that a higher ventilation (Vent) is achieved by a higher respiration volume (ViVol and VeVol) as well as by a faster inspiration (lower Ti and thus lower Tt), and is associated with an increase in heart rate (higher HR and lower RR) (see Figure 4).

The first component of the recovery phase can be read as a ‘ventilation by total respiration time’ component, indicating that the decrease in ventilation that shows up in the

phase means is accomplished by slower breathing (higher T_i , T_e and T_t). The second component of the recovery phase is a heart rate component, which is very similar to the third component of the baseline.

The finding that more components are needed to describe cardio-respiratory behavior in a resting state (baseline) than in a state where the respiratory system is heavily challenged by an imperative metabolic stimulus (i.e., CO_2), is interesting. From a dynamic system perspective, a healthy organism in a resting state should be able to flexibly generate appropriate responses to constantly changing demands. Such flexibility is assured by a variable system in which different components are only loosely coupled. However, when confronted with a demand that poses a threat to survival (e.g., excess levels of CO_2), an organism leaves this flexible state and moves into a more specific constellation that aims to compensate for the distortion. Consistent with this idea, the above described results indicated that cardiorespiratory parameters covaried less and were rather loosely coupled during a baseline resting state, supporting the capacity to flexibly respond to changing demands. However, when the system was challenged by excessive CO_2 (i.e., a threat to survival), all parameters became more coupled and behaved in synchrony in function of the imperative goal of increasing ventilation to get rid of the excess of CO_2 . As a consequence, less components are needed to describe cardio-respiratory behavior in this phase. In the recovery phase, cardiac and respiratory components became more loosely coupled again.

A data set with experimentally induced states makes a good test for the Switching PCA model. Indeed, when we expect that the experimental manipulations influence the means and/or covariances of the variables to some extent, we can apply Switching PCA to see whether it finds phases that are somehow related to the experimental manipulations. From the results above, we can conclude that Switching PCA was indeed able to distinguish between three phases in the physiological time series data which are clearly related to the experimental

manipulations. Finally, the information that was obtained from this exploratory analysis, could be useful to design more confirmatory models for similar data from another subject. For instance, one could specify a HMM model with three states in which the probability of a transition to a previously experienced state is zero and in which the covariance matrices in the different states are specified in accordance with the found component structures.

Discussion

Many behavioral research questions pertain to how the level and/or the covariance structure of a set of variables changes across time. Indeed, these phenomena are the corner stone of many psychological theories, where they are referred to as response patterning and synchronicity or coherence, respectively. Yet, by lack of appropriate models and associated algorithms most of the work on response patterning and synchronicity remains theoretical (Gross, 2010). Therefore, the development of methods like Switching PCA, which are able to capture both phenomena, is crucial to put such theories to the test.

In the future, the Switching PCA model could be further improved by taking into account or even modeling serial dependencies among subsequent time points. On the one hand, whenever the data can be assumed to consist of smooth curves perturbed by measurement error, smoothness restrictions (e.g., Timmerman & Kiers, 2002) could be imposed on the component scores of each Switching PCA phase to obtain better interpretable within-phase models. On the other hand, one could explicitly model autocorrelations and time-lagged covariances, as they are often theoretically interesting. For instance, in emotion theory, autocorrelations are used to measure emotional inertia or the tendency for affective states to be resistant to change (Koval & Kuppens, 2012), and cross-lagged covariances to measure augmentation or blunting of a certain emotion by another emotion at a previous time

point (Pe & Kuppens, 2012). Building on the seminal ideas in dynamic factor analysis (Molenaar, 1985) and dynamic PCA (Ku, Storer, & Georgakis, 1995), this could be modeled by means of lagged loading matrices (e.g., Timmerman, 2001), capturing the cross-lagged effects of the underlying dimensions on the observed variable scores, or by allowing autoregressive effects among the scores on a component as well as cross-lagged covariances between components (e.g., McArdle, 2009; Molenaar & Nesselroade, 2001; Rato & Reis, 2013; Zhang, Hamaker, & Nesselroade, 2008).

Secondly, our objective function (see Appendix A) implies, amongst others, the assumption that error variances are identical across phases and variables. For empirical data this assumption will often not hold, however. The simulation studies of the current paper give some information on the effect of violations of this assumption, as random fluctuations in error variances occurred between the variables and phases but hardly influenced the Switching PCA performance. In future research, it would be useful to manipulate such differences more systematically and thoroughly examine the robustness of Switching PCA to differences in residual variance as well as to other violations of model assumptions.

