Recovering Bistable Systems from Psychological Time Series

Jonas Haslbeck^{*†1} and Oisín Ryan^{*‡2}

¹Psychological Methods Group, University of Amsterdam

²Department of Methodology and Statistics, Utrecht University

A revised version of this manuscript with the title "Recovering Within-Person Dynamics from Psychological Time Series" is available on PsyArXiv: <https://psyarxiv.com/dymhw>

Abstract

Conceptualizing mental disorders as complex dynamical systems has become a popular framework to study mental disorders. Especially bistable dynamical systems have received much attention, because their properties map well onto many characteristics of mental disorders. While these models were so far mostly used as stylized toy models, the recent surge in psychological time series data promises the ability to recover such models from data. In this paper we investigate how well popular (e.g., the Vector Autoregressive model) and more advanced (e.g., differential equation estimation) data analytic tools are suited to recover bistable dynamical systems from time series. Using a simulated high-frequency time series (measurement every six seconds) as an ideal case we show that while it is possible to recover global dynamics (e.g., position of fixed points, transition probabilities) it is difficult to recover the microdynamics (i.e., moment to moment interactions) of a bistable system. Repeating all analyses with a sampling frequency typical for Experience Sampling Method studies (measurement every 90 minutes) showed that the recovery of the global dynamics was still successful, but no microdynamics could be recovered. These results raise two fundamental issues involved in studying mental disorders from a complex systems perspective: first, it is generally unclear what to conclude from a statistical model about an underlying complex systems model; and second, if the sampling frequency is too low, it is impossible to recover microdynamics. In response to these results we propose a new modeling strategy based on substantively plausible dynamical systems models.

1 Introduction

Conceptualizing mental disorders as complex dynamical systems has become a popular framework to study mental disorders (e.g., Borsboom, 2017; Cramer et al., 2016; Wichers, Wigman, & Myin-Germeys, 2015). This framework is attractive because it acknowledges the fact that many mental disorders are massively multifactorial (e.g., Kendler, 2019), and because it allows one to specify powerful within-person dynamical systems models that capture many of the characteristics hypothesized for mental disorders. The central goal of this framework is to obtain such models to further our understanding of mental disorders, and allow us to develop and test more successful interventions.

The class of dynamic systems that has received most attention in this emerging literature is the class of *bistable systems* (e.g., Borsboom, 2017; Cramer et al., 2016; Kalisch et al., 2019; Nelson, McGorry, Wichers, Wigman, & Hartmann, 2017; van de Leemput et al., 2014; Wichers, Schreuder, Goekoop, & Groen, 2019; Wichers et al., 2015). The reason is that its behavior maps well on many phenomena observed in mental disorders: bistable systems describe variables that have two stable states, which can be interpreted as different psychological states such as “healthy” or “unhealthy” (e.g., depressed). The stability landscape reflecting the dynamics of the system determines how easy it is to transition from one state to the other, and thereby offers a possible formalization of

^{*}Both authors contributed equally to this work.

[†]jonashaslbeck@gmail.com | www.jonashaslbeck.com

[‡]ryanoisin@gmail.com | <https://ryanoisin.github.io>

properties of the mental disorder, such as vulnerability or resilience to developing it (Scheffer et al., 2018). Bistable systems can also show sudden transitions from one state to another, thereby mapping well on, for example, bipolar disorder or the phenomenon of sudden gains and losses in psychotherapy (Lutz et al., 2013; Stiles et al., 2003).

In parallel, the realization that inferences from between-subjects data to within-person data are only possible under stringent assumptions (Hamaker, 2012; Molenaar, 2004) together with the increasing availability of psychological time series collected from mobile devices has led to a surge in studies aiming at recovering the within-person dynamics associated with mental disorders (e.g., Bringmann et al., 2013; Fisher, Reeves, Lawyer, Medaglia, & Rubel, 2017; Groen et al., 2019; Pe et al., 2015; Snippe et al., 2017). This is an exciting development, because within-person time series potentially allow one to recover bistable systems and other dynamical systems from empirical data. This would be a major step forward for studying mental disorders as complex systems, because so far these models were only used as stylized toy models.

However, so far there has been no systematic investigation into to what extent dynamical systems models can in fact be recovered from psychological time series. To investigate this question for a given dynamical system and data analytic method, it has to be broken down into two parts: The first question is whether the method at hand can recover (some aspect of) a dynamical system *in principle*, that is, with “ideal” data (long time series, extremely high sampling frequency). If this is the case, the second question is whether the method also works with realistic data (shorter time series, much lower sampling frequency). In the present paper, we investigate both questions for a bistable dynamical system and a selection of the most popular (e.g., the Vector Autoregressive (VAR) model; Hamilton, 1994) and some more advanced (e.g., differential equation estimation; Boker, Deboeck, Edler, & Keel, 2010) methods. Specifically, we use a basic bistable dynamical system for emotion dynamics to simulate both an ideal time series with extremely high sampling frequency (measurement every six seconds) and a more realistic time series with a sampling frequency common for Experience Sampling Method (ESM) studies (measurement every 90 minutes). Using these time series, we evaluate how useful each method is for recovering bistable dynamical systems *in principle*, and how useful it can be *in practice* when analyzing realistic ESM time series.

We will show that the popular VAR model (and the Gaussian Graphical Model fitted on its residuals; Epskamp, Waldorp, Möttus, & Borsboom, 2018) is in principle unable to recover the global dynamics (e.g., location and variance of stable states, frequency of transitions) and succeeds only in recovering some of the microdynamics (moment-to-moment interactions) of the true bistable system. However, descriptive statistics, data visualization and more flexible statistical models are able to capture the global dynamics. The only method that recovered the complete bistable system is an iterative model building procedure that directly estimates the system of differential equations (DEs). Reducing the sampling frequency from every six seconds to every 90 minutes affects the considered methods differently: the VAR model and its extensions no longer recover any microdynamics, and the DE-estimation procedure fails. However, descriptives, data visualization and appropriate statistical models still recover the global dynamics. These results raise two fundamental issues involved in studying mental disorders from a complex systems perspective: first, it is generally unclear what to conclude from a statistical model about an underlying complex systems model; and second, if the sampling frequency is too low, it is impossible to recover microdynamics. In response to these findings, we outline a different research strategy to arrive at dynamical systems models for mental disorders: Proposing initial formal models which can subsequently be scrutinized and developed by deriving data implications that can be tested empirically. We will show that in this process many of the presented methods are instrumental to testing predictions of the formal model and thereby triangulating the formal model that captures the true dynamical system best.

Our paper is structured as follows. In Section 2 we introduce a simple bistable dynamical system for emotion dynamics, discuss its dynamics and characteristics, and describe how we generate the ideal and the more realistic time series from it. We use the ideal data (measurement every six seconds) in Section 3 to evaluate for each method to which extent it can recover a bistable dynamical system. Next, in Section 4 we evaluate the same methods but using the time series with a sampling frequency that matches typical ESM studies (measurement every 90min). Finally, in Section 5 we discuss the implications of our results for the framework of empirically studying mental disorders from a complex systems perspectives, and outline a new research strategy based on formal modeling, which avoids shortcomings of a purely data analytic approach.

2 Bistable Emotion System as Data Generating Model

In this section we present a bistable dynamical system and describe its dynamics (Section 2.1), show how we generate data from this system (Section 2.2) and discuss its qualitative characteristics (Section 2.3).

2.1 Model Specification

Bistable dynamical systems are typically formalized within the framework of differential equations (e.g., [Hirsch, Smale, & Devaney, 2012](#); [Strogatz, 2015](#)) and we therefore also choose this framework. Our goal is to provide an accessible first investigation of how well bistable systems can be recovered from psychological time series and therefore use the one of the simplest multivariate bistable systems. Specifically, we choose a system with four variables that is a generalization of the classic Lotka Volterra model for competing species (e.g., [Freedman, 1980](#)) to four variables; a similar model was used by [van de Leemput et al. \(2014\)](#) who interpreted the four variables as positive and negative emotion variables, an interpretation we also adopt here. In an appropriate parameter regime, this model exhibits two stable states: one in which positive emotions are high and negative emotions are low (the “healthy” state); and one in which the positive emotions are low and the negative emotions are high (the “unhealthy” state).

Note that different types of (bistable) dynamical systems will differ in how difficult they are to recover with a given method and type of time series, and much research is needed to map out the space of dynamical systems model classes, data analytic methods and types of time series. However, the intuition we rely on in the present paper is that if there are fundamental problems in recovering one of the simplest multivariate bistable systems, then these problems will be at least equally severe when recovering a more complicated bistable dynamical system.

The bistable system we use throughout this paper consists of two emotions with positive valence (Cheerful (x_1) and Content (x_2)) and two emotions with negative valence (Anxious (x_3) and Sad (x_4)). The dynamics of the system is defined by the stochastic differential equations

$$\frac{dx_i}{dt} = r_i x_i + \sum_{j=1}^4 C_{i,j} x_j x_i + a_i + \epsilon_i, \quad (1)$$

where r_i can be thought of as the main effect of an emotion on itself over time, that is, the effect of x_i on its own rate of change. This parameter is set to 1 for positive emotions, and will be varied between $r_3, r_4 \in [0.9, 1.1]$ for negative emotions. We interpret the variations in r_3, r_4 as being related to stress: higher stress means that the effects of a high degree of negative emotion stays in the system longer. The matrix C represents the dependencies between emotions in the form of interaction effects

$$C = \begin{bmatrix} -0.2 & 0.04 & -0.2 & -0.2 \\ 0.04 & -0.2 & -0.2 & -0.2 \\ -0.2 & -0.2 & -0.2 & 0.04 \\ -0.2 & -0.2 & 0.04 & -0.2 \end{bmatrix}.$$

The interactions in the matrix C show that emotions of the same valence reinforce each other, while emotions of different valence suppress each other. For example, $C_{1,2} = 0.04$ indicates that the rate of change of x_1 (Cheerful) depends on the product of x_1 and x_2 (Content) weighted by 0.04. Similarly, $C_{1,3} = -0.2$ indicates that the rate of change of x_1 depends on the product of x_1 and x_3 (Anxious) weighted by -0.2 . The diagonal elements are quadratic effects: for example, $C_{1,1} = -0.2$ indicates that the rate of change of x_1 depends on the product $x_1 x_1 = x_1^2$. Note that the matrix C is symmetric because we aim at choosing the simplest possible bistable system. In general, however, it does not have to be symmetric.

We interpret 0 as the absence of positive/negative emotion, and therefore do not allow emotions to become negative. We ensure this with high probability by setting the constant $a_i = 1.6$ for all i . The Gaussian noise term ϵ_i has a mean of zero and a fixed standard deviation σ and represents short-term fluctuations in emotions due to the environment the system interacts with. Note that we used the same parameterization as [van de Leemput et al. \(2014\)](#), except that in our model we use an additive noise term instead of a multiplicative noise term for simplicity and set all $a_i = 1.6$.

Figure 1: The dynamics of the bistable system we will use as the data generating model throughout the paper. Panel (a) shows the fixed points of the deterministic part of the model as a function of stress, operationalized by the rate of change of the negative emotions. Solid lines indicate stable fixed points and dashed lines indicate unstable fixed points. Panels (b), (c) and (d) show the vector fields of the system for the stress values $r_3, r_4 = 0.9, 1$ and 1.1 . Solid points indicate stable fixed points and empty points indicate unstable fixed points. The solid lines indicate the values at which derivative of positive emotion (orange) and negative emotion (light blue) is equal to zero. At the points at which the two lines meet, both derivatives are equal to zero and the system remains in this (stable) state.

Due to the symmetries in C , r and a , emotions with the same valence are exchangeable. We can therefore describe the dynamics of the 4-dimensional system using a 2-dimensional system consisting of one dimension for positive emotions and one dimension for negative emotions (for details see Appendix A). Figure 1 illustrates the dynamics of the deterministic part (i.e., with $\epsilon_i = 0$) of this model: Panel (a) displays the stable (solid lines) and unstable (dashed lines) fixed points for positive (green) and negative (red) emotions, as a function of stress. For example, for a low stress level of 0.9 there is only a single fixed point: the positive emotions (PE) have the value 5.28 and the negative emotions (NE) have the value 1.15. We therefore also refer to this fixed point as the healthy state. If the stress level remains unchanged, the system will always end up at this fixed point, no matter how one chooses the starting values. This dynamic is illustrated in the corresponding vector field in panel (b). The arrows depict the partial derivatives with respect to the two emotions and therefore describe the linearized dynamics at a given point in the 2-d space. The vector field shows us that whichever initial values we choose, the system will always end up at the fixed point at (PE = 5.28, NE = 1.15). Thus, the system with stress = 0.9 describes a person whose emotions can be changed by external influences, but eventually always returns to the healthy state of having strong positive emotions and weak negative emotions. The solid lines in panel (b) indicate the values of positive and negative emotions for which the two differential equations are zero. At the intersections of those lines both differential equations are equal to zero, which means that the system does not change anymore, which is the definition of a fixed point.

Panel (a) of Figure 1 shows that when increasing stress from 0.9 until around 0.95, the stable fixed point changes *quantitatively*: the value of positive emotion value decreases, and the value of negative emotion value increases. However, from around 0.95 on the dynamics of the system change *qualitatively*: the system now has three fixed points. For example, at stress = 1, the fixed points are (PE = 4.89, NE = 1.36), (PE = 2.80, NE = 2.80), and (PE = 1.36, NE = 4.89). The first fixed point is the stable healthy fixed point we also observed for values smaller than 0.9. The second fixed point is an unstable fixed point. Specifically, it is a saddle point, because the arrows in the vector field flow towards this fixed point in one direction, but flow away in the other direction (Strogatz, 2015). The third fixed point is again stable, however, now negative emotions have a high value and positive emotions have a low value. We could call this fixed point the unhealthy fixed point.

The presence of these three fixed points means that, if the system is initialized anywhere except on the diagonal, the system will end up at one of the two stable fixed points. This behavior is illustrated in panel (c), which shows the vector field of the system for stress = 1. We see that eventually all arrows point away from the unstable fixed point at (PE = 2.80, NE = 2.80) and towards one of the two stable fixed points. Thus, the system will never converge to this point except if it is initialized exactly on the diagonal. For all other starting values, the system will converge to one of the two stable fixed points. For the particular case of stress = 1, starting values above the diagonal line will converge to the unhealthy fixed point (PE = 1.36, NE = 4.89), whereas starting values below the diagonal line will converge to the healthy fixed point (PE = 4.89, NE = 1.36). This system describes a person that starts out in the healthy (unhealthy) state, and always returns to the healthy (unhealthy) state after small outside influences. However, a large influence can push the person into the unhealthy (healthy) state, and now the person remains there until a large enough influence pushes her back into the healthy (unhealthy) state.

When increasing stress further until around 1.06, we observe again a quantitative change of the three fixed points: the negative emotions go up, and the positive emotions go down. However, from around 1.06 on the system changes again qualitatively. It now again exhibits only one fixed point, which is now the unhealthy fixed point. Thus, when stress is larger than around 1.06, the system will always converge to the unhealthy fixed point. This behavior is illustrated in panel (e), which depicts the vector field for the system with stress = 1.1. We see that there is only a single fixed point at (PE = 1.03, NE = 5.98) and the arrows show that the system will always converge to this point. This system describes a person that will always return to the unhealthy state, no matter how large of an outside influence is applied.

So far, we only discussed the deterministic part of the model, that is, our model with noise set to zero (i.e., with $\epsilon_i = 0$). Introducing noise changes the dynamics of the system, and how exactly it changes depends on the stress level. For low stress (below 0.95), the system will fluctuate around the healthy fixed point. For high stress (above 1.06), the system will fluctuate around the unhealthy fixed point. The interesting behavior is observed for stress values between 0.95 and 1.06: then, the system will fluctuate around one of the two fixed points, but occasionally the noise will be large enough to push the system to the other fixed point. The frequency of switching is a function of the distance between the two fixed points, the vector field between the two fixed points, and the variance of the Gaussian noise process ϵ_i . If the variance is low, the probability of a noise draw that is large enough to “push” the system to the other fixed point is small, and consequently the frequency of switching is low. In contrast, if the variance is high, the probability of a large enough noise draw to switch to the other fixed point is high, and consequently the switching frequency is high.

2.2 Generating Time Series from Bistable System

In the previous section we have shown that our dynamical system is bistable for stress values $(r_3, r_4) \in [0.96, 1.06]$. In the remainder of this manuscript we keep stress constant at stress = 1, and therefore study the bistable dynamical system with the dynamics displayed in panel (c) of Figure 1 and the fixed points described above. Apart from stress we chose all parameters as indicated in the previous section.

To obtain a plausible switching frequency for emotion dynamics we set the standard deviation of the Gaussian noise term $\sigma = 4.5$. Note that a system can be bistable, but the outside influences (the noise term) are so weak that the system switches very infrequently or not at all. In such cases the bistable system is more difficult (infrequent) or impossible (no switches) to recover. Thus, our

choice of σ represents an ideal situation, and all presented methods will perform worse with a lower switching frequency.

In the remainder of this section we describe how we generated the two time series that we will use throughout the paper: an “ideal” time series with an extremely high sampling frequency of 1 measurement every six seconds (Section 2.2.1); and a more realistic time series with measurements every 90 minutes, a sampling frequency typical for ESM studies (Section 2.2.2).

2.2.1 Ideal Time Series

We generated data by computing the numerical solution to the model in Section 2.1 with $\text{stress} = 1$ on the interval $[0, 20160]$, using Euler’s method (e.g., Atkinson, 2008) with a step size of 0.01. We interpret a time step of 1 as one minute, and therefore the time series spans two weeks ($60 \times 24 \times 14 = 20160$). We obtain a time series by sampling the numerical solution obtained via Euler’s method 10 times per minute (or every six seconds). We therefore obtain the ideal time series with $20160 \times 10 = 201600$ measurements, which appears to switch between fixed points around 17 times¹. Figure 2 displays this time series:

Figure 2: Panel (a) shows the ideal time series of the four emotion variables Cheerful, Content, Anxious and Sad. We see that the system switches 17 times between healthy and unhealthy state. Panel (b) displays the twelfth switch, which is a transition from the unhealthy to the healthy state, which occurs on day 9.