Thirdly, in the current implementation of Switching PCA, each phase pertains to a different state or ‘regime’ (i.e., with different means and/or components), making it difficult to unravel whether some response patterns or synchronization processes re-occur. Indeed, if we would have added an additional CO₂ inhalation and recovery phase to the application, it would be more parsimonious to fit a model with three regimes, two of which occur twice, instead of a model with five distinct phases. Therefore, in future research, it would be worthwhile to develop a Switching PCA variant which allows people to switch back and forth between a few regimes. This extension of the model would require some intricate adaptations, however, regarding both model estimation (i.e., simultaneously updating the phase transitions and the regime-specific means and component structures) and selection (i.e., how many

regimes and how many phases – reoccurrences – per regime). Until then, the following post-hoc approach may be useful to investigate the recurrence of means and covariance structures: Whenever one deems that two phases are very alike in terms of means and component loadings, one may collapse the data of these phases and re-estimate the PCA model for the thus combined data. Subsequently, one can compare the fit of the resulting model with that of the original Switching PCA model. If both values are very similar the collapsed model is more parsimonious and thus indicated. However, this comparison of the fit is only fair when the collapsed model does not imply a change in the number of components for one of the phases; e.g., when the third (recovery) phase of the application is collapsed with the first (baseline) phase and the combined data is modeled by three components, the fit of the model will most likely increase because the recovery phase is fitted with three components instead of two.

Fourthly, in Switching PCA, all time points within a phase are modeled with the same vector of means and the same components, which resembles the local weak stationarity assumption (e.g., Francq & Zakoïan, 2001; Liu, 2006) that is used in many time series models including hidden markov models and regime-switching state-space models. Nonetheless, Switching PCA has the advantage that more gradual changes in the mean level of the variables will be picked up quite nicely as changes in the covariance and thus by the within-phase components. For example, in our application (see Figure 4), the peak inspiratory acceleration and the respiratory volume and ventilation measures increase gradually from time point 120, while respiration time and heart rate simultaneously decrease, and around time point 155 everything starts to normalize again. These gradual changes in the mean levels of the variables are captured by the component of the second phase that explains quite a lot of variance.

Finally, Switching PCA could be extended to incorporate the data of multiple subjects at once. Specifically, when time series data are available on the same set of variables for a number of subjects, it is possible that the same structural phases underlie the data of all the subjects. For instance, when some individuals are exposed to the experimental manipulation of a certain emotion eliciting stimulus, it is reasonable to assume that the covariational changes in their physiological data may be similar. To gain insight into these structural changes, the same number of phases could be imposed on all subjects. Stated differently, the data of multiple subjects could be modeled by developing a switching variant of simultaneous component analysis (SCA; Kiers, 1990; Kiers & ten Berge, 1994; Timmerman & Kiers, 2003).

References

- Aach, J., & Church, G. M. (2001). Aligning gene expression time series with time warping algorithms. *Bioinformatics*, *17*, 495-508.
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*, 716-723.
- Aue, A., Hörmann, S., Horváth, L., & Reimherr, M. (2009). Break detection in the covariance structure of multivariate time series models. *The Annals of Statistics*, *37*, 4046-4087.
- Bauer, D. J. & Curran, P. J. (2004). The integration of continuous and discrete latent variable models: potential problems and promising opportunities. *Psychological Methods*, *9*, 3-29.
- Baum, L. E., & Petrie, T. (1966). Statistical inference for probabilistic functions of finite state Markov chains. *Annals of Mathematical Statistics*, *37*, 1554-1563.
- Besse, P. & Ramsay, J. O. (1986). Principal components analysis of sampled functions. *Psychometrika*, *51*, 285-311.
- Boiten, F. (1993). Component analysis of task-related respiratory patterns. *International Journal of Psychophysiology*, *15*, 91-104.
- Boker, S. M., Neale, M. C., & Rausch, J. (2004). Latent differential equation modeling with multivariate multi-occasion indicators. In K. van Montfort, H. Oud, & A. Satorra (Eds.), *Recent developments on structural equation models: Theory and applications* (pp. 151-174). Dordrecht, Netherlands: Kluwer Academic Publishers.
- Borsboom, D., Mellenbergh, G. J., & van Heerden, J. (2003). The theoretical status of latent variables. *Psychological Review*, *110*, 203-219.
- Bulteel, K., Wilderjans, T. F., Tuerlinckx, F., & Ceulemans, E. (2013). CHull as an alternative to AIC and BIC in the context of mixtures of factor analyzers. *Behavior Research Methods*. Advance online publication. doi:10.3758/s13428-012-0293-y

- Butler, E. A. (2011). Temporal interpersonal emotion systems: The “TIES” that form relationships. *Personality and Social Psychology Review*, 15, 367–393.
- Cattell, R. B. (1952). The three basic factor analytic designs: Their interrelations and derivatives. *Psychological Bulletin*, 49, 499–520.
- Ceulemans, E., & Kiers, H. A. L. (2006). Selecting among three-mode principal component models of different types and complexities: A numerical convex hull based method. *British Journal of Mathematical and Statistical Psychology*, 59, 133–150.
- Ceulemans, E., Timmerman, M. E., & Kiers, H. A. L. (2011). The CHULL procedure for selecting among multilevel component solutions. *Chemometrics and Intelligent Laboratory Systems*, 106, 12–20.
- Christie, I. C., & Friedman, B. H. (2004). Autonomic specificity of discrete emotion and dimensions of affective space: a multivariate approach. *International Journal of Psychophysiology*, 51, 143–153.
- Cotton, D. J., Sheiban, M., & Engel, L. A. (1988). Volume acceleration as an index of neuromuscular output. *Respiration Physiology*, 71, 117–130.
- De Roover, K., Ceulemans, E., Timmerman, M. E., Nezlek, J. B., & Onghena, P. (2013). Modeling differences in the dimensionality of multiblock data by means of clusterwise simultaneous component analysis. *Psychometrika*. Advance online publication. doi:10.1007/s11336-013-9318-4
- De Roover, K., Ceulemans, E., & Timmerman, M. E. (2012). How to perform multiblock component analysis in practice. *Behavior Research Methods*, 44, 41–56.
- Diebold, F. X., Lee, J. H., & Weinbach, G. C. (1994). Regime switching with time-varying transition probabilities. In: P. Hargreaves (Ed.), *Nonstationary time series analysis and cointegration* (pp. 283–302). Oxford: Oxford University Press.
- DiFiore, J. M., Martin, R. J., Dreshaj, I., Abu-Shaweesh, J., Neuman, M. R., Adams, J. A., &

- Sackner, M. A. (1997). Peak acceleration of respiratory waveforms at onset of inspiration: A new noninvasive measure of respiratory drive. *Pediatric Research*, 41, 251.
- Dolan, C. V., Schmittman, V. D., Lubke, G. H., & Neale, M. C. (2005). Regime switching in the latent growth curve mixture model. *Structural Equation Modeling*, 12, 94–119.
- Fannes, S., Van Diest, I., Meulders, A., De Peuter, S., Vansteenwegen, D., Van den Bergh, O. (2008). To inhale or not to inhale: Conditioned avoidance in breathing behavior in an odor - 20% CO₂ paradigm. *Biological Psychology*, 78, 87–92.
- Filardo, A. J. (1998). Choosing information variables for transition probabilities in a time-varying transition probability model. Federal Reserve Bank of Kansas City RWP 98-09.
- Flavell, J. H. (1971). Stage-related properties of cognitive development. *Cognitive Psychology*, 2, 421–453.
- Floyd, F. J., & Widaman, K. F. (1995). Factor analysis in the development and refinement of clinical assessment instruments. *Psychological Assessment*, 7, 286–299.
- Fraley, R. C. (2007). A connectionist approach to the organization and continuity of working model of attachment. *Journal of Personality*, 75, 1157–1180.
- Francq, C., & Zakoïan, J.-M. (2001). Stationarity of multivariate markov-switching ARMA models. *Journal of Econometrics*, 102, 339–364.
- Friston, K. J., Frith, C. D., Liddle, P. F., & Frackowiak, R. S. J. (1993). Functional connectivity: The principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow and Metabolism*, 13, 5–14.
- Gales, M. J. F. (1999). Semi-tied covariance matrices for hidden markov models. *IEEE Transactions on Speech and Audio Processing*, 7, 272–281.

- Graham, J. W., Collins, L. M., Wugalter, S. E., Chung, N. K., & Hansen, W.B. (1991). Modeling transitions in latent stage-sequential processes: A substance use prevention example. *Journal of Consulting and Clinical Psychology*, 59, 48–57.
- Gross, J. J. (2010). The future's so bright, I gotta wear shades. *Emotion Review*, 2, 212–216.
- Hair, J. F., Black, W. C., Babin, J. B., Anderson, R. E., & Tatham, R. L. (2009). *Multivariate Data Analysis* (6th ed.). Upper Saddle River, NJ: Pearson Education Inc.
- Hamilton, J. D. (1994). *Time series analysis*. Princeton, N.J.: Princeton University Press.
- Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, 2, 193–218.
- Hudson, D. J. (1966). Fitting segmented curves whose join points have to be estimated. *Journal of the American Statistical Association*, 61, 1097–1129.
- Jolliffe, I. T. (2002). *Principal component analysis* (2nd ed.). New York: Springer.
- Jones, C. J., & Nesselroade, J. R. (1990). Multivariate, replicated, single-subject, repeated measures designs and P-technique factor analysis: a review of intraindividual change studies. *Experimental aging research*, 16, 171–183.
- Kiefer, N. M. (1978). Discrete parameter variation: Efficient estimation of a switching regression model. *Econometrica*, 46, 427–434.
- Kiers, H. A. L. (1990). *SCA. A program for simultaneous components analysis of variables measured in two or more populations*. Groningen, The Netherlands: iec ProGAMMA.
- Kiers, H. A. L., & ten Berge, J. M. F. (1994). Hierarchical relations between methods for Simultaneous Components Analysis and a technique for rotation to a simple simultaneous structure. *British Journal of Mathematical and Statistical Psychology*, 47, 109–126.
- Kim, C.-J., & Nelson, C. R. (1999). *State-space models with regime-switching: classical and Gibbs-sampling approaches with applications*. Cambridge, MA: MIT Press.