We choose this unrealistically ideal time series (two weeks, one measurement every six seconds, continuous response scale, no measurement error or missing values) with 201600 measurements to be able to study the usefulness of different data analytic methods *in principle*. That is, we study the usefulness of all methods on the *population level*, which is the situation in which we have infinitely many observations and sampling variation does not exist. With 201600 observations, in this setting we approximate “infinitely many” for all practical purposes.

Note that we would not be able to investigate how well different methods perform in principle, if we made the time series more realistic by choosing a shorter time interval or sampling it with a lower sampling frequency: in such a case we would not know whether a method cannot recover (an aspect of) the bistable system for fundamental reasons, or because the time series is too short or the sampling frequency too low. We therefore first study all methods with the ideal time series in order to identify their fundamental limitations. In the second part of the paper, we make the time series more realistic by taking measurements at a sampling frequency that is typical for ESM studies. This will allow us to investigate the impact of sampling frequency on all methods. In the following section we describe how we generate this ESM time series.

¹The code to generate data and reproduce all analyses and results shown in this paper can be found at <https://github.com/jmbh/RecoveringBistableSystems/>.

2.2.2 Experience Sampling (ESM) Time Series

Clearly, the ideal time series is very different from time series data sets obtained from typical ESM studies. The perhaps two most important differences between the ideal time series and realistic time series are the measurement scale and the sampling frequency. With respect to the measurement scale, most ESM studies do not use a continuous response but, for example, a 7-point Likert scale. Regarding the sampling frequency, ESM studies investigating psychological variables typically do not measure more frequently than every 90 minutes (e.g., [Bringmann et al., 2013](#); [Fisher et al., 2017](#); [Groen et al., 2019](#); [Pe et al., 2015](#); [Snippe et al., 2017](#)). Thus, ESM time series have a much lower sampling frequency (every 90 minutes) than the ideal time series used above (every six seconds). To be able to explain possible drops in performance of certain methods when making the time series more realistic, we are only allowed to change one aspect of the time series. While the measurement scale can possibly be made near-continuous, there are certainly hard limits on how many times one can notify a person each day with an ESM questionnaire. We therefore consider the sampling frequency the more fundamental constraint in realistic data, and thus make it the focus of our investigation in Section 4.

Taking a measurement every 90 minutes in the two weeks of the original ideal data leads to 224 measurements. This would mean that we would compare the "ideal" time series with 201600 measurements which essentially implies the absence of sampling variation to an ESM time series with 224 measurement which implies a lot of sampling variation. Thus, any comparison would be confounded by the difference in the number of measurement points (i.e., sample size). To avoid this confound, we increase the measurement interval of the ESM time series to 1800 weeks, which ensures that the new ESM time series has exactly the same sample size as the ideal data ($\frac{224}{2} \times 1800 = 201600$). Thereby, we provide that any drop in performance is a function of the lowered sampling frequency and cannot be explained by lower sample size (and higher sampling variation). Note that studying the performance of methods as a function of sample size (sampling variation) is of paramount importance to evaluate how useful a given method is in a realistic application. However, here we study the more fundamental question of the impact of reducing the sampling frequency to a level that is typical for ESM studies. We do this because if we find that a method is ill-suited to recover (some aspect of) a bistable system with a realistic sampling frequency on the population level (i.e., with infinite sample size), then it does not make sense to investigate the performance of the method in the less ideal scenario with realistic (small) sample sizes.

So far, we only discussed that we sample every 90 minutes. However, to emulate ESM measurements, we also need to formalize how exactly ESM questions measure the four emotion variables. This is far from trivial: questions in some ESM studies refer to the very moment of measurement and are phrased along the lines of "How cheerful do you feel right now?". Such measurements could be formalized by defining the measurement as the set of current values of the system (a "snapshot" of the system) at the measurement time. In contrast, other ESM studies refer to the time period since the last measurement. A question of this type could be phrased "How cheerful did you feel in the time since the last notification?". Such measurements could be formalized by defining the measurement as the average values of the system since the last measurement. However, many other measurement functions are also possible. In this paper we analyze the first type of ESM question, because its measurement function is the simplest. However, we also performed all analyses with the second kind of ESM question, and all our main conclusions also hold in this situation.

Figure 3 displays the two week long original time series (see also Figure 2 panel (a)) next to the ESM time series which was obtained by taking "snapshots" of the process at 90 min intervals:

Figure 3: Panel (a) shows the original time series that was already shown in panel (a) of Figure 2. Panel (b) shows the ESM time series which was obtained by taking snapshots every 90 minutes in the series. Note that the ESM time series we analyze in Section 4 is much longer (1800 weeks) than the 14 day ESM time series shown here.

The ESM time series in panel (b) appears less dense, which is what we would expect since it contains only 1/900 of the time points of the ideal time series in panel (a). However, we see that the system is still bistable and that the location of and variance around the fixed points is largely the same. In Section 4 we will use this emulated “snapshot” ESM time series to try to recover the true bistable system using the same array of methods as in Section 3, in which we analyze the ideal time series.

2.3 Qualitative Characteristics of the Model

In this section we discuss the key qualitative characteristics of the bistable system introduced in the previous section. We list these characteristics because most considered methods are models that are misspecified (i.e., they do not contain the true system as a special case). In such a situation one can only hope to recover some characteristics of the true system, and we therefore evaluate how well a method recovers the bistable system based on how well it recovers the following seven characteristics:

Global dynamics

1. Bistability (two stable fixed points)
2. Position of fixed points
3. Variability around fixed points
4. Frequency of transitions

Microdynamics

5. Suppressing effects between valences, reinforcing effects within valences
6. Relative size of suppressing/reinforcing effects
7. All parameters are independent of time and independent of variables outside the model

The first four characteristics describe the global dynamics of the dynamical system. The first is bistability, which means that the data generating mechanism exhibits two stable fixed points. This is the case for the data generating mechanism with stress set to 1, which we use to generate data from and aim to recover throughout the paper (see Figure 2, panel (a) and (c)). The second characteristic is the position of the fixed points, which are at ($PE = 4.89$, $NE = 1.36$) for the

healthy fixed point, and ($PE = 1.36$, $NE = 4.89$) for the unhealthy fixed point. Third, we consider the variability around the different fixed points. Figure 2 shows that, for both fixed points, the variability of the emotions with lower values is smaller than the variability of the emotions with larger values. The fourth characteristic is the frequency of transitions between the area around the healthy fixed point and the area around the unhealthy fixed point. In the time series shown in Figure 2 we see that the system switches around 17 times.

The remaining three characteristics describe the microdynamics of the dynamical system. The fifth characteristic is that emotions of the same valence reinforce each other, while emotions of different valence suppress each other. The sixth characteristic is the fact that the size (absolute value) of the reinforcing effects (0.04) are smaller than the suppressing effects (0.2). The last characteristic is that all parameters in the system of differential equations are independent of time and independent of variables outside the model.

3 Recovering the Bistable System from Ideal Data

In this section we analyze the ideal time series to evaluate how well different methods recover the data generating bistable system. The methods considered here primarily consist of the most popular models used in analyzing time series in clinical psychology and psychiatry, and some extensions thereof. In all but one instance, this entails the estimation of misspecified models, that is, models which do not contain the true bistable system as a special case. Thus we will focus our investigation on whether these models allow one to recover some of the characteristics of the true model, as outlined in Section 2.3.

We analyse each method in order of increasing complexity, moving from methods which may be helpful in recovering global characteristics alone to methods which are typically used with the aim of characterising the microdynamic structure. We begin by inspecting the time series using descriptive statistics (Section 3.1); in Section 3.2 we characterize the switching behaviour in the system using a mean-switching Hidden Markov Model (Hamaker, Grasman, & Kamphuis, 2016; Hosenfeld et al., 2015). Next, in Section 3.3 we analyze the multivariate lag-0 (same time point) relationships using correlations and partial correlations with the popular Gaussian Graphical Model (GGM) (Epskamp et al., 2018). In Section 3.4 we use the most popular approach to modeling microdynamics in experience sampling settings, the lag-1 Vector Auto-Regressive (VAR(1)) model (e.g. Bringmann et al., 2013; Fisher et al., 2017; Groen et al., 2019; Pe et al., 2015). Next, we evaluate the Threshold VAR model, an extension of the VAR model that allows the modeling of state-switching behaviour using time-varying parameters (Hamaker & Grasman, 2012; Hamaker, Zhang, & van der Maas, 2009) (Section 3.5). While all models so far are misspecified, we include one final method that is capable of recovering the full bistable system: a two-step model building approach based on direct estimation of differential equations from data, following the dynamic systems modeling approach of Boker et al. (2010) and Chow (2019).

3.1 Descriptive Statistics

To get a rough overview of the behavior of the system, we inspect the time series plot of all four emotion variables shown above in Figure 2 panel (a). We see that at almost every time point the two positive emotion variables have high values around 5, and the two negative emotion variables have small values around 1, or the other way around. At the remaining time points, the variables seem to transition between those two states (see panel (b) in Figure 2). In addition, we see that the variables switch between states 17 times.

We can extract a considerable amount of information from simply inspecting the time series plot. There seem to be two stable states (fixed points), one in which positive emotions are high and negative emotions are low, and one in which the reverse is true. Further, we see that the variance is higher for the emotion with higher values, and we saw that the system switches around 17 times in the two week window. To get a more direct picture of the distribution of variables and possible fixed points we show the histograms for each variable in Figure 4:

Figure 4: The histograms of the emotion intensity of the four modeled emotions Cheerful, Content, Anxious and Sad, for the ideal data.

We see that at most time points in the time series, each emotion either takes on values around 1 or around 5. This is what we would expect from inspecting the time series plot, however, the histograms give a more precise picture of the distributions and allow one to guess possible fixed points with greater precision. For instance, we could separate the two distributions (using a fixed threshold, or clustering algorithm) and take their means as estimates for the fixed points.

While eyeballing the data should be the first step in any time series analysis, the conclusions are subjective and do not allow us to quantify how certain we are about bistability and the switching frequency. We can quantify the observation that there are two states and that the system is switching between them by fitting a Hidden Markov Model (HMM) (e.g., [Rabiner, 1989](#)) to the data, which we will do in the following section. Such quantification is especially valuable in more realistic situations, in which the two states are probably harder to separate than in our ideal simulated data.

3.2 Hidden Markov Model

In this section we fit a mean-switching Hidden Markov Model (HMM) in order to scrutinize the intuition that the system switches between two states and to quantify the switching frequency. The HMM models the observed data as consisting of K latent states or components, characterized by K multivariate Gaussian distributions, which may differ in their means μ_k and variances σ_k .² Each observation over time is drawn from one or other of these distributions, and the switching between these states is governed by a matrix of transition probabilities A . For more details about this model see [Appendix B.1](#).

Here we choose $K = 2$ components, and fit this model to our time series using the R-package *depmixS4* ([Visser & Speekenbrink, 2010](#)), obtaining the following parameter estimates

$$\hat{\mu}_1 = \begin{pmatrix} 1.47 \\ 1.46 \\ 4.71 \\ 4.71 \end{pmatrix}, \quad \hat{\sigma}_1 = \begin{pmatrix} 0.41 \\ 0.40 \\ 0.63 \\ 0.62 \end{pmatrix}, \quad \hat{\mu}_2 = \begin{pmatrix} 4.75 \\ 4.76 \\ 1.45 \\ 1.45 \end{pmatrix}, \quad \hat{\sigma}_2 = \begin{pmatrix} 0.63 \\ 0.62 \\ 0.40 \\ 0.40 \end{pmatrix}, \quad \hat{A} = \begin{matrix} & C_1 & C_2 \\ \begin{matrix} C_1 \\ C_2 \end{matrix} & \begin{pmatrix} 0.9996 & 0.0004 \\ 0.0004 & 0.9996 \end{pmatrix} \end{matrix}.$$

We can see from the estimate $\hat{\mu}_1$ that in state 1 the means of positive emotions are low, and the means of negative emotions are high. We can therefore identify state 1 as the unhealthy state. We also see that the standard deviations of positive emotions are lower than for negative emotions in the unhealthy state which is what we already observed in the time series plot in [Figure 2](#). Similarly, state 2 can be identified as the healthy state, with high means and standard deviations for positive emotions, and low for negative emotions. The transition matrix A indicates the probabilities of switching between states. We see that there is a very high probability for remaining in the same state ($A_{11} = A_{22} = 0.9996$), and a correspondingly low probability to switch states ($A_{12} = A_{21} = 0.0004$). This is what we would expect, because we take one measurement every six seconds, but the system changes states only a couple of times within the two week window. Multiplying the number of time points of the time series with the switching probability we obtain

²In principle, the distributions may also differ with respect to their covariances, but in this analysis, we set all covariances to zero due to limitations of the software package used in estimation.

$201600 \times .0004 \approx 81$ switches, which is in the same order of magnitude of the eyeballed number of switches (17) reported in Section 2.2.

In addition to obtaining estimates of means and standard deviations of the two fixed points and the transition matrix, the HMM allows to predict the most likely state for each time point. We show the predicted states for the entire time series in Figure 5:

Figure 5: Time series of the four emotion variables, also shown in panel (a) of Figure 2, with background color indicating whether a given time point is assigned to the first or second component of the mean-switching HMM.

When inspecting the predicted states visually, it seems that the HMM captured the switches well. Next to the larger blocks in which the system stays at the same fixed point, it also identifies switches in which the system switches back and forth within only a few time points. These switches might have been missed when inspecting the time series visually alone.

Taking all results together, which characteristics of the bistable system did we recover with the HMM (see list in Section 2.3)? We got an estimate of the location (characteristic 2) and variance (characteristic 3) around two fixed points, which are very close to the healthy and unhealthy fixed points in the true bistable system. We also quantified the frequency of transitions in the transition matrix A . Since the transition frequency (characteristic 4) is not explicit in the true bistable model, there is no clear way to evaluate this estimate. However, the number of predicted transitions (81) is at least in the same ballpark as the number of transitions eyeballed from the entire time series (at least 17). Note that while a bistable HMM seems to fit the data well, we provided $K = 2$ as an *input* to the model, and therefore bistability (characteristic 1) cannot be considered a characteristic we recovered with this model. Instead of fixing a particular K , an optimal K can be obtained via model selection. However, in Appendix B.2 we show that at least the standard approach to selecting K in mean-switching HMMs performs poorly since the data was not generated from a mean-switching HMM.

One additional way to visualize or ascertain how much of the true systems behaviour a given model is able to capture is by generating new data from the estimated model parameters. In Figure 16 in Appendix C.1 we generate a two week time series from the estimated mean-switching HMM and compare it to the original time series. We find that the data generated from the HMM is similar to the original data, except for two features: First, the system tends to switch between states somewhat more frequently, and second, there are no observations on the transitions between states.

The remaining three characteristics (5-7) are about the microdynamics of the true bistable system, that is, about how the components are related to each other. Clearly, the mean-switching HMM we used here cannot elucidate these characteristics since it does not model any dependencies between the four emotion variables. In the following sections we fit models that include such dependencies.

3.3 Lag-0 Relationships / Gaussian distribution

In this section we analyse the relationships between variables at the same time point. Figure 6 panel (a) displays the relationship between Content and Cheerful, two emotions of the same valence. We see that the observations cluster around two points, one close to (1, 1) with smaller variance, and one close to (5, 5) with larger variance. The red line indicates the best fitting regression line (correlation 0.98).

Figure 6: Panel (a) shows the relationship between Content and Cheerful, two emotions with the same valence, at the same time point. The red line indicates the best fitting regression model. Similarly, panel (b) shows the relationship between Anxious and Content, two emotions with different valence. Panel (c) displays the correlation matrix as a network, and panel (d) displays the partial correlation matrix as a network.

Panel (b) displays the relationship between Cheerful and Anxious, two emotions of different valence. We see that the observations cluster around two points, one close to (1, 5) and the other one close to (5, 1). The red line indicates the best fitting regression line (correlation $\rho = -0.97$).

Panel (c) displays the correlation network for all four emotion variables. As we have already seen in panel (a) and (b) there is a positive correlation ($\rho = 0.98$) between Content and Cheerful, and a negative correlation ($\rho = -0.97$) between Cheerful and Anxious. Due to the symmetry in the true bistable system, all correlations between emotions with the same valence are equal to $\rho = 0.98$ and all correlations between emotions with different valences are equal to $\rho = -0.97$. Panel (d) shows the partial correlation network (i.e., GGM). We see that the partial correlations between emotions with the same valence are equal to $\theta = 0.51$, and the partial correlations between emotions with different valences is equal to $\theta = -0.24$ or $\theta = -0.25$.

What can we learn from these results about the underlying bistable system? From inspecting the pairwise relationships of emotions with same and different valence in panels (a) and (b) one could guess the location and variance of possible fixed points, similarly to inspecting the histograms

in Section 3.1. However, the 2-dimensional representation offers additional information about the stability landscape, for example the shape around the fixed points and the most likely paths to transition between them. When interpreting the correlations in panel (b) as “contemporaneous” relationships, we would conclude that there are strong positive linear relationships between emotions with the same valence, and similarly strong negative linear relationships between variables with different valences at a relatively short time scale. The partial correlations in (d) are smaller than the correlations, which is what one would expect since all correlations are high.

Using our knowledge about the true bistable system, which characteristics did we correctly recover? From inspecting the scatter plots in panels (a) and (b) one sees that most observations fall in one of two clusters indicating bistability (characteristic 1). Also, one can obtain rough estimates of the position of the fixed points (characteristic 2) and sees that the variances around the fixed points is different (characteristic 3). Note that the shape of the scatter plot in panel (b) is determined by the vector field in Figure 1 (c). The two clusters are exactly at the location of the two fixed points, and the observations between the clusters are both due to variance around the fixed points and switches between fixed points.

From the correlation and partial correlation network, we correctly find that there are reinforcing effects within valences, and suppressing effects between valences (characteristic 5). However, the correlation network suggests that their relative size is equal, and the partial correlation network suggests that the reinforcing effects are stronger. In the true bistable system, however, the suppressing effects between valences are larger than the reinforcing effects within valences. Thus, judging the relative size of suppressing/reinforcing effects within/between valences from (partial) correlation would lead to incorrect conclusions.

In sum, inspecting scatter plots of pairwise relationships indicated bistability, and allowed us to obtain a rough estimate of the location of and variances around the fixed points and. The scatter plots also allowed one to get a projection of the stability landscape on two dimensions and thereby provide more information than histograms. While inspecting the scatter plots allows one to recover global dynamics of the true bistable system, one cannot infer the coupling between the emotion variables in the true bistable system from (partial) correlations. This is not too surprising since the Gaussian distribution is very restrictive in that it only models pairwise linear relationships (opposed to e.g., 3-way, 4-way, etc. interactions). In addition, it does not model any dependencies across time, which are the types of dependencies that constitute the microdynamics (characteristics 5-7) of the true model. In the next section, we inspect those dependencies across time and fit a Vector Autoregressive (VAR) model to the data, which captures temporal linear dependencies.

3.4 Lag-1 Relationships / VAR Model

In this section we aim to characterize the microdynamics between the four emotion variables by modeling the lagged linear relationships between them. Panels (a) and (b) of Figure 7 show the *marginal* relationship of Content at time t with Cheerful at the previous time point $t - 1$, and Anxious at time t with Content at time $t - 1$, respectively. Although the data is generated from a dynamic model, the marginal lagged relationships look very similar to the contemporaneous relationships shown in Figure 6: when averaging over all other variables, the lagged relationships within-valence are positive, and between-valence are negative. The reason is that the system largely stays around the two stable fixed points, and relative to the length of the time series, switches are infrequent. As such, these marginal relationships are largely driven by the relative location of the two fixed points.

Figure 7: Panel (a) shows the relationship between Content and Cheerful, two emotions with the same valence, spaced one time point apart (at a lag of one). The red line indicates the best fitting regression model. Similarly, panel (b) shows the relationship between Anxious and Content, two emotions with different valence, at a lag of one. Panel (c) displays the matrix of lagged regression parameters, estimated from a VAR(1) model, as a network, and panel (d) displays the partial correlation matrix of the residuals of the VAR(1) model as a network. This latter network is often referred to as the contemporaneous network.

We can gain further insight into the dependencies between variables in our model by examining the *conditional* lagged relationships between pairs of emotions, that is, when keeping the other emotion variables at the previous time point(s) fixed. A popular model for such conditional lagged relationships is the first-order vector auto-regressive (VAR(1)) model. The VAR(1) model is one of the simplest multivariate dynamic models which can be fit to repeated measurement data, allowing linear relationships between all pairs variables observed at consecutive measurement occasions t and $t - 1$:

$$X_t = b + \Phi X_{t-1} + e_t \quad (2)$$

where b is a vector of intercepts, Φ is a matrix containing the auto-regressive (Φ_{ii}) and cross-lagged ($\Phi_{ij}, i \neq j$) effects, that is, conditional linear dependencies, and e_t is a vector of normally distributed residuals $e_t \sim \mathcal{N}(0, \Psi)$, which are independent across time, with residual variance-covariance matrix Ψ .

The VAR(1) has been used widely to analyze experience sampling data in psychopathology research, particularly in the form of dynamic network analysis, wherein the Φ and Ψ matrices are used to construct directed and undirected network structures, respectively (e.g., [Bringmann et al., 2013](#); [Epskamp et al., 2018](#); [Pe et al., 2015](#)). The VAR(1) model describes a system which fluctuates around a *single* stable fixed point: Stochastic input in the form of a residual term pushes the system away from this fixed point, and the system returns to the fixed point with an exponential decay ([Hamilton, 1994](#)). The location of the stable fixed point is given by the mean vector μ , a function of the intercepts and the lagged relationships $\mu = (I - \Phi)^{-1}b$, where I is the identity matrix.

Panel (c) of Figure 7 displays the network of estimated lagged regression coefficients ($\hat{\Phi}$) between Cheerful, Content, Anxious and Sad. We can see that the auto-regressive parameters are large and positive for all four variables ($\hat{\Phi}_{ii} = .91$). Furthermore, there are positive cross-lagged relationships between variables of the same valence ($\hat{\Phi}_{12} = \hat{\Phi}_{21} = \hat{\Phi}_{43} = \hat{\Phi}_{34} = .05$) and weaker, negative cross-lagged effects between variables of opposite valence ($\hat{\Phi}_{13} = \hat{\Phi}_{31} = \dots = -.02$). All within-valence effects, and all between-valence effects, are of roughly equal magnitude, respectively. In panel (d) we show the partial correlations of the residuals (i.e., standardized $\hat{\Psi}^{-1}$), sometimes referred to as the “contemporaneous” network or the residual GGM (Epskamp et al., 2018). Here we see a similar pattern as above: the residuals have negative conditional relationships between-valence, and slightly greater in magnitude positive conditional relationships within-valence. Note that the residual variance in this case is quite low for each variable ($\hat{\Psi}_{ii} \approx 0.0185$ with in-sample explained variance of approximately 99 percent). The estimated fixed point (that is, the mean) is approximately $\hat{\mu}_1 = \hat{\mu}_2 = 3.16$ for Cheerful and Content, and $\hat{\mu}_3 = \hat{\mu}_4 = 3.04$ for Anxious and Sad.

Which characteristics of the bistable system can we recover based on the VAR(1) estimates? The strong auto-regressive effects correctly capture the strong linear auto-effects present in the true system, defined by the r parameters in Equation 2. The lagged regression parameters suggest that there are suppressing effects between valences, and reinforcing effects with valences, capturing characteristic number five of the data generating mechanism (Section 2.3). However, the relative size of the suppressing and reinforcing effects is flipped in the VAR(1): The suppressing effects are in fact larger in absolute value than the reinforcing ones (see Section 2.1).

It is unclear what conclusions we can draw from the weak relationships present in the residual network — as there are no such additional instantaneous relationships present in the data generating system. We assume a-priori when fitting the VAR(1) model that these parameters are independent of (i.e., constant across) time. Finally, the VAR(1) model describes a uni-stable system, precluding us from capturing any characteristics related to bistability. The dynamics implied by the VAR(1) are illustrated by generating new data from the estimated parameters, displayed in Appendix C.2. The estimated location of the single fixed point is not equivalent to either of the two stable fixed points or the unstable fixed point in the true system.

Importantly, we can use our knowledge of the true bistable system to determine how we arrived at these observed parameter estimates. First, the estimated position of the fixed point (given by $\hat{\mu}$) is roughly halfway between the positions of the two stable fixed points, and is approximately equal to the sample mean for each variable, reflecting that the system spends roughly the same amount of time around each of the two stable fixed points. The main counter-intuitive result from the VAR(1) model is that the order of magnitude of the between- and within-valence relationships is different than in the true bistable system. In the true system, these pairwise relationships are non-linear, taking the form of interaction effects, and that the VAR(1) model captures the best linear approximation of these non-linear relationships. As we can see from panel (a) of Figure 7, the linear approximation of the within-valence relationship is largely driven by the strong positive relationship present when both variables take on a high value, e.g., when Cheerful and Content are both near the healthy fixed point. From panel (b) we can see that, in contrast, the linear between-valence relationship of Content on *Anxious* is in fact a mixture of the strong negative effect near the unhealthy fixed point (Content is low, Anxious is high) and the weaker effect near the healthy fixed point (Content is high, Anxious is low). Combined, this results in higher linear within-valence relationships and lower linear between-valence relationships.

Finally, the residual covariances displayed in panel (d) of Figure 7 are produced by a combination of model-misspecification (linear approximation of non-linear relationships) and paths between each process at a shorter time scale than observed, due to the Euler steps used in data generation. We stress here that, even in the current idealized situation, it is not trivial to derive an exact explanation for the residual covariance structure, and so its utility in drawing conclusions about the underlying system should be approached with great caution.

In summary, the VAR(1) model gives us rather limited information regarding the core characteristics of the bistable model. In principle, the VAR(1) model is unable to capture any features which relate to bistability (characteristics 1-4), as one would expect from a model that exhibits only a single fixed point. What is perhaps more surprising is that, while the sign of the lagged relationships (characteristic 5), and their symmetries are captured, their relative ordering (characteristic 6) is not. This observation is critical: while we could expect that the VAR(1) model would not reproduce the global dynamics of the system, even when we have ideal data, the linear relationships in the VAR(1) model also fail to appropriately capture the local microdynamics in

this instance. Fundamentally, this is due to the non-linear relationships which must be present in the underlying system in order to induce bi-stability: In general we would not expect that linear approximations of non-linear effects would preserve the same rank ordering. This observation has potentially major implications for the analysis of dynamic network structures, because many network derivatives such as centrality metrics are strongly dependent on the relative ordering of effects. In the following we will examine if extending the VAR(1) model to allow for bi-stability brings us closer to recovering a more accurate characterization of the data generating bistable system.

3.5 Threshold VAR Model

Regime-switching VAR(1) models extend the VAR(1) model to allow for observations to be drawn from two different conditional distributions for X_t given X_{t-1} , that is, two different *regimes*, described by two different sets of model parameters. These extensions in principle allow us to directly capture a notion of multi-stability, by interpreting the mean vector of each conditional distribution as a separate fixed point. Different extensions allow for different mechanisms by which to model the switch between these regimes.

One popular regime-switching VAR(1) model is the Threshold TVAR(1) model, where the system enters a different regime whenever a threshold value or values τ of an a-priori specified threshold variable z_t is crossed, written

$$\begin{aligned} X_t &= b^{(1)} + \Phi^{(1)} X_{t-1} + e_t^{(1)} && \text{if } z_t \leq \tau \\ X_t &= b^{(2)} + \Phi^{(2)} X_{t-1} + e_t^{(2)} && \text{if } z_t > \tau \end{aligned}$$

for a two-regime model with a single threshold, where the VAR(1) parameters are indexed by regime, with $e_t^{(r)} \sim \mathcal{N}(0, \Psi^{(r)})$, and mean vectors $\mu^{(r)} = (I - \Phi^{(r)})^{-1} b^{(r)}$ (Hamaker, Grasman, & Kamphuis, 2010; Tong & Lim, 1980). The threshold variable z_t may be an exogenous variable, or one of the variables in the VAR model. Here we choose to use Cheerful ($z_t = x_{1,t-1}$) as the thresholding variable. The threshold value τ is a hyper-parameter that is estimated. Here, we estimate the TVAR(1) model using the R-package *tsDyn* (Fabio Di Narzo, Aznarte, & Stigler, 2009), which estimates τ using a grid search which selects the model with minimum summed squared residuals.

Figure 8 displays the main results from the estimated TVAR(1) model, in which the threshold is estimated as $\hat{\tau} = 2.811$. In panel (a) of Figure 8 we show the time-series with shading indicating which observations are below (grey) or above (white) the threshold. We can see that the estimated threshold nicely separates the time series into periods in which the system is in an unhealthy state (based on Cheerful values below the threshold) and a healthy state (Cheerful values above the threshold).

Inspecting the lagged networks for each regime in panels (b) and (c) of Figure 8 we can see that the auto-regressive effects, and the within-valence cross-lagged effects are pretty similar across both regimes. However, the cross-lagged effects between variables of opposite valence are different. In the healthy regime, negative valence emotions have much stronger cross-lagged effects on positive emotions ($\hat{\Phi}_{13}^{(2)} = \hat{\Phi}_{14}^{(2)} = \hat{\Phi}_{23}^{(2)} = \hat{\Phi}_{24}^{(2)} = -.08$), and vice versa for the unhealthy regime ($\hat{\Phi}_{31}^{(1)} = \hat{\Phi}_{41}^{(1)} = \hat{\Phi}_{32}^{(1)} = \hat{\Phi}_{42}^{(1)} = -.08$). Residual partial correlation networks for both regimes are shown in Appendix D, which display a similar pattern to the regular VAR(1) model of weak positive residual partial correlation within-valence and weak negative residual partial correlation between-valence. For the TVAR, however, the residual covariance matrix is not symmetric across regimes: In the healthy regime there is a slightly higher covariance between positive emotions than negative emotions, and vice versa. The estimated means are given as $\hat{\mu}_2 = \{4.74, 4.75, 1.45, 1.46\}$ for the healthy state and $\hat{\mu}_1 = \{1.49, 1.48, 4.69, 4.69\}$ for the unhealthy state. Data generated by the TVAR(1) model estimates is shown in Figure 18 in Appendix C.3. From this figure we can see that most of the global dynamics are well reproduced, although the system contains fewer switches between regimes than we would expect and there are fewer observations on the switches between states compared to the original time series.

Figure 8: Panel (a) shows the two weeks of the time series, with observations shaded in either grey or white as a function of whether $x_{1,t-1}$ is above or below the threshold $\hat{\tau} = 2.811$. Panels (b) and (c) show the estimated VAR(1) parameters as lagged networks in the healthy (white) and unhealthy (grey) regimes respectively.

Which characteristics of the bistable system do we recover on the basis of the TVAR(1) parameter estimates? First, the model picks up a number of characteristics related to the bistability of the system: The estimated mean vectors capture approximately the position of the two stable fixed points (characteristic 2), and the estimated threshold correctly captures the position of the unstable fixed point in the Cheerful dimension. However, note that bistability (characteristic 1) has been specified a-priori and therefore cannot be considered to be recovered by the model. Second, although the simulated data in Figure 18 (Appendix C.3) exhibits less frequent switches between states than we would expect, we can see that the combination of state-dependent lagged parameters and residual variances does reproduce higher variability of positive emotion in the healthy state in comparison to the unhealthy state, and vice versa for negative emotions (characteristic 3). Finally, the lagged regression parameters in each regime correctly capture that there are reinforcing effects within valence, and suppressing effects between valence (characteristic 5).

The result that stands out in this analysis is the asymmetry in lagged regression coefficients across both regimes. This asymmetry would appear to indicate that the parameters relating processes either change over time or are all explicitly a (step) function of the Cheerful variable. This last result is striking because this intuitive interpretation does not correctly characterize the relationship between variables of different valences in the true bistable system. This is because we know that the dependencies in C are invariant over time and fully symmetric. However, the dependencies in C relate to pairwise interaction effects rather than linear dependencies in the VAR(1) model. For example, the relationship between Anxious, denoted x_3 , and the rate of change of Content, $\frac{dx_2}{dt}$, depends both on the value of C_{23} and on the current value of x_2

$$\frac{dx_2}{dt} = r_2 x_2 + (C_{23} \times x_2) x_3 + \dots \quad (3)$$

If we view x_2 as a moderator, we can see that, when x_2 is high, the effect of x_3 on the rate of change, given by $C_{23} \times x_2$, is relatively greater than when x_2 is low. In our system, separating the time-series into two regimes based on a threshold of 2.811 for the Cheerful emotion essentially means we condition on high values of x_1 and x_2 in the healthy regime, and low values in the unhealthy regime. This leads to the relatively stronger linear relationship from negative emotions to positive emotions in the healthy regime, and vice versa in the unhealthy regime. As such, we can see that the asymmetry in lagged relationships over time picked up by the TVAR(1) model is a characteristic of the true bistable system. Notably, however, the mechanism by which this asymmetry occurs is entirely due to non-linear relationships between the observed variables and the similarity of variables that share the same valence, while the TVAR(1) modeller might be tempted to ascribe this entirely to the effect of the level of Cheerful.

To summarize, the TVAR(1) model allows us to recover global dynamics, and it recovers some aspects of the microdynamics. However, we saw that a naive interpretation of the TVAR parameter estimates may easily lead to the incorrect conclusion that there is one time-varying variable which moderates the relationships between all variables. In addition, we provided bistability as an input to the model, and therefore cannot be considered a characteristic recovered from data. In principle one could perform model selection between TVAR(1) models with different numbers of components, however compared to the Mean switching HMM in Section 3.2, the run time for such a model comparison was unfeasible for the large data set used in our paper.

Furthermore, note that the threshold VAR(1) model does remarkably well for this specific system for the following reason: While TVAR(1) models have frequently been discussed in the literature (e.g. De Haan-Rietdijk, Gottman, Bergeman, & Hamaker, 2016; Hamaker et al., 2010, 2009; Warren, 2002) a major limitation of this method is the difficulty in choosing a thresholding variable. In our data generating mechanism, we know there to be an unstable fixed point defined in multivariate space, $x_1 = x_2 = x_3 = x_4 = 2.8$. It just so happens that in this case, almost always when we pass this position in one dimension (e.g., $x_1 > 2.8$) we also do so in all other dimensions (e.g., $x_2 > 2.8, x_3 < 2.8, x_4 < 2.8$). This means that the true mechanism of state-switching behaviour is very well approximated by the univariate mechanism in the TVAR(1) model, for this choice of parameter values. In more general situations, the choice of thresholding variable(s), and number of thresholds, is likely to be less trivial. While the TVAR(1) model does capture that there are suppressing and reinforcing effects between and within valences, it does not capture the relative size of these effects, and it may easily lead to the incorrect conclusion that there is a single time-varying variable which moderates all of the relationships between other variables in the system.