- Klausmeier, H. J., & Sipple, T. S. (1982). Factor structure of the Piagetian stage of concrete operations. *Contemporary Educational Psychology*, 7, 161–180.
- Koval, P., & Kuppens, P. (2012). Changing emotion dynamics: Individual differences in the effect of anticipatory social stress on emotional inertia. *Emotion*, 12, 256–267.
- Ku, W., Storer, R. H., & Georgakis, C. (1995). Disturbance detection and isolation by dynamic principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, 30, 179–196.
- Lavielle, M., & Teyssière, G. (2006). Detection of multiple change-points in multivariate time series. *Lithuanian Mathematical Journal*, 46, 287–306.
- Lawley, D. N., & Maxwell, A. E. (1962). Factor analysis as a statistical method. *The Statistician*, 12, 209–229.
- Lawson, A. E., & Nordland, F. H. (1976). The factor structure of some Piagetian tasks. *Journal of Research in Science Teaching*, 13, 461–466.
- Li, F., Duncan, T. E., Duncan, S. C., & Acock, A. (2001). Latent growth modeling of longitudinal data: A finite growth mixture modeling approach. *Structural Equation Modeling*, 8, 493–530.
- Liu, J., Wu, S., & Zidek, J. V. (1997). On segmented multivariate regression. *Statistica Sinica*, 7, 497–525.
- Liu, J.-C. (2006). Stationarity of a markov-switching GARCH model. *Journal of Financial Econometrics*, 4, 573–593.
- Lodewyckx, T., Tuerlinckx, F., Kuppens, P., Allen, N. B., & Sheeber, L. B. (2011). A hierarchical state space approach to affective dynamics. *Journal of Mathematical Psychology*, 55, 68–83.
- Lorenzo-Seva, U., & ten Berge, J. M. F. (2006). Tucker's congruence coefficient as a meaningful index of factor similarity. *Methodology*, 2, 57–64.

- Lu, R.-P., Smith, E. P., & Good, I. J. (1989). Multivariate measures of similarity and niche overlap. *Theoretical Population Biology*, 35, 1–21.
- Mauss, I. B., Levenson, R. W., McCarter, L., Wilhelm, F. H., & Gross, J. J. (2005). The tie that binds? Coherence among emotion experience, behavior, and physiology. *Emotion*, 5, 175–190.
- McArdle, J. J. (2009). Latent variable modeling of differences and changes with longitudinal data. *Annual Review of Psychology*, 60, 577–605.
- McArdle, J. J., & Epstein, D. (1987). Latent growth curves within developmental structural equation models. *Child Development*, 58, 110–133.
- McLachlan, G. J., & Peel, D. (2000). *Finite mixture models*. New York: Wiley.
- Meredith, W., & Millsap, R. E. (1985). On component analyses. *Psychometrika*, 50, 495–507.
- Milligan, G. W., Soon, S. C., & Sokol, L. M. (1983). The effect of cluster size, dimensionality, and the number of clusters on recovery of true cluster structure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5, 40–47.
- Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological Review*, 102, 246–268.
- Molenaar, P. C. M. (1985). A dynamic factor model for the analysis of multivariate time series. *Psychometrika*, 50, 181–202.
- Molenaar, P. C. M., & Nesselroade, J. R. (2001). Rotation in the dynamic factor modeling of multivariate stationary time series. *Psychometrika*, 66, 99–107.
- Molenaar, P. C. M., & von Eye, A. (1994). On the arbitrary nature of latent variables. In A. von Eye & C. C. Clogg (Eds.), *Latent variable analysis: Applications for developmental research* (pp. 226–242). Thousand Oaks, CA: Sage

- Muthén, B., Brown, C. H., Booil Jo, K. M., Khoo, S.-T., Yang, C.-C., Wang, C.-P. et al. (2002). General growth mixture modeling for randomized preventive interventions. *Biostatistics*, 3, 459–475.
- Nagin, D. S. (1999). Analyzing developmental trajectories: a semi-parametric, group-based approach. *Psychological Methods*, 4, 139–157.
- Pappens, M., de Peuter, S., Vansteenwegen, D., Van den Bergh, O., & Van Diest, I. (2012). Psychophysiological responses to CO₂ inhalation. *International Journal of Psychophysiology*, 84, 45–50.
- Pe, M. L., & Kuppens, P. (2012). The dynamic interplay between emotions in daily life: Augmentation, blunting, and the role of appraisal overlap. *Emotion*, 12, 1320–1328.
- Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2, 559–572.
- Piaget, J. (1972). Intellectual evolution from adolescence to adulthood. *Human Development*, 15, 1–12.
- Raftery, A. E. (1985). A model for high-order Markov chains. *Journal of the Royal Statistical Society. Series B (Methodological)*, 47, 528–539.
- Ramadan, Z., Jacobs, D., Grigorov, M., & Kochlar, S. (2006). Metabolic profiling using principal component analysis, discriminant partial least squares, and genetic algorithms. *Talanta*, 68, 1683–1691.
- Rao, C. R. (1958). Some statistical methods for the comparison of growth curves. *Biometrics*, 14, 1–17.
- Rato, T. J., & Reis, M. S. (2013). Defining the structure of DPCA models and its impact on process monitoring and prediction activities. *Chemometrics and Intelligent Laboratory Systems*, 125, 74–86.