Finally, the TVAR(1) is only one of a variety of different regime-switching dynamic models which could be fitted to the data at hand. Another alternative would be the Markov-Switching (MS-)VAR model (Chow et al., 2018; Hamaker & Grasman, 2012; Hamaker et al., 2010; Hamilton, 1989), a combination of the HMM and VAR models, in which the regime-switching behaviour is determined by a random Markov process operating between latent categorical variables. While this model is more flexible than the Threshold VAR model, we show here the TVAR results for two reasons. First, in this instance the switching behaviour will be less well approximated by the MS-VAR model, leading to even less straightforward conclusions about the data generating process, but otherwise likely highly similar lagged parameter estimates. Second, while recent advances such as the *dynR* package (Ou, Hunter, & Chow, 2019) have made this model easier to estimate, it is still prohibitively difficult and time consuming to fit to data.³

Now that we have shown the capabilities and limitations of the TVAR model in recovering the bistable system, there are a few different avenues we could pursue to further increase our model complexity in the hope of recovering more and more of the features of underlying system. For example, both the TVAR and MS-VAR can be considered special cases of *time-varying parameter* models, that assume the true time-varying model is a partition between a finite set of components. Other types of time-varying VAR models assume the parameters are a smooth function of time (e.g., Haslbeck, Bringmann, & Waldorp, 2017). However, we would not expect these models to outperform the threshold VAR in this instance for two reasons: First, the threshold VAR model is already able to capture the major source of variation in parameters over time, that is, the step-like switches between stable states. Second, since these models are still based on fitting locally stationary VAR models, the fundamental limitations of approximating the dynamics with linear

³Despite numerous attempts and correspondence with the authors of the package, we were unable to get the model estimates for the dataset described here to converge.

relationships remain. As such, in the next section we examine an approach which aims to recover the exact system of differential equations (DEs) from data, by allowing non-linear terms to enter into a step-wise model building procedure.

3.6 Differential Equation Model Building

In the previous sections we have shown that some models have been able to recover some characteristics of the true model, but that it is generally difficult to make inferences about the characteristics of the true system from these models. Also, since all of these models were misspecified they were fundamentally unable to recover the exact true bistable system. In this section, we aim to recover the exact system of differential equations (DEs) directly from the ideal time series.

3.6.1 Model Building Procedure

The structure of the true model is typically unknown in practice, and therefore has to be learned from the data. [Chow \(2019\)](#) describes a general methodology for building dynamic systems models which consists of two steps: In the first step, we approximate the first-order derivatives by taking difference scores between consecutive measurement occasions, divided by the length of the time-interval between those occasions

$$\frac{dx_{i,t}}{dt} \approx \frac{x_{i,t+1} - x_{i,t}}{\Delta t} \quad (4)$$

where in each case, $\Delta t = .1$, as described in Section 2.2 (cf. [Boker et al., 2010](#)).

In the second step, we use this approximate derivative as the outcome variable, and try to find a regression model that predicts this outcome variable as well as possible with as few parameters as possible. Here, we use results obtained from the statistical models in the previous sections as a starting point, and then follow the standard model-building approach of fitting models with increasing complexity and evaluating the improvement in out-of-sample fit.

From the descriptive statistics (Section 3.1) the marginal lag-0 (Section 3.3) and lag-1 (Section 3.4) relationships and the mean-switching Hidden Markov Model (Section 3.2), we saw that the system is bistable. From dynamical systems theory we know that bistability is only possible in the presence of non-linear terms ([Strogatz, 2015](#)). Similarly, we saw from the TVAR(1) model (Section 3.5) that the linear relationships between pairs of variables is dependent on where in the state-space other variables are located (i.e., below or above a given threshold). Both of these pieces of information suggest the presence of interaction effects between variables: However, we have no information about what specific interaction terms, or what other linear or non-linear dependencies should be in the model. Here, we start out with a main effects-only model, and then add more and more non-linear terms (interactions, quadratic effects, etc.). We evaluate the fit using the mean out-of-bag proportion of explained variance (R^2) obtained from a 10-fold cross-validation scheme (see Appendix E for details), and we choose the model that maximizes this value. If two models result in the same fit, we choose the model with less parameters.

The fit of the all models considered here is shown in Table 1. First, we test a baseline model (Model A) where each derivative is a linear function of all other variables. As we described above, the absence of interaction effects makes it unlikely that this is a suitable candidate model, but it gives us a baseline explained variance value of $R^2 = 0.04664$. In Model B, we add to the baseline model all pairwise interactions between the outcome process (e.g., x_1 when the DV is dx_1/dt) and the other variables in the model x_j (i.e., $x_1 \times x_j$, $\forall j \in p$). Adding these pairwise interactions increases the variance explained to $R^2 = 0.06874$. In Model C, we further extend this model by adding all possible pairwise interactions between all variables $x_i \times x_j$, $\forall (i, j) \in p$. However, we see that adding these parameters in fact leads to a slight decrease in explained variance, $R^2 = 0.06870$, indicating overfitting. For brevity, we display only these three models, but adding further complexity to model in terms of additional interaction terms, quadratic or cubic terms also fails to increase the out-of-bag R^2 (see Appendix E). As such, we can take Model B to be our final model.

Model	$\frac{dx_{i,t}}{dt} \sim a + r_i x_i + \dots$	q	R^2
A	$\sum_{j \neq i} r_j x_j$	5	0.04464
B	$\sum_{j \neq i} R_{ij} x_j + \sum_j^p C_{ij} x_j x_i$	9	0.06874
C	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k$	15	0.06870

Table 1: Model fit results for each of the four models described in text. The second column gives the model equation for each variable, q denotes the number of parameters estimated per univariate regression model, and the final column indicates the mean proportion of explained variance R^2 , calculated on the hold-out sets of a 10-fold cross-validation scheme (for details see Appendix E)

3.6.2 Dynamics and Data Generated by Final Model

We can see that the structure of Model B is highly similar to the structure of our data generating model, with additional main effects between variables, that is, the linear effects denoted by the $p \times 1$ vector r in the true model is replaced by a $p \times p$ matrix R in our chosen model. Furthermore, we can see from the left panel of Figure 9 that the parameter estimates are highly similar, but not exactly equal to the data generating parameters:

$$\begin{aligned} \hat{a} &= \begin{bmatrix} 1.40 & 1.37 & 1.25 & 1.27 \end{bmatrix}^T \\ \hat{\sigma} &= \begin{bmatrix} 1.35 & 1.34 & 1.34 & 1.34 \end{bmatrix}^T \\ \hat{R} &= \begin{bmatrix} 0.88 & 0.02 & -0.01 & 0.01 \\ 0.03 & 0.95 & 0.01 & 0.00 \\ -0.02 & 0.05 & 0.96 & 0.04 \\ 0.04 & -0.01 & 0.08 & 0.91 \end{bmatrix} \\ \hat{C} &= \begin{bmatrix} -0.18 & 0.04 & -0.17 & -0.18 \\ 0.03 & -0.19 & -0.19 & -0.19 \\ -0.18 & -0.19 & -0.19 & 0.03 \\ -0.19 & -0.18 & 0.02 & -0.18 \end{bmatrix} \end{aligned}$$

Figure 9: Left panel: the parameters estimated from the ideal data. Right panel: the vector field defined by the estimated parameters in the left panel. Solid points indicate stable fixed points and empty points indicate unstable fixed points. The solid lines indicate the values at which derivative of positive emotion (orange) and negative emotion (light blue) is equal to zero. At the points at which the two lines meet, both derivatives are equal to zero and the system remains in this (stable) state.

While we would not expect to recover the exact parameters of the true model with a different functional form, we see that the signs, size and relative orderings of parameters in the estimated $\hat{C}$ matrix are quite accurate. Based on these parameter, we recover that there are suppressing effects between valences and reinforcing effects within valences (characteristic 5), that the reinforcing effects are smaller in absolute value than the suppressing effects (characteristic 6), and by capturing approximately the correct functional form, we capture that the microdynamic parameters are dependent only on variables inside the model (characteristic 7). Furthermore, we can see that false positive (i.e., off-diagonal) elements of $\hat{R}$ are of a much smaller order of magnitude than the true positive diagonal elements. The full parameter estimates, with standard errors and p-values are shown in Appendix E.

Beyond inspecting the estimated parameters, we can judge how good of an approximation of the true bistable system our estimated model represents by comparing the dynamics implied by that model to that of the true system. The dynamics of a differential equation model are described by its vector field, which we depict for Model B in the right panel of Figure 9. To construct this

vector field we use the same two-dimensional approximation (positive and negative emotion) as we did in Section 2.1 (see Appendix A for details). The orange and light blue lines are solution lines which indicate the locations where the rate of change in one dimension (orange for no change in positive emotion, light blue for no change in negative) is zero. The points at which these solutions line cross determine the fixed points. We can see that our model correctly identifies three fixed points in this range of values: one stable healthy ($x_1 = x_2 = 4.91, x_3 = x_4 = 1.34$), one stable unhealthy ($x_1 = x_2 = 1.39, x_3 = x_4 = 4.84$), and one unstable fixed point approximately halfway between those two ($x_1 = x_2 = 2.79, x_3 = x_4 = 2.82$). If we compare these global dynamics to the global dynamics of the true bistable system depicted in Figure 1 in Section 2.1, we see that Model B very accurately reproduces these dynamics, approximating the position of the fixed points in the true system closely. From this we can say that the estimated DE model captures characteristics 1 (bistability) and 2 (location of the fixed points) in the true system.

An additional way to evaluate whether the dynamics of the estimated model are similar to the dynamics of the true model is to generate data from the estimated model and compare this data to the original data. Figure 10 shows a time series generated from Model B using a step size of .1:

Figure 10: Data generated from the estimated DE model, with the same initial values as the observed data

We can see that the data looks very similar to the original data generated from the true bistable system: the fixed points are at roughly the same location, there is a difference in variance across the high and low emotion value fixed points (characteristic 3), and there is a similar number transitions (around 14) between the healthy and unhealthy state (characteristic 4). Thus, even though we did not exactly recover the set of true parameters exactly, we seem to have recovered a model that is equal to the true model in all relevant aspects, capturing all of the seven characteristics listed in Section 2.3.

3.6.3 Exact Recovery of Model Parameters

While this model building procedure performed extremely well in this scenario, the findings here should be approached with a note of caution. Observe that, despite negligible sampling error, we do not succeed in recovering the exact parameter estimates. The reason for this is that, while data is generated using an Euler step of $\Delta t = .01$, our ideal time series is created by sub-sampling with $\Delta t = .1$. While this is an unrealistically high sampling frequency, it still means that we cannot estimate the derivative perfectly: As the sampling frequency becomes lower, we would expect the quality of this approximation to degrade.

In theory, to recover the data generating parameters, we would need to fit the *integral solution* form of the differential equation (Strogatz, 2015), as this describes the relationships between observed variables spaced Δt apart, as implied by the differential equation. It is well known that this integral solution may contain a seemingly different set of dependency relationships than the differential equation: variables which are independent in the DE form may be dependent in the

integral form, and the signs and relative orderings of these dependencies may change depending on the value of Δt (Aalen, Røysland, Gran, & Ledergerber, 2012; Kuiper & Ryan, 2018; Ryan, 2018). Because methods based on approximating integral solutions are expected to suffer from similar problems as the two-step DE estimation procedure, and because these methods are difficult to apply in practice, we limit ourselves to the two-step approach in this paper (see Discussion section 5.3 for further details).

3.7 Summary: Analysis of Ideal Time Series

In this section we aimed to recover characteristics of the true bistable system from the ideal time series with 10 measurements each minute using a number of time series analysis tools. Table 2 provides a rough summary of which method recovered which characteristics of the true bistable system:

	Bistability (1)	Position (2)	Variance (3)	Transitions (4)	Suppr./Reinf. (5)	Relative Size (6)	Time-constant (7)
Data Visualization	✓	✓	✓	×	×	×	×
HMM	✓*	✓	✓	✓	×	×	×
Lag-0 / GGM	×	×	×	×	✓	×	✓*
Lag-1 / VAR(1)	×	×	×	×	✓	×	✓*
TVAR(1)	✓*	✓	✓	✓	✓	×	×
DE-Estimation	✓	✓	✓	✓	✓	✓	✓*

Table 2: Summary of which method recovered which of the seven qualitative characteristics listed in Section 2.3 from the ideal time series. The first four characteristics are global dynamics, the last three are microdynamics. The check marks with asterisk indicate that the method includes the characteristic as a model assumption, and can therefore not be considered recovered from the time series.

We showed that data visualization (Histograms and the pairwise marginal relationships in Sections 3.3 and 3.4) revealed bistability and provided a rough estimate of the position and showed that the variances around the fixed points differ. However, when comparing the eye-balled number of switches with the estimates of the HMM, we saw that we missed instances in which the system quickly switched back and forth. The Mean switching HMM recovered all global dynamics, however, we provided bistability as a model assumption, which is why we mark the check mark at the first characteristic with an asterisk.

Turning to methods that capture dependencies between variables, the analysis of lag-0 relationships with the GGM and the analysis of lag-1 relationships with the VAR(1) model (and a GGM on its residuals) fundamentally cannot recover any global dynamics of the bistable system, but they recovered some microdynamics: the characteristic that within valence effect are reinforcing, and between valence effects are suppressing; and that the parameters are constant across time, however this is again an assumption of the model and therefore cannot be considered recovered from the data. The TVAR(1) model was able to recover all global dynamics with the same caveat as in the HMM, that bistability is a model assumption and not recovered from data. Similarly to the VAR model, it recovered the reinforcing/suppressing characteristic. However, a naive interpretation of the model parameters would lead one to conclude that the parameters are time-varying. Finally, the DE-estimation method was able to recover all microdynamics reasonably well, which implies that it also recovered all global dynamics.

The purpose of this section was to establish whether or not each method can recover, in principle, some aspect of the bistable system. To do this we used a highly idealized dataset, with an unrealistically high sampling frequency. As such, the performance of each method described above can be considered an upper bound on its performance in any more realistic scenario. It remains to be seen exactly how the performance of each method, and in general our ability to recover global

and microdynamic characteristics of the system, changes when a more realistic sampling frequency is used.

4 Recovering the Bistable Systems from ESM Data

In this section we analyze a time series that is similar to the ideal time series in all aspects, except that the system is sampled every 90 minutes instead of every six seconds (see Section 2.2). This allows us to investigate how the ability of each method to recover (some characteristic of) the bistable system is affected by having only a low sampling frequency time series, as it is typical for ESM studies.

4.1 Descriptive Statistics, HMM and Lag-0 Relationships

The descriptive statistics, such as histograms, and the lag-0 relationships obtained from the ESM times series (Figures 20 and 21 respectively in Appendix F) are essentially identical to those obtained from the ideal time series, depicted in Figures 4 and 6 in Section 3.1. This makes sense: we have exactly the same amount of data points, sampled from the same system as in the ideal time series case. The only difference is that in the ESM data set 900 time points are “missing” between each measurement of the ESM time series. However, because lag-0 relations do not pick up on any temporal dependence, the lower sampling frequency does not affect the lag-0 relations. While this suggests that lag-0 relations are robust against low sampling frequency, it also puts their utility to infer the dynamics of an underlying dynamical system into question.

The parameter estimates obtained by fitting a two-component mean-switching Hidden Markov Model on the ESM dataset were

$$\mu_1 = \begin{pmatrix} 1.47 \\ 1.46 \\ 4.71 \\ 4.71 \end{pmatrix}, \quad \sigma_1 = \begin{pmatrix} 0.41 \\ 0.41 \\ 0.63 \\ 0.63 \end{pmatrix}, \quad \mu_2 = \begin{pmatrix} 4.71 \\ 4.71 \\ 1.47 \\ 1.47 \end{pmatrix}, \quad \sigma_2 = \begin{pmatrix} 0.64 \\ 0.64 \\ 0.41 \\ 0.41 \end{pmatrix}, \quad A = \begin{matrix} & C_1 & C_2 \\ \begin{matrix} C_1 \\ C_2 \end{matrix} & \begin{pmatrix} 0.915 & 0.085 \\ 0.090 & 0.910 \end{pmatrix} \end{matrix}.$$

and the predicted states for two weeks of the time series are shown in Appendix F Figure 22.

We see a very similar pattern of results as obtained from the HMM fit to the ideal time series in Section 3.2, with the means and standard deviations of state 1 and state 2 reflecting the unhealthy and healthy states respectively. However, the parameters of the estimated transition matrix A for this time series show substantially higher switching probabilities, $a_{12} = .085$ and $a_{21} = .090$. As we can see from Figure 3, although the sub sampled ESM time series contains only 224 observations for a two-week period, rather than 201600, the sampling frequency is still high enough to capture each of the 17 switches between states in this period. That means that, although the amount of transitions that occur over a period of time remains the same, the number of measurement occasions between any two transitions is lower, which results in a higher transition probability. We can see that this higher transition probability captures the number of transitions over two weeks quite accurately — the model predicts between $224 \times .090 \approx 20$ and $224 \times .085 \approx 18$ switches on average over a two week period. As such, the HMM fitted on the ESM time series still allows us to estimate the location of and variance around the two fixed points (characteristics 2 and 3), and approximate the frequency of transitions between these two fixed points (characteristic 4). In fact, the transition probabilities appear to be even more accurate than the ideal case — most likely this numerical imprecision in the ideal case is because the number of transitions relative to total time series was so low that slight changes in the transition probability value lead to very different prediction about the number of transitions over 201600 time points.

4.2 Lag-1 Relationships and VAR model

When analysing the lagged relationships in the ESM time series, we begin to see some striking differences from the analysis of the ideal time series: panels (a) and (b) of Figure 11 show the *marginal* relationship of Content observed at time t with Cheerful at the previously observed time point $t - 1$, and Anxious at time t with Content at time $t - 1$, respectively. Focusing on panel (a), we see that the density of the lagged variables takes on a square-like shape, and each quadrant

seems to be filled with a roughly circular density. This is in contrast to the density of the lagged relationships in the ideal data displayed in Figure 6 in Section 3.4, which was described by two elliptical shapes at the two fixed points.

How can we explain this pattern? In the ideal data, the two elliptical shapes indicate that Content_{t-1} and Content_t tend to be near the same fixed point (two shapes), and that the two variables are positively correlated (elliptical shape). Now, in the ESM time series, we still have most of the density in the upper-right and the bottom-left quadrant, indicating that if Content_{t-1} is at the healthy (unhealthy) fixed point, it is very likely that Content_t is also at the healthy (unhealthy) fixed point (noting that $t-1$ now reflects a 90 min instead of 6 second time interval). However, we now also observe density at the top-left and bottom-right quadrant. These densities represent the situation in which Content_{t-1} is in the healthy (unhealthy) state, but Content_t is in the unhealthy (healthy) state. This situation is created when a switch between states falls within the 90min interval between two ESM measurements. Next, we focus on the shape of the density *within* each of the quadrants: we see that each of the densities have roughly a circular shape, which indicates that Content_{t-1} and Content_t are uncorrelated at each fixed point. This makes sense: in the ESM time series 900 time points are missing between each pair of measurement, which means that there is essentially no temporal dependence anymore between the variables. The relationship between Anxious_{t-1} and Content_t can be explained in an analogous way. Before fitting the VAR model below, this already shows us that it is futile to recover the microdynamics of the bistable system from these data.

Figure 11: Panel (a) shows the relationship between Content and Cheerful, two emotions with the same valence, spaced one measurement occasion apart (i.e., at a lag of one but with 90 minutes between measurements) for the ESM dataset. The red line indicates the best fitting regression model. Similarly, panel (b) shows the relationship between Anxious and Content, two emotions with different valence, at a lag of one. Panel (c) displays the matrix of lagged regression parameters, estimated from a VAR(1) model, as a network, and panel (d) displays the partial correlation matrix of the residuals of the VAR(1) model as a network. This latter network is often referred to as the “contemporaneous” network.

Panel (c) of Figure 11 displays the network of estimated lagged regression coefficients from a VAR(1) model fit to the data. If we were to use these to infer the microdynamic characteristics of the system, we would manage to recover the signs of effects between variables: Negative, suppressing effects between-valence, and positive, reinforcing effects within-valence (characteristic 5, see Section 2.3). However, we again fail to recover the relative size of the suppressing and reinforcing effects. In fact, in this case, all of the estimated auto-regressive and cross lagged effects have approximately equal absolute value $|\hat{\phi}_{ij}| \approx .2$. This means that we also fail to recover the strong auto-regressive relationships encoded by the r parameters in the data generating model, and reflected by the strong auto-regressive effects estimated by the VAR(1) model in the ideal setting. In panel (d) we can see that, as was the case for the ideal time series, we obtain positive residual partial correlations within-valence and negative residual correlations between-valence, although the magnitude of these correlations is now quite high, $\hat{\theta} = 0.51$ and $\hat{\theta} = -0.23$ respectively. In addition, note that the residual variances of each variable in the model is considerably higher than the ideal case, and approximately equal for all variables ($\hat{\Psi}_{ii} \approx 1.1$, explained variance ≈ 0.62). As we would expect, taking only every 900th measurement from the ideal time series means that the predictive power of the VAR model decreases.

How can we reconcile these parameter estimates with what we know of the underlying bistable system? Although we may be tempted to interpret the VAR(1) parameters as reflecting the microdynamic structure, we have already seen in the analysis of the marginal relationships above that the large time-interval between observed measurements means that such an interpretation would be incorrect. Instead, the VAR(1) parameter values in the present situation are fully determined by the global characteristics of the system. Essentially, the estimated lagged relationships reflect that, at a time-scale of 90 minutes, the dynamics of the system from one observation to the next can be boiled down to four possibilities: Either the entire system stays in the same state or moves from one state to the other. Since 1) the most likely behaviour is that the system stays near the same fixed point, and 2) those two fixed points are defined as high-positive and low-negative emotions, or low-positive and high-negative emotions, we end up with positive within-valence relationships (e.g. if Cheerful now is near the high fixed point, it's likely that Content later will be too) and negative between-valence relationships (e.g. if Anxiety now is high, it's likely that Content later will be low). All of the auto-regressive and cross-lagged relationships are of equal value, as essentially all variables have the same value in predicting what fixed point each other variable will be near at the next measurement occasion: enough time elapses between measurements occasions that even the auto-regressive effect is only as predictive as the cross-lagged effects. As noted above, this interpretation is also reflected in the joint densities in panels (a) and (b) of Figure 11: Each density takes the appearance of four quadrants of uncorrelated Gaussian distributions, indicating that the microdynamic dependencies present in the ideal times series are totally absent from the ESM time series.

In summary, having a realistic sampling frequency results in the VAR(1) model providing even less information about the characteristics of the bistable model than in the ideal scenario. The longer time-interval between observations implies that interpreting VAR(1) parameters as reflecting truly microdynamic behaviour would be incorrect: parameters interpreted as reflecting microdynamics in fact have to be interpreted as reflecting the global characteristics of the system. Although the sign of the microdynamic relationships (characteristic 5) is recovered, in this instance it happens that the pattern of microdynamic relationships has the same valence as the pattern of relationships at a longer time-scale, that is, the movement of the process between fixed points. Thus, while in the ideal case the VAR parameters were a mixture of the microdynamics (around each fixed point), and global characteristics (i.e., position of the two fixed points), in the ESM time series these parameters are only reflective of the global characteristics.

4.3 Threshold VAR Model

We saw in the previous sections that inferring global characteristics using ESM data was somewhat successful, but that inferring microdynamic characteristics using a VAR(1) model was impossible. For the ideal time series, we saw that the threshold VAR(1) model was in principle able to capture some microdynamic and some global characteristics, and so in this section we examine how well that performance generalises to our emulated ESM data. We use the same thresholding variable (Cheerful, X_1) and model specification as described in Section 3.5.

Figure 12 displays the main results from the TVAR(1) model estimated on the ESM data. The

estimated threshold is $\hat{\tau} = 2.796$, very close to the estimated threshold in the ideal case, and from panel (a) of Figure 12 we can see that this threshold value does well in separating the time-series into the healthy and unhealthy states. Inspecting the lagged networks for each regime in panels (b) and (c) of Figure 12 we see a similar general pattern of results as the lagged networks for the ideal time series in Figure 8 in section 3.5: In the healthy regime, the negative variables have much stronger cross-lagged effects on the positive variables, and vice versa for the unhealthy regime. However, we see even more differences between regimes in this case than we did for the ideal time series. For instance, in the healthy regime, the within-valence and auto-regressive relationships for the negative variables is much stronger than for the positive variables, a pattern which is flipped for the unhealthy regime. In both regimes, the within-valence cross-lagged parameters are roughly equal to the auto-regressive effects of the variables involved. The estimated means of each regime are given as $\hat{\mu}_2 = \{4.31, 4.31, 1.87, 1.87\}$ for the healthy state and $\hat{\mu}_1 = \{1.74, 1.71, 4.44, 4.62\}$ for the unhealthy state.

Figure 12: Panel (a) shows the first two weeks of the time series, with observations shaded in either grey or white as a function of whether $x_{1,t-1}$ is above or below the threshold $\hat{\tau} = 2.796$. Panels (b) and (c) show the estimated VAR(1) parameters as lagged networks in the healthy (white) and unhealthy (grey) regimes respectively.

We can see from this that the TVAR(1) model for the ESM data succeeds in recovering some global characteristics of the system. Specifically, the estimated mean vectors capture approximately the position of the two stable fixed points (characteristic 2), and the estimated threshold correctly captures the position of the unstable fixed point in the Cheerful dimension. However, the recovery of this characteristic comes with the same caveats as described in Section 3.5: The use of a univariate threshold for this particular configuration of the true system happens to be a good approximation of the unstable fixed point in multi-dimensional space.

Regarding the microdynamics, the lagged parameters in each regime approximately capture that there are reinforcing effects within valence and suppressing effects between valence (characteristic 5). Otherwise, however, the recovery of microdynamic relationships performs worse than

for the ideal time series, as expected. As was the case for the VAR(1) model, the regime-specific lagged parameters here again reflect global characteristics at the 90 minute time-scale rather than microdynamics: Partitioning the joint densities in panels (a) and (b) of Figure 11 using a threshold does not aid us in any way to reproduce microdynamic dependencies which are absent due to the low sampling frequency. Thus, the asymmetry in parameter values across regimes has to be a function of the global characteristics, influenced by both the different variances around the fixed points in each state (i.e., high variance for positive emotions, low variance for negative emotions in the healthy state, and vice versa) and those observations which jump from one fixed point to the other across consecutive measurement occasions, as discussed in the previous section.

As we did throughout Section 3, we could evaluate how well this model describes the bistable system by generating data from it. Notably, the dynamics defined by $\Phi^{(1)}$ and $\Phi^{(2)}$ reflect an *unstable* system in both regimes: The eigenvalues of both contain a value outside the unit circle (i.e., with absolute value greater than one) (Hamilton, 1994). This means that, if we were to generate data using these parameters, the time series would always diverge towards infinity. This instability also precludes us from making any statement regarding the variance of positive and negative emotions in each regime (characteristic 3), as the long run variances implied by the model are infinite. As such, we can say that overall, the set of estimated parameters for the TVAR(1) based on the ESM time series are a poor characterisation of the microdynamics of the model at any time-scale.

In summary, the TVAR(1) model fitted on the ESM time series still picks up a global characteristic of the system, but the recovery of microdynamic characteristics fails. In fact, the relationship between the estimated lagged parameters and the characteristics of the system was much more opaque than in the ideal data case, and our ability to generalize from the estimated parameters to the behaviour of the system at any time scale was considerably worse than in the ideal case. Again here, we should note that the only difference between the ideal and ESM time series is the sampling frequency. Fundamentally, the results here indicate that, if we do not have a sufficiently high sampling frequency, then fitting increasingly complex models, or extensions to simpler models such as the TVAR(1), does not aid us in recovering the characteristics we are interested in: Even when we have an arbitrarily large *number* of observations, we fail to recover basic characteristics of the microdynamics due to the spacing between measurements.

4.4 Differential Equation Model Building

In this section we will examine whether the DE model-building procedure described in Section 3.6 also succeeds in recovering the bistable system when applied to the emulated ESM dataset. Recall from Section 3.6 that for the ideal time series, this method succeeded in recovering the microdynamics of the system: The global characteristics were also considered to be recovered as the global characteristics implied by the estimated model (bistability, position of fixed points) matched up with the actual global characteristics of the underlying system.

4.4.1 Model Building Procedure

Similarly to Section 3.6 we first estimate the derivatives directly from the data by differencing the time series, and then search for the best fitting model by fitting a series of regression models with increasing complexity. Table 3 displays the fit of seven increasingly complex regression models, evaluated using the mean out-of-bag proportion of explained variance R^2 :

Model	$\frac{dx_{i,t}}{dt} \sim a + r_i x_i + \dots$	q	R^2
A	$\sum_{j \neq i} r_j x_j$	5	0.13991
B	$\sum_{j \neq i} R_{ij} x_j + \sum_j^p C_{ij} x_j x_i$	9	0.16827
C	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k$	15	0.16928
D	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k + \sum_j^p \gamma_j x_j^3$	19	0.19455
E	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k + \sum_j^p \gamma_j x_j^3 + \sum_{j \neq k \neq l}^p \zeta_j(x_j x_k x_l)$	23	0.19801
F	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k + \sum_j^p \gamma_j x_j^3 + \sum_{(j,k,l)}^p \zeta_j(x_j x_k x_l)$	35	0.19940
G	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k + \sum_j^p \gamma_j x_j^3 + \sum_{(j,k,l)}^p \zeta_j(x_j x_k x_l) + \sum_{(j,k,l,m)}^p \eta_j(x_j x_k x_l x_m)$	70	0.20420

Table 3: Model fit results for each of the seven models described in text, for the emulated snapshot ESM data. The second column gives the model equation for each variable, q denotes the number of parameters estimated per univariate regression model, and the final column indicates the mean proportion of explained variance R^2 , calculated on the hold-out sets of a 10-fold cross-validation scheme (for details see Appendix E)

Model A ($R^2 = 0.13991$), Model B ($R^2 = .16827$) and Model C ($R^2 = 0.16928$) are the same models as introduced in Section 3.6. However, since we did not observe a clear drop in R^2 as we increased model complexity from Model A to Model C, we also assess the fit of four additional models. Model D adds cubic main effects $x_1^3, \dots, x_4^3$ as predictors, increasing the model fit to $R^2 = 0.19455$. Model E adds four three-way interactions $(x_i \times x_j \times x_k, i \neq j \neq k)$ to this, further increasing the model fit to $R^2 = 0.19801$. Adding yet more three-way interactions $(x_i \times x_j \times x_k, \forall (i, j, k) \in p)$ in Model F still increases model fit ($R^2 = 0.19940$), as does adding all possible four-way interactions in Model G ($R^2 = 0.20420$). As it is not possible to specify more unique product interaction terms, we consider Model G to be our final model.

4.4.2 Dynamics and Data Generated by Final Model

Clearly, the model-building procedure for the emulated ESM data failed to recover the functional form of the true bistable system. Furthermore, we have arrived at a final model which is so complex ($4 \times 70 = 280$ vs. $4 \times 6 = 24$ parameters in the true model) that it is close to uninterpretable. In theory we could continue adding complexity to the model in the form of non-linear transformations or spline functions, which we know to be absent from the data generating mechanism, but which may increase fit. However, this would make the model even more difficult to interpret.

In the left panel of Figure 13 we present the estimated parameters that are also contained in the true model, with full parameter estimates and standard errors shown in Appendix E.3. We can see that the estimates deviate widely from the parameters in the true bistable system. In addition, the estimated parameters in the C matrix fail to capture the sign and relative ordering of all parameters in the true C matrix, though a full evaluation of whether suppressing and reinforcing effects of different sizes are present (i.e., characteristics 5 and 6) is infeasible due to the large number of parameters present in the model. Thus, we can say that this approach fails to recover the microdynamics of the system at least to the degree that they can be interpreted.

$$\hat{a} = \begin{bmatrix} -0.07 & -0.16 & 0.31 & 0.20 \end{bmatrix}^T$$

$$\hat{\sigma} = \begin{bmatrix} 0.01 & 0.01 & 0.01 & 0.01 \end{bmatrix}^T$$

$$\hat{R} = \begin{bmatrix} 0.19 & 0.07 & -0.06 & -0.01 \\ 0.16 & 0.18 & 0.01 & -0.04 \\ -0.09 & -0.29 & -0.05 & 0.01 \\ -0.12 & -0.22 & -0.07 & 0.10 \end{bmatrix}$$

$$\hat{C} = \begin{bmatrix} -0.02 & -0.11 & -0.02 & -0.05 \\ -0.10 & -0.02 & -0.01 & -0.05 \\ 0.00 & 0.05 & -0.01 & 0.01 \\ 0.00 & 0.02 & -0.04 & -0.03 \end{bmatrix}$$

Figure 13: Left panel: the parameters estimated from the snapshot ESM time series. Right panel: the vector field defined by the estimated parameters. Solid points indicate stable fixed points and empty points indicate unstable fixed points. The solid lines indicate the values at which derivative of positive emotion (orange) and negative emotion (light blue) is equal to zero. At the points at which the two lines meet, both derivatives are equal to zero and the system remains in this (stable) state.

While the system did not recover the microdynamics in the sense that it captures the qualitative characteristics of the true bistable system, it could still be the case that this more complex system exhibits global characteristics that are similar to the true bistable system. Similarly to Section 3.6, we can evaluate these dynamics by inspecting its vector field, shown on the right-hand side of Figure 13. As in the vector field obtained from the ideal data (Figure 9), the intersections of the two solution lines (blue and orange) indicate the position of the different fixed points in the shown range of the state space. These fixed points are further denoted by dots, with filled dots indicating a stable fixed point, and empty dots indicating an unstable fixed point.

We can immediately see from Figure 13 that the stability landscape is much more complex than the one of the true bistable system, with high-degree polynomial solution lines, and with four rather than three fixed points. Interestingly, the system correctly identifies (1) that there are two stable fixed points relating to healthy state ($x_1 = x_2 = 4.51, x_3 = x_4 = 1.67$) and unhealthy state ($x_1 = x_2 = 1.71, x_3 = x_4 = 4.47$) and (2) that there is an unstable fixed point approximately half-way between those two ($x_1 = x_2 = 2.99, x_3 = x_4 = 3.19$). Despite having an entirely different functional form, the estimated model *does* capture two stable fixed points (characteristic 1) and the approximate position of those fixed points (characteristic 2). This shows that Model G performs well in capturing the characteristics of the system for emotion values that were observed in the time series, that is, near the two stable fixed points.

Crucially, however, we cannot say that this system recovers the global dynamics of the true system, not least because the system contains an additional unstable fixed point at ($x_1 = x_2 = 5.31, x_3 = x_4 = 7.94$), which is not present in the true bistable system. The presence of this unstable fixed point means that if, for instance, both negative and positive emotions take on a high value simultaneously, then the system enters an unstable region and diverges to infinity. If we examine the behaviour of the system even further outside the range of observed values ($-\infty > X > 0$ and $10 < X < \infty$) even more fixed points and regions of stability and instability can be found. We can further demonstrate these dynamics by generating data from Model G. Figure 14 shows a time series generated from the difference-form of Model G (i.e., with a step size equal to that of the observed data).⁴ We see that the process moves between the healthy and unhealthy fixed point

⁴This is obtained by re-fitting the differential equation using the unscaled difference $x_{i,t+1} - x_{i,t}$ as the outcome variable, leading to equivalent results with parameters approximately scaled by $dt = 90$. The residual variance used is the estimated residual variance scaled down to .65 the magnitude, to account for the non-normal residual distribution. Using the estimated residual standard deviation results in shocks which immediately move the system into an unstable region.

for the first ten days, exhibiting the bistable behaviour we see in the true system. However around the eleventh day, the stochastic input is large enough to move the system to an unstable region in the vector field and which leads the system to diverge.