- Repp, B. H. (2005). Sensorimotor synchronization: A review of the tapping literature. *Psychonomic Bulletin and Review*, 12, 969–992.
- Rice, J. A., & Silverman, B. W. (1991). Estimating the mean and covariance structure nonparametrically when the data are curves. *Journal of the Royal Statistical Society. Series B (Methodological)*, 53, 233–243.
- Robert, P., & Escoufier, Y. (1976). A unifying tool for linear multivariate statistical methods: the RV-coefficient. *Applied Statistics*, 25, 257–265.
- Rosti, A-V. I., & Gales, M. J. F. (2004). Factor analysed hidden markov models for speech recognition. *Computer Speech and Language*, 18, 181–200.
- Sander, D., Grandjean, D., & Scherer, K. R. (2005). A systems approach to appraisal mechanisms in emotion. *Neural Networks*, 18, 317–352.
- Schmittmann, V. D., Dolan, C. V., van der Maas, H. L. J., & Neale, M. C. (2005). Discrete latent markov models for normally distributed response data. *Multivariate Behavioral Research*, 40, 461–488.
- Sullivan, J. H., & Woodall, W. H. (2000). Change-point detection of mean vector or covariance matrix shifts using multivariate individual observations. *IIE Transactions*, 32, 537–549.
- Timmerman, M. E. (2001). Component analysis of multisubject multivariate longitudinal data. Unpublished doctoral dissertation. University of Groningen, Groningen, The Netherlands.
- Timmerman, M. E., Ceulemans, E., De Roover, K., & Van Leeuwen, K. (2013). Subspace K-means clustering. *Behavior Research Methods*. Advance online publication. doi:10.3758/s13428-013-0329-y

- Timmerman, M. E., & Kiers, H. A. L. (2000). Three-mode principal component analysis: Choosing the numbers of components and sensitivity to local optima. *British Journal of Mathematical and Statistical Psychology*, 53, 1–16.
- Timmerman, M. E., & Kiers, H. A. L. (2002). Three-way component analysis with smoothness constraints. *Computational Statistics and Data Analysis*, 40, 447–470.
- Timmerman, M. E., & Kiers, H. A. L. (2003). Four simultaneous component models of multivariate time series from more than one subject to model intraindividual and interindividual differences. *Psychometrika*, 86, 105–122.
- Tucker, L. R. (1951). *A method for synthesis of factor analysis studies* (Personnel Research Section Report No. 984). Washington, DC: Department of the Army.
- Tucker, L. R. (1958). Determination of parameters of a functional relation by factor analysis. *Psychometrika*, 23, 19–23.
- Velicer, W. F., Martin, R. A. & Collins, L. M. (1996). Latent transition analysis for longitudinal data. *Addiction*, 91, S197–S209.
- Wilderjans, T. F., Ceulemans, E., Kiers, H. A. L., & Meers, K. (2009). The LMPCA program: A graphical user interface for fitting the linked-mode PARAFAC-PCA model to coupled real-valued data. *Behavior Research Methods*, 41, 1073–1082.
- Wilderjans, T. F., Ceulemans, E., Meers, K. (2013). CHull: A generic convex-hull-based model selection method. *Behavior Research Methods*, 45, 1–15.
- Wilderjans, T. F., Ceulemans, E., Van Mechelen, I., & van den Berg, R. A. (2011). Simultaneous analysis of coupled data matrices subject to different amounts of noise. *British Journal of Mathematical and Statistical Psychology*, 64, 277–290.
- Wohlwill, J. F. (1973). *The study of behavioural development*. New York: Academic Press.
- Woods, S. W., Charney, D. S., Loke, J., & Goodman, W. K. (1986). Carbon dioxide sensitivity in panic anxiety. *Archives of General Psychiatry*, 43, 900–909.

- Yung, Y. F. (1997). Finite mixtures in confirmatory factor-analysis models. *Psychometrika*, 62, 297–330.
- Zhang, Z., Hamaker, E. L., & Nesselroade, J. R. (2008). Comparison of four methods for estimating a dynamic factor model. *Structural Equation Modeling*, 15, 377–402.

Appendix A: Derivation of an AIC-based objective function for Switching PCA

We start from the *AIC* (Akaike, 1974) to derive an objective function for Switching PCA that penalizes the complexity of the within-phase models. The *AIC* is given by:

$$AIC = -2\text{loglik}(\mathbf{X}|\mathbf{M}) + 2fp, \quad (8)$$

where $\text{loglik}(\mathbf{X}|\mathbf{M})$ denotes the loglikelihood of data $\mathbf{X}$ given model $\mathbf{M}$ and fp is the number of free parameters to be estimated. Assuming the residuals e_{nj} to be independent and identically distributed as $e_{nj} \sim N(0, \sigma^2)$, the loglikelihood of a Switching PCA model reads as follows:

$$\text{loglik}(\mathbf{X}|\mathbf{M}) = \log \left[\left(\frac{1}{2\pi\sigma^2} \right)^{\frac{NJ}{2}} \exp \left(-\frac{SSE}{2\sigma^2} \right) \right] = -\frac{NJ}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} SSE. \quad (9)$$

Using $\hat{\sigma}^2 = \frac{SSE}{NJ}$ as a post-hoc estimator of the error variance σ^2 (Wilderjans, Ceulemans, Van Mechelen, & van den Berg, 2011), the loglikelihood boils down to:

$$\text{loglik}(\mathbf{X}|\mathbf{M}) = -\frac{NJ}{2} \left[1 + \log(2\pi) - \log(NJ) + \log(SSE) \right], \quad (10)$$

where the first three terms are invariant across solutions and thus can be discarded when minimizing the objective function.