Figure 14: Data generated from the estimated DE model, with the same initial values as the "ideal" data

Note that the complexity of the final model here is not a result of over-fitting the data because we performed model selection based on the out-of-bag R^2 , an approximation of the out-of-sample R^2 . Rather, the complexity of this model can be attributed to two factors. First, due to the low sampling frequency, our approximation of the derivative at each point in time is poor. The second, as we discussed in Section 3.6.3, is that given the spacing between observations, the best one can hope for is to approximate the integral solution to the data generating equation, which is likely of a highly complex functional form. The ability of the misspecified Model G to reproduce some characteristics in regions where we have observed data can be attributed to the high flexibility afforded by the many non-linear terms. In that sense, this behaviour is highly comparable to the problem of using a high-degree polynomial regression model to make predictions outside of the range of observed values. The vector field in Figure 14 is constructed by obtaining predicted values for the derivatives across a grid of input values: As such, it is unsurprising that the vector field is accurate where the input values are close to the observed data, and inaccurate elsewhere.

In summary, we do not at all recover the functional form or parameters of the system; we do recover some of the global characteristics and behaviour of the system in the region where we have observations, capturing that there are two stable fixed points and one unstable fixed point, and their locations. However, the estimated model also implies the presence of at least one extra unstable fixed point, which has major implications for the dynamics of the model, implying divergent behaviour. Thus, the estimated model implies fundamentally different microdynamic and global characteristics. Based on the simulated data in Figure 14, it does not seem that we correctly capture the variability around these fixed points, or the frequency of transitions, as any reasonable simulation of data from this model eventually leads the system to diverge. Crucially, we fail in recovering an interpretable approximation of the data generating model. As such, it is not feasible to assess whether there are truly suppressing effects between valences and reinforcing effects within valences, or the relative size of these effects (characteristics 5 and 6).

4.5 Summary: Analysis of ESM Time Series

In this section, we aimed to investigate to which extent lowering the sampling frequency affects the the ability of our considered methods to recover the bistable system. Our findings are summarized in Table 4:

	Bistability (1)	Position (2)	Variance (3)	Transitions (4)	Suppr./Reinf. (5)	Relative Size (6)	Time-constant (7)
Data Visualization	✓	✓	✓	×	×	×	×
HMM	✓*	✓	✓	✓	×	×	×
Lag-0 / GGM	×	×	×	×	×	×	✓*
Lag-1 / VAR(1)	×	×	×	×	×	×	✓*
TVAR(1)	✓*	✓	✓	✓	×	×	×
DE-Estimation	×	×	×	×	×	×	×

Table 4: Summary of which method recovered which of the seven qualitative characteristics listed in Section 2.3 from the ESM time series. The first four characteristics are global dynamics, the last three are microdynamics. The check marks with asterisk indicate that the method includes the characteristic as a model assumption, and can therefore not be considered recovered from the time series.

Our main findings are that, in general, we remain able to recover global characteristics of the system using simple methods, but that we are completely unable to recover any of the microdynamics: We saw that each approach which aimed to capture microdynamic characteristics either deteriorated dramatically in performance (for the VAR and TVAR approaches) or broke down altogether (for the DE model building approach) as soon as we applied them to a time series obtained with a more realistic sampling frequency. This is despite the fact that the time series we used in this section can be considered a highly idealized approximation to ESM time series, in terms of the number of observations and the quality of measurements, suggesting that sampling frequency is a fundamental barrier to inference which needs further investigation. As a side result, we can say that lowering the sampling frequency typically made it much more difficult to interpret and understand the results of different methods: In particular for methods which involved lagged relationships of some kind (i.e. the VAR, TVAR and DE approaches), it was difficult to ascertain precisely what features of the system the estimated parameters reflected.

The recovery of global characteristics was more successful: Using data visualization and the Hidden Markov Model it was still possible in principle to learn about the position, variance around and frequency of transitions between fixed points, and the threshold estimate from the TVAR model succeeded in capturing the unstable fixed point. Finally, the predictions made by the best-fitting differential equation model did allow us to get some tentative indication of bistable behaviour, and the possible location of stable fixed points. However, the resulting model suffered from a high degree of complexity, limiting both substantive interpretation and our ability to extrapolate the model parameters to predict the behaviour of the system under different conditions.

In summary, the results in this section call into question to what extent it is possible to investigate moment-to-moment microdynamics using data sampled at a rate typical of ESM studies. We have showed that interpreting model estimates from ESM time series as reflecting the microdynamics can be highly misleading, when the process of interest is varying at a higher frequency than the sampling frequency. Although the recovery of global characteristics is more promising, we remind the reader that the time series considered here is still highly idealized, with essentially infinite sampling size, and so the performance of these methods should be considered an upper bound on performance in any realistic situation.

5 Discussion

In this paper we explored to what extent dynamical systems models can be recovered from psychological time series by investigating two successive questions: first, how well does a set of popular and more advanced methods recover (characteristics of) a basic bistable system with an ideal data set sampled at extremely high sampling frequency (every six seconds)? And second, how is the performance of each method affected when reducing the sampling frequency to one measurement every 90min, which is typical for ESM studies.

When analyzing the ideal time series we found that the popular VAR model (and the GGM fitted on its residuals) can in principle not recover the global dynamics of the true bistable system, and only recovers some of its microdynamics. However, we showed that descriptive statistics, data visualization and statistical models which are based on mixtures (the HMM and threshold VAR) were able to capture the global dynamics of the bistable system. The only method that recovered the full bistable system was a differential equation (DE) model building procedure. Reducing the sampling frequency from every six seconds to every 90 minutes affected the considered methods differently: The VAR model and its extensions no longer recover any microdynamics, and the DE-estimation procedure fails. However, descriptives, data visualization and appropriate statistical models still recover the global dynamics. Overall, our analysis therefore suggests that it is neither possible to estimate dynamical systems directly from realistic time series, nor is it possible to reliably infer its microdynamics from the parameter estimates of statistical models.

5.1 Implications for Studying Mental Disorders as Complex Systems

Our results raise fundamental questions about how to study mental disorders from a complex systems perspective. First, they show that it is unclear what exactly one can in principle conclude from statistical models estimated from psychological time series about an underlying dynamical system. Clearly, these models are always misspecified (i.e., do not include the true system as a special case), so one cannot hope to directly recover the underlying dynamical system. More surprisingly, however, recovering the qualitative characteristics of the true system also turned out to be difficult: while it was possible to recover the global dynamics, no statistical model correctly recovered the microdynamics. For example, the VAR model fundamentally cannot capture the global characteristics (e.g. location of fixed points) of the true bistable system and only recovered some of its microdynamics (e.g., reinforcing vs. suppressing effect between two variables). This is a problem for the emerging framework of studying mental disorders as complex systems, because one is typically interested in the microdynamics (the “mechanics”) of a disorder because one hopes to intervene on them. In contrast, it is usually less clear how interventions can target global dynamics, since they can be seen as the aggregate behavior implied by the microdynamics. Especially the failure of the popular VAR model to correctly recover the qualitative nature of the microdynamics in the true model is concerning, because it calls into question whether it allows any reliable conclusions about an underlying dynamical system. It therefore seems to be an open question how useful VAR models and other statistical models are to studying mental disorders from a complex systems perspective.

Second, the analysis of the ESM time series raises the question of which process can be recovered with which sampling frequency. While we were still able to recover the global dynamics of the system, each method that provides some approximation of the microdynamics was strongly affected by sampling only every 90min instead of every six seconds: The qualitative characteristics of the VAR and TVAR models were even less in agreement with those of the true model, and the DE-estimation method, which was the only fully successful method in the ideal data case, returned a model with uninterpretable parameters and incorrect global- and microdynamics. Thus, our results suggest what also seems intuitive: it is impossible — or at least extremely difficult — to recover microdynamics at a time scale that is much smaller than the sampling frequency. This intuition is also in line with sampling theorems from the field of signal processing: for example, the Nyquist-Shannon sampling theorem states that a sine wave (a process much simpler than our bistable system) that completes one cycle within, say, 2 minutes, has to be sampled at least every minute to be recovered (e.g., [Marks, 2012](#); [Papoulis & Pillai, 2002](#)). This suggests that it is futile to try use a time series sampled every 90 minutes to directly recover dynamics of emotions that operate on a time scale of seconds or minutes ([Houben, Van Den Noortgate, & Kuppens, 2015](#)) or even from moment to moment ([Wichers et al., 2015](#)). However, this also means that ESM time series can certainly be used to recover processes that unfold at a time scale of several hours or days.

To summarize, we identified two fundamental problems to studying mental disorders from a complex systems perspective: first, even with extremely high sampling frequency it is generally unclear how to make inferences from a statistical model to an unspecified dynamical systems model. Second, the sampling frequency of the data collection constrains the type of processes one can recover. Specifically, a process can only be recovered if the sampling frequency is sufficiently high. Clearly, these are profound problems every empirical discipline struggles with and no simple

answers can be expected. Indeed, they might imply that studying some phenomena of mental disorders will always remain out of reach. That said, we believe that that much progress *can* be made by studying mental disorders as complex systems and that acknowledging and studying the above issues allows one to do so more efficiently. As a way forward, in the following section we suggest a new research strategy based on proposing substantively plausible dynamical systems, which opens up avenues to creatively tackle the two problems identified in this section.

5.2 Moving Forward: Proposing Plausible Dynamical Systems Models

A more abstract perspective on the first problem identified in the previous section is the following: we have parameters of a statistical model which we estimated from a time series sampled from some system, and we hope to infer some characteristics (e.g., global or microdynamics) of the data generating system from them. The problem, however, is that the mapping from parameters of statistical model to the parameters and structure (and the implied dynamics) of the true model is unknown. Thus, this inference cannot be made. The main reason this mapping is unknown is the trivial reason that no true dynamical system model is specified.

We propose that, in order to overcome this fundamental problem, researchers must begin the research process by proposing a “first guess” model of the dynamical system. While this is clearly difficult and the validity of this model should certainly be questioned, this approach has one major advantage: it is much clearer how to draw conclusions from descriptive statistics, data visualizations or statistical models about the underlying dynamical systems model. This is because one can generate time series from the “first guess” model and fit a statistical model of choice; that way, one always knows which statistical model is implied by the dynamical systems model. This implied model can then be compared to the model fitted to corresponding empirical data. If the implied model and the empirical model are in agreement, we have tentative evidence that the dynamical system model is correct; if not, we can use the nature of the disagreement to improve the dynamical system model. Clearly, this modeling approach, which is typical to more quantitative disciplines such as physics, chemistry and biology, is different to the statistical modeling framework most psychological researchers are familiar with. On the one hand these formal dynamical systems models are harder to build, since they cannot be estimated directly from the data. On the other hand, they are powerful enough to be plausible for complex phenomena such as mental disorders, and have additional benefits such as synthesizing knowledge, revealing unknowns, laying open hidden assumptions and the checking the internal consistency of a model (Epstein, 2008; Lewandowsky & Farrell, 2010; Smaldino, 2017).

This modeling approach also allows to tackle the problem of sampling frequencies that are too low to recover the process of interest directly, as one can generate a time series from the specified dynamical systems model and reduce the sampling frequency to a level that is also available in empirical data. Then, similarly to above, one can again compute the statistical model of choice that is implied by the dynamical systems model with a given sampling frequency, compare it to the corresponding model fit on empirical data, and in the case of disagreement adapt the dynamical systems model accordingly. Of course, this approach is not a panacea: less information is available when the sampling frequency is low, which makes model identification more difficult. However, specifying an initial dynamical systems model allows one to gauge how difficult it is to recover a given type of process on a given time scale with a given sampling frequency.

In addition, starting out with a dynamical systems model also allows to study the *measurement function* that defines the mapping from the variables in the dynamical systems to the obtained measurements, a topic we only touched on briefly in this paper. In our emulated ESM time series we took the measurement function to return the exact values of variables at the time point of measurement. However, different questions imply different measurement functions. For example, if the phrasing of a particular question refers to the entire period since the last measurement, one could instead define the measurement as a function of the variable values since the last measurement, such as the average. Next to formalizing which experiences an ESM question refers to exactly, defining a measurement function also allows to formalize known response and memory biases, such as the recency effect (Ebbinghaus, 1913/2013).

Finally, having a plausible dynamical systems model allows one to explicitly address a behavior that has been largely ignored in the psychological time series modeling literature: the fact that humans sleep. Sleep interacts with essentially everything physiological and psychological, is part of the definition of several mental disorders (e.g., Major Depression) and related to many more

(e.g., Walker, 2017). Thus, for many mental disorders, it seems necessary for a plausible model to include sleep. This may also allow using existing data in new ways, because data around the “day-night shift” does not have to be excluded anymore, but instead can be used to test hypotheses about the sleep-related assumptions of the dynamical systems model.

Clearly, this brief outline of the proposed modeling approach leaves many important questions unanswered: Where should the initial “first guess” dynamical system come from? How to formalize different substantive aspects in a dynamical systems model? Which statistical models should one choose to test which implications of the dynamical systems model? Given some disagreement between predicted and observed statistics, how should one adapt the existing dynamical systems model? These and other questions are difficult ones and answering them requires the combined creativity of a large research community. Nonetheless, a more detailed account of our proposed new modeling approach would be desirable. However, since such a detailed account is beyond the scope of the present work, we address it in a forthcoming paper.

5.3 Limitations

Several limitations of our work require discussion. First, our goal was to explore to which extent one can recover (bistable) dynamical systems for mental disorders from psychological time series. However, we only studied a single bistable system. Therefore, it could be that the fundamental problems identified in the paper and summarized in Section 5.1 are in fact a particularity of the chosen bistable system. This, however, seems extremely unlikely: First, because we identify the problems in our paper as examples of well-known issues such as model misspecification and sampling systems with a sampling frequency that is sufficient for recovery. Second, the bistable system we chose is arguably the simplest bistable system for four variables one can find. Choosing a different model therefore results most likely in choosing a more complex model, and our intuition is that the methodological difficulties discussed in this paper become more and not less relevant in such models.

Second, a more specific criticism of our bistable system could be that the time scale of the process is unrealistically small, and we therefore exaggerated the problem of recovering dynamics of psychological processes from ESM time series. We agree that it is possible that some psychological processes are easier to recover from ESM data than the dynamical system used in this paper. Thus, strictly speaking, we only showed that it is impossible to recover a system if the sampling frequency does not appropriately match the time scale of the system. In principle, it is therefore an open question whether there is a mismatch between the time scale of the system of interest and the available sampling frequency. However, intuition — and the sampling theorems such as the one mentioned in Section 5.1 — strongly suggest that it is impossible (or at least very difficult) to recover a process that operates at a time scale of seconds or minutes from an ESM time series that is measured every 1.5 hours. Clearly, however, our investigation is only a first treatment of the important topic of sampling frequency, and much work on it is required to establish a tight connection between psychological time series and dynamical systems models.

Third, one could reverse the argument in the previous two paragraphs and argue that our model is so ideal that many analyses perform better than in most realistic applications. This is certainly the case for the Threshold VAR model, which performs well only because of the simple dynamics of the bistable system as we discussed in Section 3.5. Other examples are the descriptive statistics and data visualization which may not be as insightful if fixed points are closer to each other and if there is more noise in the system. Also, the two-step approach to estimating the differential equations in Section 3.6 may work less well for a more complicated model. Thus, we would agree with this assessment, however chose to use a simple bistable system in order to make the paper more accessible to applied researchers.

Fourth, we analyzed a bistable system whose structural parameters do not change over time. However, much of the framework of considering mental disorders as complex systems is based on the idea that pathology is defined with respect to a structural change in the underlying system, and therefore structural change is of central interest. We expect that structural change renders the recovery of a system more difficult, and we therefore did not include this feature in order to keep the paper at a reasonable length. However, we believe that future methodological research into how to recover such structural changes both in principle and with realistic time series would be extremely helpful to better understand phenomena such as early warning signals (Scheffer et al., 2009; van de Leemput et al., 2014) and more generally structural change in mental disorders.

Fifth, in order to estimate a differential equation from data, we took a rather simple two-step approach based on local linear approximation of the derivative (cf. [Boker et al., 2010](#)). This approach involves first estimating the derivative itself using scaled difference scores, and then using this derivative as an outcome variable in a regression model. While this method benefits from being extremely simple to implement, we could expect that it would perform poorly in the presence of low sampling frequency as the quality of the derivative approximation degrades (as we noted in Section 3.6.3 and observed in Section 4.4). There are multiple alternative approaches to estimating DE equations which we did not consider here. For example, approaches based on numerical integration of the DE equation during estimation, such as implemented in *dynR* ([Ou et al., 2019](#)) and *stan* ([Carpenter et al., 2017](#)) (with additional functionality in the *ctsem* package [Driver, Oud, & Voelkle, 2017](#)) may in general perform better than the two-step procedure when the sampling frequency is low. However, for the analysis shown in the present paper, neither the *ctsem* nor *dynr* package performed better than the two-step approach. In general, however, more research is needed to map out which method deals best with the problem of low sampling frequencies.

Lastly, throughout our paper we studied how well certain analysis methods can recover the true bistable system *in principle*. We did this by studying the population properties of these methods, that is, the situation in which one has essentially infinite sample size, which we approximated with a huge number (201600) of measurements. This was necessary in order to study the more fundamental questions of (1) whether a given method can recover our bistable system in principle and (2) whether a given method can recover our bistable system based on a time series with realistic sampling frequency. We did this because it would be meaningless to study the performance of a method as a function of sample size, if the method already fails with infinite sample size. Clearly, however, to apply any of the methods we studied in practice, one has to know how reliable they are with which sample size, and much more research is necessary to map out these sample size requirements (e.g., [Dablander, Ryan, & Haslbeck, 2019](#)).