The number of free parameters fp is given by:

$$fp = (C-1) + \sum_{c=1}^C fp^c = (C-1) + \sum_{c=1}^C (J + N^c Q^c + JQ^c - (Q^c)^2), \quad (11)$$

where $(C-1)$ reflects the number of phase separations and fp^c is the number of free parameters within each phase c , with the first three terms of fp^c indicating the number of phase-specific

variable means, component scores⁷, and loadings (Ceulemans & Kiers, 2006). The fourth term of fp^c corrects for the rotational freedom. When estimating a model given specific values of C and Q^c , all terms of fp are invariant except for the number of component scores, and can therefore be discarded. Thus, the following *AIC*-based objective function L is obtained:

$$L = NJ \log(SSE) + 2 \sum_{c=1}^C N^c Q^c. \quad (12)$$

⁷ The number of free parameters includes the number of component scores because component analysis always provides estimates of the component scores, as opposed to factor analysis, wherein factor scores are usually not explicitly estimated.

Appendix B: The Switching PCA algorithm

To estimate the Switching PCA parameters, the ALS algorithm performs the following steps:

- 1) Randomly segment the time series into C phases: The algorithm selects at random $C-1$ phase transition points among the N time points, provided that the resulting phases contain at least J time points. This initializes the u_{nc} scores in Equation 2.

- 2) Estimate the phase-specific means and PCA models of the initial phases: The phase

means are calculated as⁸ $\bar{\mathbf{x}}^c = \frac{1}{N^c} \sum_{n=1}^N u_{nc} (\mathbf{x}_n - \mathbf{m})$. Given the singular value

decomposition of the centered data of phase c (i.e., $(\mathbf{X}^c - \mathbf{1}_{N^c}'(\mathbf{m} + \bar{\mathbf{x}}^c)) = \mathbf{U}^c \mathbf{S}^c \mathbf{V}^c$), the

component scores $\mathbf{F}^c$ of all N^c time points in the c -th phase and the corresponding

loadings $\mathbf{B}^c$ are estimated as $\mathbf{F}^c = \sqrt{N^c} \mathbf{U}_{Q^c}^c$ and $\mathbf{B}^c = \sqrt{\frac{1}{N^c}} \mathbf{V}_{Q^c}^c \mathbf{S}_{Q^c}^c$ respectively. $\mathbf{U}_{Q^c}^c$

and $\mathbf{V}_{Q^c}^c$ are the first Q^c columns of $\mathbf{U}^c$ and $\mathbf{V}^c$, and $\mathbf{S}_{Q^c}^c$ consists of the first Q^c rows

and columns of $\mathbf{S}^c$. Without loss of fit, this step implies some orthogonality

constraints (i.e., the covariance matrix among the components of phase c equals an

identity matrix) that are not strictly necessary and can be removed by applying an

oblique rotation⁹.

⁸ The overall means of the variables in vector $\mathbf{m}$ are calculated at the start of the algorithm and remain unchanged throughout the rest of the analysis.

⁹ Note that in EFA it is possible to estimate the relevant parameters without imposing orthogonality constraints, which may be considered an advantage of this approach.

- 3) Find the best possible transition points: The algorithm searches for the optimal position for each phase transition, while keeping the other phase transitions fixed (see Figure 1) and subject to the restriction that each phase contains at least J time points. With each relocation of a phase transition, the affected mean vectors and PCA models are updated (as described in step 2), and the resulting objective function value L is computed. This procedure (i.e., step 3) is repeated until the phase transitions no longer change.

Appendix C: The Switching PCA model selection procedure

If one fits all models with one up to $C^{\max}$ clusters and with one up to $Q^{\max}$ phase-specific components, the total number of models to choose from equals:

$$S = \sum_{C=1}^{C^{\max}} \frac{(Q^{\max} + C - 1)!}{((Q^{\max} - 1)! C!)} \quad (13)$$

For example, if $C^{\max}$ and $Q^{\max}$ equal six, 923 different models are obtained. In order to avoid the computationally demanding estimation of all these models, we propose a more time-efficient, stepwise model selection strategy, which consists of the following steps:

- 1) Initial screening: Select the best model from the subset of models that impose an equal number of components Q^c in all phases (only these models have to be fitted): First, determine the best number of phases $C_{(Q)}^{\text{best}}$, among $C = 1, \dots, C^{\max}$, for each Q -value ($Q = 0, \dots, Q^{\max}$) separately, by maximizing the following scree ratio which indicates that the increase in fit with additional phases levels off (see, e.g., Figure 5):

$$sr_{(C|Q)} = \frac{\text{dif}_{(C|Q)}}{\text{dif}_{(C+1|Q)}}, \quad (14)$$

where $\text{dif}_{(C|Q)}$ equals $[\text{VAF}_{(C|Q)} - \text{VAF}_{(C-1|Q)}]$ or the increase in variance accounted for (VAF) when going from $C-1$ to C phases, given Q components. The overall best number of phases C^{best} is determined as the C -value which has the highest average scree ratio across the Q -values. Note that Q equal to zero implies that a model is used that takes only differences in means into account, ignoring covariance differences.

Second, assess the best number of components Q^{best} , given C^{best} phases, by maximizing:

$$sr_{(Q|C^{\text{best}})} = \frac{\text{dif}_{(Q|C^{\text{best}})}}{\text{dif}_{(Q+1|C^{\text{best}})}}, \quad (15)$$

with $\text{dif}_{(Q|C^{\text{best}})}$ equal to $[\text{VAF}_{(Q|C^{\text{best}})} - \text{VAF}_{(Q-1|C^{\text{best}})}]$ or the increase in VAF when going from $Q-1$ to Q components, given C^{best} phases. Note that $\text{VAF}_{(0|C^{\text{best}})}$ (i.e., with zero components within the phases) is usually larger than zero due to the phase means.

- 2) Re-evaluate the best number of components $Q^{c,\text{best}}$ for each phase separately: Using the segmentation found in step 1, compute PCA solutions with one up to Q^{max} components per phase (e.g., see Figure 7). For each phase, determine $Q^{c,\text{best}}$ as the number of components that maximizes the phase-specific scree ratio:

$$sr_{(Q^c)}^c = \frac{\text{dif}_{(Q^c)}^c}{\text{dif}_{(Q^c+1)}^c}, \quad (16)$$

where $\text{dif}_{(Q^c)}^c$ equals $[\text{VAF}_{(Q^c)}^c - \text{VAF}_{(Q^c-1)}^c]$ or the increase in VAF within phase c when adding a Q^c th component. Note that $\text{dif}_{(1)}^c$ equals $\text{VAF}_{(1)}^c$.

- 3) Estimate the Switching PCA model with the selected numbers of phases and components: Perform a Switching PCA analysis with 25 random starts and a 26th ‘rational’ start using the segmentation obtained in step 1, and retain the best solution.
- 4) Check convergence: Repeat step 2 using the phase segmentation that results from step 3, to evaluate whether changes in this segmentation affect the selected Q^c s. If this is the case, repeat step 3 and 4.

Table 1

The nine variables of the physiological time series data set.

Abbreviation	Variable Description	Measurement Unit
ViVol	inspiratory volume: volume of air inhaled during inspiration	milliliter
VeVol	expiratory volume: volume of air inhaled during expiration	milliliter
Vent	minute ventilation: amount of air displaced by breathing on a one minute basis	liter/minute
Ti	inspiration time: duration of inspiration	seconds
Te	expiration time: duration of expiration	seconds
Tt	total respiration cycle time: duration of total breath cycle	seconds
PiaAB	peak inspiratory acceleration of abdominal signal: maximum acceleration of the airflow during inspiration derived from the abdominal respiratory signal, an estimate of neural drive to inspire	liter/seconds ²
HR	heart rate: number of heart beats on a one minute basis	beats/minute
RR	RR interval or cardiac interbeat interval: time between subsequent R-peaks of the electrocardiogram	milliseconds

Table 2

Scree ratios for the numbers of phases C given the fixed numbers of components Q and averaged over the numbers of components (above), and for the fixed numbers of components Q given three phases (below), for the physiological time series data. The maximal scree ratio in each column is highlighted in bold face.

	0 comp.	1 comp.	2 comp.	3 comp.	4 comp.	average
2 phases	0.43	1.47	1.31	2.69	1.52	1.48
3 phases	6.27	2.81	1.47	1.18	2.05	2.76
4 phases	2.06	1.48	1.89	1.85	1.17	1.69
5 phases	1.73	0.98	2.59	1.46	0.94	1.54

	3 phases
1 component	1.60
2 components	1.43
3 components	1.89

Table 3

Scree ratios for the numbers of components Q^c for phase 1, phase 2 and phase 3, for the physiological time series data. The maximal scree ratio in each column is highlighted in bold face.

	Phase 1	Phase 2	Phase 3
1 component	1.50	2.64	1.37
2 components	1.26	1.33	2.07
3 components	1.91	1.20	1.61

Table 4

The overall means $\mathbf{m}$ (of the rescaled data) and the phase means $\bar{\mathbf{x}}^c$ (in terms of deviations from the overall offsets) of the variables for the Switching PCA model with three phases for the physiological time series data.

	Overall means	Phase means		
		Phase 1: Baseline	Phase 2: CO ₂ inhalation	Phase 3: Recovery
ViVol	3.17	-0.20	1.82	-0.53
VeVol	3.13	-0.24	1.82	-0.47
Vent	2.65	-0.39	1.94	-0.27
Ti	3.18	0.39	-0.55	-0.40
Te	3.79	0.17	-0.21	-0.19
Tt	5.08	0.42	-0.56	-0.44
PiaAB	1.39	-0.37	2.08	-0.38
HR	10.77	-0.22	0.04	0.35
RR	11.71	0.24	-0.01	-0.39

Table 5

Between- and within-phase variances (proportions to the total variance between brackets)
for the Switching PCA model with three phases for the physiological time series data.*