5.4 Summary

In the present paper we identified two fundamental problems involved in studying mental disorders from a complex systems perspective: first, it is generally unclear what to conclude from a statistical model about an unspecified underlying complex systems model. Second, if the sampling frequency of a time series is not high enough, it is futile to attempt to recover the microdynamics of the underlying complex system. In response to these problems, we proposed a new modeling strategy that takes an initial substantively plausible dynamical systems model as a starting point, and develops the dynamical systems model by testing its predictions. In this approach it is much clearer what we can learn from data and statistical models about an underlying dynamical system, and in addition it provides avenues to move the field forward by formalizing the sampling process, measurement, response and memory biases, measurement reactivity and the influence of sleep.

Acknowledgements

We would like to thank Denny Borsboom, Fabian Dablander, Charles Driver, Peter Edelsbrunner, Jens Lange, Don Robinaugh, Noémi Schuurman, Arnout Smit and Lourens Waldorp for their feedback on an earlier version of this manuscript. Especially, we would like to thank Don Robinaugh for many great discussions about a related project, which influenced the present paper.

JMBH has been supported by the European Research Council Consolidator Grant no. 647209. OR has been supported by the Netherlands Organisation for Scientific Research Onderzoekstalent Grant 406-15-128.

References

- Aalen, O. O., Røysland, K., Gran, J. M., & Ledergerber, B. (2012). Causality, mediation and time: a dynamic viewpoint. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 175(4), 831–861.
- Atkinson, K. E. (2008). *An introduction to numerical analysis*. John Wiley & sons.

- Boker, S. M., Deboeck, P. R., Edler, C., & Keel, P. K. (2010). Generalized local linear approximation of derivatives from time series. In *Statistical methods for modeling human dynamics: An interdisciplinary dialogue* (pp. 161–178). Routledge.
- Borsboom, D. (2017). A network theory of mental disorders. *World Psychiatry*, 16(1), 5–13.
- Bringmann, L. F., Vissers, N., Wichers, M., Geschwind, N., Kuppens, P., Peeters, F., ... Tuerlinckx, F. (2013). A network approach to psychopathology: new insights into clinical longitudinal data. *PloS one*, 8(4), e60188.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1).
- Chow, S.-M. (2019). Practical tools and guidelines for exploring and fitting linear and nonlinear dynamical systems models. *Multivariate Behavioral Research*, 1–29.
- Chow, S.-M., Ou, L., Ciptadi, A., Prince, E. B., You, D., Hunter, M. D., ... Messenger, D. S. (2018). Representing sudden shifts in intensive dyadic interaction data using differential equation models with regime switching. *Psychometrika*, 83(2), 476–510.
- Cramer, A. O., van Borkulo, C. D., Giltay, E. J., van der Maas, H. L., Kendler, K. S., Scheffer, M., & Borsboom, D. (2016). Major depression as a complex dynamic system. *PloS one*, 11(12), e0167490.
- Dablander, F., Ryan, O., & Haslbeck, J. M. B. (2019, Jan). *Choosing between AR(1) and VAR(1) models in typical psychological applications*. PsyArXiv. Retrieved from psyarxiv.com/qgewy doi: 10.31234/osf.io/qgewy
- De Haan-Rietdijk, S., Gottman, J. M., Bergeman, C. S., & Hamaker, E. L. (2016). Get over it! a multilevel threshold autoregressive model for state-dependent affect regulation. *Psychometrika*, 81(1), 217–241.
- Driver, C. C., Oud, J. H. L., & Voelkle, M. C. (2017). Continuous time structural equation modeling with R package ctsem. *Journal of Statistical Software*, 77(5), 1–35. doi: 10.18637/jss.v077.i05
- Ebbinghaus, H. (1913/2013). Memory: A contribution to experimental psychology. *Annals of Neurosciences*, 20(4), 155.
- Epskamp, S., Waldorp, L. J., Möttus, R., & Borsboom, D. (2018). The Gaussian Graphical Model in Cross-Sectional and Time-Series Data. *Multivariate Behavioral Research*, 1–28. doi: 10.1080/00273171.2018.1454823
- Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11(4), 12.
- Fabio Di Narzo, A., Aznarte, J. L., & Stigler, M. (2009). tsDyn: Time series analysis based on dynamical systems theory [Computer software manual]. Retrieved from <https://cran.r-project.org/package=tsDyn/vignettes/tsDyn.pdf> (R package version 0.7)
- Fisher, A. J., Reeves, J. W., Lawyer, G., Medaglia, J. D., & Rubel, J. A. (2017). Exploring the Idiographic Dynamics of Mood and Anxiety via Network Analysis. *Journal of Abnormal Psychology*, 126(8), 1044–1056. doi: 0.1037/abn0000311
- Freedman, H. I. (1980). *Deterministic mathematical models in population ecology* (Vol. 57). Marcel Dekker Incorporated.
- Groen, R. N., Snippe, E., Bringmann, L. F., Simons, C. J., Hartmann, J. A., Bos, E. H., & Wichers, M. (2019). Capturing the risk of persisting depressive symptoms: A dynamic network investigation of patients’ daily symptom experiences. *Psychiatry Research*, 271, 640–648. doi: 10.1016/j.psychres.2018.12.054
- Hamaker, E. L. (2012). Why researchers should think “within-person”: A paradigmatic rationale. *Handbook of Research Methods for Studying Daily Life*, 43, 61.
- Hamaker, E. L., & Grasman, R. (2012). Regime switching state-space models applied to psychological processes: Handling missing data and making inferences. *Psychometrika*, 77(2), 400–422.
- Hamaker, E. L., Grasman, R. P., & Kamphuis, J. H. (2010). Regime-switching models to study psychological process. In P. C. M. Molenaar & K. M. Newell (Eds.), *Individual pathways of change: Statistical models for analyzing learning and development* (p. 155–168). Washington, DC: American Psychological Association.
- Hamaker, E. L., Grasman, R. P., & Kamphuis, J. H. (2016). Modeling bas dysregulation in bipolar disorder: Illustrating the potential of time series analysis. *Assessment*, 23(4), 436–446.
- Hamaker, E. L., Zhang, Z., & van der Maas, H. L. (2009). Using threshold autoregressive models

- to study dyadic interactions. *Psychometrika*, 74(4), 727.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society*, 357–384.
- Hamilton, J. D. (1994). *Time series analysis* (Vol. 2). Princeton, NJ: Princeton University Press.
- Haslbeck, J., Bringmann, L. F., & Waldorp, L. J. (2017). How to estimate time-varying vector autoregressive models? a comparison of two methods. *arXiv preprint arXiv:1711.05204*.
- Hirsch, M. W., Smale, S., & Devaney, R. L. (2012). *Differential equations, dynamical systems, and an introduction to chaos*. Academic press.
- Hosenfeld, B., Bos, E. H., Wardenaar, K. J., Conradi, H. J., van der Maas, H. L., Visser, I., & de Jonge, P. (2015). Major depressive disorder as a nonlinear dynamic system: bimodality in the frequency distribution of depressive symptoms over time. *BMC Psychiatry*, 15(1), 222.
- Houben, M., Van Den Noortgate, W., & Kuppens, P. (2015). The relation between short-term emotion dynamics and psychological well-being: A meta-analysis. *Psychological Bulletin*, 141(4), 901.
- Kalisch, R., Cramer, A. O., Binder, H., Fritz, J., Leertouwer, I., Lunansky, G., ... others (2019). Deconstructing and reconstructing resilience: a dynamic network approach. *Preprint*.
- Kendler, K. S. (2019). From many to one to many—the search for causes of psychiatric illness. *JAMA Psychiatry*. doi: 10.1001/jamapsychiatry.2019.1200
- Kuiper, R. M., & Ryan, O. (2018). Drawing conclusions from cross-lagged relationships: Reconsidering the role of the time-interval. *Structural Equation Modeling: A Multidisciplinary Journal*, 25(5), 809–823.
- Leroux, B. G. (1992). Consistent estimation of a mixing distribution. *The Annals of Statistics*, 1350–1360.
- Lewandowsky, S., & Farrell, S. (2010). *Computational modeling in cognition: Principles and practice*. SAGE publications.
- Lutz, W., Ehrlich, T., Rubel, J., Hallwachs, N., Röttger, M.-A., Joras, C., ... Tschitsaz-Stucki, A. (2013). The ups and downs of psychotherapy: Sudden gains and sudden losses identified with session reports. *Psychotherapy Research*, 23(1), 14–24.
- Marks, R. J. I. (2012). *Introduction to shannon sampling and interpolation theory*. Springer Science & Business Media.
- Molenaar, P. C. (2004). A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. *Measurement*, 2(4), 201–218.
- Nelson, B., McGorry, P. D., Wichers, M., Wigman, J. T., & Hartmann, J. A. (2017). Moving from static to dynamic models of the onset of mental disorder: a review. *JAMA Psychiatry*, 74(5), 528–534.
- Ou, L., Hunter, M. D., & Chow, S.-M. (2019). dynr: Dynamic modeling in r [Computer software manual]. Retrieved from <https://CRAN.R-project.org/package=dynr> (R package version 0.1.14-9)
- Papoulis, A., & Pillai, S. U. (2002). *Probability, random variables, and stochastic processes*. Tata McGraw-Hill Education.
- Pe, M. L., Kircanski, K., Thompson, R. J., Bringmann, L. F., Tuerlinckx, F., Mestdagh, M., ... others (2015). Emotion-network density in major depressive disorder. *Clinical Psychological Science*, 3(2), 292–300.
- Rabiner, L. R. (1989). A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), 257–286.
- Ryan, O. (2018, May). *Interventions in dynamic systems: A causal approach to continuous-time mediation analysis*. PsyArXiv. Retrieved from psyarxiv.com/n2fwt doi: 10.31234/osf.io/n2fwt
- Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., ... Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature*, 461(7260), 53.
- Scheffer, M., Bolhuis, J. E., Borsboom, D., Buchman, T. G., Gijzel, S. M., Goulson, D., ... others (2018). Quantifying resilience of humans and other animals. *Proceedings of the National Academy of Sciences*, 115(47), 11883–11890.
- Schwarz, G., et al. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461–464.
- Smaldino, P. E. (2017). Models are stupid, and we need more of them. In *Computational social psychology* (pp. 311–331). Routledge.
- Snippe, E., Viechtbauer, W., Geschwind, N., Klippel, A., De Jonge, P., & Wichers, M. (2017).

- The impact of treatments for depression on the dynamic network structure of mental states: Two randomized controlled trials. *Scientific Reports*, 7, 46523.
- Steele, R. J., & Raftery, A. E. (2010). Performance of bayesian model selection criteria for gaussian mixture models. *Frontiers of Statistical Decision Making and Bayesian Analysis*, 2, 113–130.
- Stiles, W. B., Leach, C., Barkham, M., Lucock, M., Iveson, S., Shapiro, D. A., ... Hardy, G. E. (2003). Early sudden gains in psychotherapy under routine clinic conditions: Practice-based evidence. *Journal of Consulting and Clinical Psychology*, 71(1), 14.
- Strogatz, S. H. (2015). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Reading, MA.
- Tong, H., & Lim, K. (1980). Threshold autoregression, limit cycles and cyclical data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 42(3), 245–268.
- van de Leemput, I. A., Wichers, M., Cramer, A. O., Borsboom, D., Tuerlinckx, F., Kuppens, P., ... others (2014). Critical slowing down as early warning for the onset and termination of depression. *Proceedings of the National Academy of Sciences*, 111(1), 87–92.
- Visser, I., & Speekenbrink, M. (2010). depmixS4: An R package for hidden markov models. *Journal of Statistical Software*, 36(7), 1–21. Retrieved from <http://www.jstatsoft.org/v36/i07/>
- Walker, M. (2017). *Why we sleep: Unlocking the power of sleep and dreams*. New York: Simon and Schuster.
- Warren, K. (2002). Thresholds and the abstinence violation effect: A nonlinear dynamical model of the behaviors of intellectually disabled sex offenders. *Journal of Interpersonal Violence*, 17(11), 1198–1217.
- Wichers, M., Schreuder, M. J., Goekoop, R., & Groen, R. N. (2019). Can we predict the direction of sudden shifts in symptoms? transdiagnostic implications from a complex systems perspective on psychopathology. *Psychological medicine*, 49(3), 380–387.
- Wichers, M., Wigman, J., & Myin-Germeys, I. (2015). Micro-level affect dynamics in psychopathology viewed from complex dynamical system theory. *Emotion Review*, 7(4), 362–367.
- Wolfram Research, Inc. (n.d.). *Mathematica, Version 12.0*. (Champaign, IL, 2019)

A Determine Fixed Points of Bistable System

In this section we show how to compute the fixed points of the deterministic part of our model, which we report in Section 2.1. The fixed points of a set of differential equations is found by setting all equations to zero and solving that system. In our case this means solving the nonlinear system of equations:

$$\begin{aligned} 0 &= r_1 x_1 + \sum_{j=1}^4 C_{1,j} x_j x_1 + a_1 \\ 0 &= r_2 x_2 + \sum_{j=1}^4 C_{2,j} x_j x_2 + a_2 \\ 0 &= r_3 x_3 + \sum_{j=1}^4 C_{3,j} x_j x_3 + a_3 \\ 0 &= r_4 x_4 + \sum_{j=1}^4 C_{4,j} x_j x_4 + a_4 \end{aligned}$$

Since we have $r_1, r_2 = 1$ and $a = [1.6, 1.6, 1.6, 1.6]$ in all studied situations, we fill in those values and write out the summation:

$$\begin{aligned} 0 &= x_1 + C_{1,1} x_1 x_1 + C_{1,2} x_1 x_2 + C_{3,1} x_1 x_3 + C_{4,1} x_1 x_4 + 1.6 \\ 0 &= x_2 + C_{2,1} x_2 x_1 + C_{2,2} x_2 x_2 + C_{2,3} x_2 x_3 + C_{2,4} x_2 x_4 + 1.6 \\ 0 &= r_3 x_3 + C_{3,1} x_3 x_1 + C_{3,2} x_3 x_2 + C_{3,3} x_3 x_3 + C_{3,4} x_3 x_4 + 1.6 \\ 0 &= r_4 x_4 + C_{4,1} x_4 x_1 + C_{4,2} x_4 x_2 + C_{4,3} x_4 x_3 + C_{4,4} x_4 x_4 + 1.6 \end{aligned}$$

We can exploit the symmetries in r and C to simplify finding the fixed points. The derivatives of x_1 and x_2 are actually identical, and the derivatives of x_3 and x_4 are identical. Thus also their integrals are identical. Thus, we can substitute x_1 into x_2 , and x_3 into x_4 to arrive at a simpler 2-dimensional system. Making the substitutions, and filling in the parameter values, the differential equations then reduce to

$$\begin{aligned} 0 &= 1x_1 - 0.2x_1^2 + 0.04x_1^2 - 0.4x_1x_2 + 1.6 \\ 0 &= r_3x_3 - 0.2x_3^2 - 0.4x_3x_1 + 0.04x_3^2 + 1.6 \end{aligned}$$

where r_3 is the stress level for which the fixed points should be computed.

We now solve these systems for a number of stress values (r_3) using Mathematica ([Wolfram Research, Inc., n.d.](#)). This way, we computed the following fixed points shown in Table 5, which are displayed in panel (a) of Figure 1 in Section 2.1:

Stress	Healthy:PE	Healthy:NE	Unhealthy:PE	Unhealthy:NE	Unstable:PE	Unstable:NE
0.90	5.28	1.15				
0.91	5.26	1.16				
0.91	5.24	1.17				
0.92	5.22	1.18				
0.93	5.19	1.19				
0.93	5.17	1.20				
0.94	5.15	1.22				
0.95	5.12	1.23				
0.95	5.10	1.24				
0.96	5.08	1.26				
0.90	5.28	1.15				
0.91	5.26	1.16				
0.91	5.24	1.17				
0.92	5.22	1.18				
0.93	5.19	1.19				
0.93	5.17	1.20				
0.94	5.15	1.22				
0.95	5.12	1.23				
0.95	5.10	1.24				
0.96	5.07	1.26	1.83	3.96	2.03	3.66
0.97	5.05	1.27	1.66	4.25	2.25	3.30
0.97	5.02	1.29	1.57	4.42	2.39	3.22
0.98	4.99	1.31	1.50	4.56	2.50	3.10
0.99	4.96	1.33	1.45	4.69	2.61	2.99
0.99	4.92	1.34	1.40	4.79	2.71	2.89
1.00	4.89	1.36	1.36	4.89	2.80	2.80
1.01	4.85	1.39	1.33	4.98	2.90	2.72
1.01	4.80	1.41	1.30	5.06	2.99	2.64
1.02	4.76	1.44	1.27	5.15	3.09	2.56
1.03	4.71	1.47	1.24	5.23	3.19	2.48
1.03	4.65	1.50	1.22	5.30	3.29	2.40
1.04	4.59	1.54	1.19	5.38	3.40	2.32
1.05	4.51	1.58	1.17	5.45	3.52	2.23
1.05	4.41	1.65	1.15	5.52	3.67	2.12
1.06	4.24	1.75	1.13	5.59	3.87	1.98
1.06			1.12	5.63		
1.07			1.11	5.66		
1.07			1.09	5.72		
1.08			1.08	5.79		
1.09			1.06	5.85		
1.09			1.04	5.91		
1.10			1.03	5.98		

Table 5: Fixed points of the emotion model for different values of stress (rows), rounded to two decimals. The 2nd and 3rd columns refer to the fixed points of the healthy fixed points for positive and negative emotions; the 4th and 5th columns refer to the unhealthy fixed points; and the last two columns refer to the unstable fixed point.