	Between phases	Phase 1: Baseline	Phase 2: CO ₂ inhalation	Phase 3: Recovery
ViVol	.61	0.51 (.27)	0.37 (.06)	0.22 (.07)
VeVol	.60	0.43 (.23)	0.59 (.09)	0.26 (.08)
Vent	.67	0.32 (.17)	0.56 (.08)	0.23 (.07)
Ti	.18	1.05 (.56)	0.57 (.09)	0.55 (.18)
Te	.03	1.17 (.62)	0.57 (.09)	0.82 (.26)
Tt	.20	0.93 (.50)	0.40 (.06)	0.75 (.24)
PiaAB	.77	0.20 (.11)	0.35 (.05)	0.22 (.07)
HR	.07	0.86 (.46)	1.30 (.20)	0.88 (.28)
RR	.08	0.89 (.47)	1.35 (.20)	0.78 (.25)

* Note that the proportions of the within-phase variances to the total variance are relatively

larger in case of a larger phase length.

Table 6

Normalized varimax rotated standardized loadings for the Switching PCA model with three phases for the physiological time series data. Loadings greater than +/- .40 are highlighted in bold face.

	Phase 1: Baseline			Phase 2: CO ₂ inhalation	Phase 3: Recovery	
	Vent by Ti	Vol & Te	HR	Vent by Ti+Vol & HR	Vent by Tt	HR
ViVol	.06	.71	.16	.33	-.31	.11
VeVol	.15	.69	-.02	.47	-.28	.08
Vent	.60	.21	.31	.54	.55	.10
Ti	-.96	-.16	-.01	-.47	-.56	.09
Te	-.01	.92	-.22	.02	-.85	.00
Tt	-.79	.58	-.17	-.41	-.98	.06
PiaAB	.18	.16	.16	.23	.04	-.01
HR	.14	-.05	.98	.96	-.13	.99
RR	-.15	.05	-.98	-.96	.09	-.99

Figure 1. Illustration of step 3 of the Switching PCA algorithm (see Data Analysis section) for an analysis with three phases. The dotted line represents the multivariate time series and the vertical lines represent the phase separations. The phase separation that is updated, whilst keeping the other separation fixed, is indicated by the ' $\leftarrow ? \rightarrow$ '. ' $\text{PCA}^{(1)}$ ', ' $\text{PCA}^{(2)}$ ', and ' $\text{PCA}^{(3)}$ ' indicate the within-phase PCA models (including the phase-specific vectors $\bar{\mathbf{x}}^c$), where '*' indicates that the model is updated with each tentative relocation of the phase separation being optimized. Thus, the figure illustrates that, first, the separation between phases 1 and 2 is optimized – implying updates of $\text{PCA}^{(1)}$ and $\text{PCA}^{(2)}$ – and, next, the separation between phases 2 and 3 – implying updates of $\text{PCA}^{(2)}$ and $\text{PCA}^{(3)}$.

A)

$$\begin{aligned}
\mathbf{B}^{(1)} &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} & \mathbf{B}^{(2)} &= \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} & \mathbf{B}^{(3)} &= \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} & \mathbf{B}^{(4)} &= \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}
\end{aligned}$$

B)

$$\begin{aligned}
\mathbf{B}^{(1)} &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} & \mathbf{B}^{(2)} &= \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{bmatrix} & \mathbf{B}^{(3)} &= \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} & \mathbf{B}^{(4)} &= \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}
\end{aligned}$$

Figure 2. Illustration of the simple structure component loadings (level 2 of factor 6): in case the Q^c -vector equals [4 4 4 4] and J equals 12, the simple structure phase loading matrices were constructed as in panel A, where, e.g., the set of variables constituting the first component differs in one variable when comparing subsequent phases (the same holds for the second, third and fourth component); in case the Q^c -vector equals [4 2 4 2] and J equals 12, the simple structure phase loading matrices were constructed as in panel A, where merging the first two (respectively last two) components of $\mathbf{B}^{(1)}$ gives the first (respectively last) component of $\mathbf{B}^{(2)}$ (A similar relationship exists between $\mathbf{B}^{(3)}$ and $\mathbf{B}^{(4)}$).

Figure 3. Mean *ARI* and associated 95% confidence intervals as a function of the amount of between-phase structural variance and the number of phases (panel A); and mean *GOPL* and associated 95% confidence intervals as a function of the amount of between-phase structural variance and the amount of error variance ($e \times 100\%$) (panel B).

Figure 4. Nine cardio-respiratory responses from a young, healthy man during three experiment phases. The solid vertical lines indicate the exact timing of the manipulations and the dashed ones the phases as recovered by Switching PCA.

Figure 5. Percentage of total variance accounted for by Switching PCA solutions for the physiological time series data, with the number of phases varying from one to six and the number of components varying from zero to four.

Figure 6 Scores of the time points on the first three principal components for the physiological time series data. The solid vertical lines indicate the timing of the manipulations.

Figure 7. Percentage of within-phase variance accounted for by PCA solutions within phase 1 (baseline), phase 2 (CO_2 inhalation) and phase 3 (recovery) for the physiological time series data, with the number of components varying from one to four.

Figure 8. Boxplots of the phase- and variable-specific residuals for the Switching PCA model with three phases and three, one and two components within the phases for the physiological time series data.