B Mean-Switching Hidden Markov Model

In this appendix we provide additional details with respect to the specification of the mean-switching Hidden Markov Model, and using model selection to obtain the number of components, described in Section 3.2

B.1 Model Specification

The mean-switching Hidden Markov Model is denoted

$$P(X, S | \mu, \sigma) = \pi_i \mathcal{N}(X_1) \prod_{t=1}^{T-1} A_{ji} \mathcal{N}(X_{t+1}),$$

where $X = \{X_1, \dots, X_T\}$ is a matrix of p -variate elements X_j , $S \in \{1, \dots, K\}^T$ is a vector of length T indicating the state at each time point, π_i is the probability of being in state $i \in \{1, \dots, K\}$, $A_{i,j}$ is the probability of transitioning from state i to state j , and μ, σ parameterize the multivariate Gaussian distribution $\mathcal{N}$ with zero covariances.

In Section 3.2 we chose $K = 2$ components, and fix the covariances of the Gaussian distribution to zero. Since we model four variables, this gives us 2×4 means and 2×4 standard deviations. The transition matrix A has three parameters since the last one is determined by the remaining three. Similarly, the marginal probabilities π_1, π_2 are determined by A and therefore do not count as additional parameters. We therefore fit a model with 19 freely estimated parameters.

B.2 Model Selection for Mean-Switching HMM

In Section 3.2 we inserted bistability as an assumption in the model by specifying that the HMM exhibits two states, and therefore the HMM does not provide us any evidence with respect to which number of states represents the data best. This can be done by performing model selection between HMMs with different numbers of states.

A popular way to select between mean-switching HMMs / Gaussian mixtures is the Bayesian Information Criterion (BIC) (Schwarz et al., 1978), because it has been shown to consistent in estimating Gaussian mixtures (Leroux, 1992), and has outperformed other information criteria (including the AIC) in simulations (Steele & Raftery, 2010). Here we fit HMMs with $K \in \{1, \dots, 10\}$ and report the BIC values in Figure 15:

Figure 15: The figure depicts BIC values as a function of the number of states K , for HMMs fitted to the ideal data.

We see that the BIC is highest for $K = 1$ and then decreases for larger K , however the change in BIC becomes less and less when adding additional states. Since we know from the true bistable system that the number of states is $K = 2$, we see that the BIC does not select the true number of states. The reason is that the BIC has been shown to be a consistent estimator of K if the data is generated from a Gaussian mixture. However, in the present case the data is generated from a bistable dynamical system. This failed attempt at model selection based on statistical models

again highlights the problems of using misspecified statistical models to make inferences about dynamical systems models.

C Data Generated from Estimated Models

In this Appendix we show data generated from estimated models for the time period of two weeks of the original time series.

C.1 Mean Switching Hidden Markov Model

Figure 16 displays a time series of two weeks generated from the Mean switching HMM estimated in Section 3.2:

Figure 16: A time series of two weeks generated from the HMM estimated in Section Section 3.2.

The generated time series looks similar to the original data in that it switches between the two fixed points at around (1,6) and (6,1). However, there are also differences. In the original data there are less switches that lead to a long-lived change in fixed point, but more switches that are very short-lived. Second, due to the form of the Mean-Switching HMM, there are no “intermediate” observations leading from one fixed point to the other. These observations exist in the original time series (see panel (b) in Figure 2).

C.2 First-order Vector Autoregressive (VAR(1)) model

Figure 17 displays a time series of two weeks generated from the VAR(1) model in Section 3.4:

Figure 17: A time series of two weeks generated from the VAR(1) model estimated in Section Section 3.4.

The generated data does not show bistability, which is expected because the VAR(1) model exhibits only a single fixed point. What looks approximately like oscillating behaviour is a result of the high auto-regressive effects present in the estimated VAR(1) model: given a stochastic input, the high auto-regressive effects ensure that the system is slow to eventually return to equilibrium. This oscillating behaviour is also evident in the eigenvalues of Φ , which consist of one complex conjugate pair (Strogatz, 2015).

C.3 Threshold VAR(1) Model

Figure 18 displays a time series of two weeks generated from the TVAR(1) model in Section 3.5:

Figure 18: A time series of two weeks generated from the TVAR(1) model estimated in Section 3.5.

The data generated from the TVAR(1) model looks similar to the original time series in that the position of the fixed points and the variance around them is very similar. However, the system seems to switch less often between states, and similarly to the data generated from the HMM above, there are much fewer observations on the transitions between states.

D Residual Partial Correlations of TVAR(1) Model

Figure 19: Residual partial correlation networks for both regimes in the TVAR model described in Section 3.5 in the main text.

E Differential Equation Model Building Details

In this appendix we present additional information relating to the two-step DE model building procedure utilized in Sections 3 and 4. This includes details on how model fit is computed, as well as full model fit results and parameter estimates for each of the models described in the main text.

E.1 Evaluating Model Fit

The fit of each model is evaluated with the mean out-of-bag explained variance, referred to throughout as R^2 . This metric is calculated using 10-fold cross-validation. First, the given dataset is randomly partitioned into ten mutually exclusive training and test sets. Second, for each partitioned dataset, regression models A through G, (defined by the expression in the second column of Table 6) are fit to the training set four times, once each of the four outcome variables $dx_i/dt, \forall i \in \{1, 2, 3, 4\}$. Third, the resulting parameters are then used to predict the values of the outcome variable in the test set dx_i/dt . The variance of the resulting residuals $VAR(dx_i/dt - \hat{dx}_i/dt)$ is then divided by the variance of the outcome variable in the test set, $VAR(dx_i/dt)$ yielding an out-of-bag variance explained for variable i based on model m in partition k , $R_{i,k,m}^2$. Averaging the explained variance across each of the partitions yields an average explained variance for variable i in model m , $R_{i,m}^2$, and averaging this number across all four outcome variables yields the average out-of-bag explained variance for model m .

E.2 Ideal Data

In Table 6 we show the fit of models A through G for the ideal dataset analysis in Section 3.6. In Table 7 we show the full parameter estimates, standard errors and p -values for the selected model, Model C.

Model	$\frac{dx_{i,t}}{dt} \sim a + r_i x_i + \dots$	q	R^2
A	$\sum_{j \neq i} r_j x_j$	5	.04464
B	$\sum_{j \neq i} R_{ij} x_j + \sum_j^p C_{ij} x_j x_i$	9	.06874
C	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k$	15	.06870
D	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k + \sum_j^p \gamma_j x_j^3$	19	.06871
E	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k + \sum_j^p \gamma_j x_j^3 + \sum_{j \neq k \neq l}^p \zeta_j(x_j x_k x_l)$	23	.06870
F	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k + \sum_j^p \gamma_j x_j^3 + \sum_{(j,k,l)}^p \zeta_j(x_j x_k x_l)$	35	.06860
G	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)}^p \beta_j x_j x_k + \sum_j^p \gamma_j x_j^3 + \sum_{(j,k,l)}^p \zeta_j(x_j x_k x_l) + \sum_{(j,k,l,m)}^p \eta_j(x_j x_k x_l x_m)$	70	.06846

Table 6: Model fit results for each of the seven models described in text in Section 3.6 for the ideal dataset. The second column gives the model equation for each variable, q denotes the number of parameters estimated per univariate regression model, and the final column indicates R^2 , the explained variance, as calculated based on the prediction error on a hold-out set, using 10-fold cross-validation.

	dx_1/dt			dx_2/dt			dx_3/dt			dx_4/dt		
	Est	SE	p	Est	SE	p	Est	SE	p	Est	SE	p
a	1.40	0.13	<.01	1.37	0.12	<.01	1.25	0.12	<.01	1.27	0.12	<.01
x_1	0.88	0.05	<.01	0.03	0.02	0.19	-0.02	0.02	0.33	0.04	0.02	0.05
x_2	0.02	0.02	0.34	0.95	0.05	<.01	0.05	0.02	0.02	-0.01	0.02	0.57
x_3	-0.01	0.02	0.72	0.01	0.02	0.68	0.96	0.05	<.01	0.08	0.02	<.01
x_4	0.01	0.02	0.51	<.01	0.02	0.80	0.04	0.02	0.10	0.91	0.05	<.01
$x_1 \times x_1$	-0.18	0.01	<.01	-	-	-	-	-	-	-	-	-
$x_1 \times x_2$	0.04	0.01	<.01	0.03	0.01	<.01	-	-	-	-	-	-
$x_1 \times x_3$	-0.17	0.01	<.01	-	-	-	-0.18	0.01	<.01	-	-	-
$x_1 \times x_4$	-0.18	0.01	<.01	-	-	-	-	-	-	-0.19	0.01	<.01
$x_2 \times x_2$	-	-	-	-0.19	0.01	<.01	-	-	-	-	-	-
$x_2 \times x_3$	-	-	-	-0.19	0.01	<.01	-0.19	0.01	<.01	-	-	-
$x_2 \times x_4$	-	-	-	-0.19	0.01	<.01	-	-	-	-0.18	0.01	<.01
$x_3 \times x_3$	-	-	-	-	-	-	-0.19	0.01	<.01	-	-	-
$x_3 \times x_4$	-	-	-	-	-	-	0.03	0.01	<.01	0.02	0.01	<.01
$x_4 \times x_4$	-	-	-	-	-	-	-	-	-	-0.18	0.01	<.01

Table 7: Full parameter estimates, standard errors and p-values for Model B in Section 3.6, for the DE model fit to ideal data.

E.3 ESM Data

In Table 8 we show the fit of models A through G for the emulated ESM dataset analysis, from Section 4.4. In Table 9 we show the full parameter estimates, standard errors and p -values for the selected model, Model G.

Model	$\frac{dx_{i,t}}{dt} \sim a + r_i x_i + \dots$	q	R^2
A	$\sum_{j \neq i} r_j x_j$	5	0.13991
B	$\sum_{j \neq i} R_{ij} x_j + \sum_j C_{ij} x_j x_i$	9	0.16827
C	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)} \beta_j x_j x_k$	15	0.16928
D	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)} \beta_j x_j x_k + \sum_j \gamma_j x_j^3$	19	0.19455
E	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)} \beta_j x_j x_k + \sum_j \gamma_j x_j^3 + \sum_{j \neq k \neq l} \zeta_j(x_j x_k x_l)$	23	0.19801
F	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)} \beta_j x_j x_k + \sum_j \gamma_j x_j^3 + \sum_{(j,k,l)} \zeta_j(x_j x_k x_l)$	35	0.19940
G	$\sum_{j \neq i} R_{ij} x_j + \sum_{(j,k)} \beta_j x_j x_k + \sum_j \gamma_j x_j^3 + \sum_{(j,k,l)} \zeta_j(x_j x_k x_l) + \sum_{(j,k,l,m)} \eta_j(x_j x_k x_l x_m)$	70	0.20420

Table 8: Model fit results for each of the seven models described in text, for the ESM time series, described in Section 4. The second column gives the model equation for each variable, q denotes the number of parameters estimated per univariate regression model. The final two columns indicate R^2 , the explained variance, as calculated based on the prediction error on a hold-out set, using 10-fold cross-validation, for the snapshot ESM data and the mean-aggregated ESM data, respectively. R^2 for Model G in the mean-aggregated ESM data case was not available due to multicollinearity problems encountered when fitting the model.

	dx_1/dt			dx_2/dt			dx_3/dt			dx_4/dt		
	Est	SE	p	Est	SE	p	Est	SE	p	Est	SE	p
(Intercept)	-0.07	0.73	0.92	-0.16	0.73	0.82	0.31	0.74	0.68	0.20	0.74	0.78
x_1	0.19	0.31	0.53	0.16	0.31	0.61	-0.09	0.31	0.78	-0.12	0.31	0.71
x_2	0.07	0.30	0.80	0.18	0.30	0.54	-0.29	0.30	0.34	-0.21	0.30	0.48
x_3	-0.06	0.31	0.86	0.01	0.31	0.96	-0.05	0.31	0.88	-0.07	0.31	0.82
x_4	-0.01	0.31	0.97	-0.04	0.31	0.89	0.01	0.31	0.98	0.10	0.31	0.74
$x_1 \times x_1$	-0.02	0.06	0.78	-0.01	0.06	0.87	-0.03	0.06	0.65	-0.01	0.06	0.81
$x_1 \times x_2$	-0.11	0.09	0.21	-0.10	0.09	0.22	0.14	0.09	0.10	0.13	0.09	0.13
$x_1 \times x_3$	-0.02	0.10	0.86	-0.04	0.10	0.66	<.01	0.10	0.99	0.02	0.10	0.80
$x_1 \times x_4$	-0.05	0.10	0.59	<.01	0.10	0.99	0.01	0.10	0.94	<.01	0.10	0.98
$x_2 \times x_2$	0.01	0.06	0.83	-0.02	0.06	0.76	0.02	0.06	0.72	0.01	0.06	0.82
$x_2 \times x_3$	<.01	0.10	0.98	-0.01	0.10	0.94	0.05	0.10	0.58	0.06	0.10	0.57
$x_2 \times x_4$	-0.01	0.10	0.95	-0.05	0.10	0.60	0.06	0.10	0.57	0.02	0.10	0.86
$x_3 \times x_3$	0.02	0.06	0.68	0.01	0.06	0.87	-0.01	0.06	0.87	0.02	0.06	0.73
$x_3 \times x_4$	0.01	0.09	0.92	0.01	0.09	0.92	0.01	0.09	0.94	-0.04	0.09	0.62
$x_4 \times x_4$	0.02	0.06	0.80	0.03	0.06	0.68	-0.03	0.06	0.69	-0.03	0.06	0.68
$x_1 \times x_1 \times x_1$	<.01	0.01	0.71	<.01	0.01	0.92	0.01	0.01	0.38	0.01	0.01	0.33
$x_1 \times x_1 \times x_2$	0.02	0.01	0.19	0.01	0.01	0.48	-0.01	0.01	0.33	-0.02	0.01	0.15
$x_1 \times x_1 \times x_3$	<.01	0.01	0.95	0.01	0.01	0.61	0.01	0.01	0.64	<.01	0.01	0.82
$x_1 \times x_1 \times x_4$	<.01	0.01	0.78	<.01	0.01	0.71	<.01	0.01	0.73	<.01	0.01	0.81
$x_1 \times x_2 \times x_2$	<.01	0.01	0.87	0.01	0.01	0.49	-0.01	0.01	0.39	<.01	0.01	0.76
$x_1 \times x_2 \times x_3$	0.01	0.02	0.58	0.01	0.02	0.69	-0.02	0.02	0.32	-0.02	0.02	0.27
$x_1 \times x_2 \times x_4$	0.02	0.02	0.25	0.02	0.02	0.22	-0.03	0.02	0.21	-0.02	0.02	0.34
$x_1 \times x_3 \times x_3$	<.01	0.01	0.79	<.01	0.01	0.96	<.01	0.01	0.81	<.01	0.01	0.98
$x_1 \times x_3 \times x_4$	0.01	0.02	0.73	<.01	0.02	0.80	<.01	0.02	0.99	<.01	0.02	0.99
$x_1 \times x_4 \times x_4$	<.01	0.01	0.98	-0.01	0.01	0.62	<.01	0.01	0.80	<.01	0.01	0.75
$x_2 \times x_2 \times x_2$	<.01	0.01	0.92	<.01	0.01	0.96	<.01	0.01	0.98	<.01	0.01	0.82
$x_2 \times x_2 \times x_3$	<.01	0.01	1.00	<.01	0.01	0.87	<.01	0.01	0.79	<.01	0.01	0.79
$x_2 \times x_2 \times x_4$	-0.01	0.01	0.70	<.01	0.01	0.90	<.01	0.01	0.89	<.01	0.01	0.96
$x_2 \times x_3 \times x_3$	<.01	0.01	0.92	<.01	0.01	0.93	<.01	0.01	0.89	-0.01	0.01	0.50
$x_2 \times x_3 \times x_4$	-0.01	0.02	0.70	<.01	0.02	0.79	<.01	0.02	0.83	0.01	0.02	0.59
$x_2 \times x_4 \times x_4$	<.01	0.01	0.91	0.01	0.01	0.63	<.01	0.01	0.83	<.01	0.01	0.78
$x_3 \times x_3 \times x_3$	<.01	0.01	0.95	<.01	0.01	0.93	<.01	0.01	0.97	<.01	0.01	0.51
$x_3 \times x_3 \times x_4$	-0.01	0.01	0.40	-0.01	0.01	0.68	0.01	0.01	0.63	0.01	0.01	0.39
$x_3 \times x_4 \times x_4$	0.01	0.01	0.44	<.01	0.01	0.76	-0.01	0.01	0.56	<.01	0.01	0.73
$x_4 \times x_4 \times x_4$	-0.01	0.01	0.36	-0.01	0.01	0.42	0.01	0.01	0.33	0.01	0.01	0.40
$x_1 \times x_1 \times x_1 \times x_1$	<.01	<.01	0.94	<.01	<.01	0.71	<.01	<.01	0.91	<.01	<.01	0.48
$x_1 \times x_1 \times x_1 \times x_2$	<.01	<.01	0.82	<.01	<.01	0.55	<.01	<.01	0.52	<.01	<.01	0.92
$x_1 \times x_1 \times x_1 \times x_3$	<.01	<.01	0.99	<.01	<.01	0.74	<.01	<.01	0.46	<.01	<.01	0.40
$x_1 \times x_1 \times x_1 \times x_4$	<.01	<.01	0.80	<.01	<.01	0.78	<.01	<.01	0.83	<.01	<.01	0.88
$x_1 \times x_1 \times x_2 \times x_2$	<.01	<.01	0.25	<.01	<.01	0.23	<.01	<.01	0.15	<.01	<.01	0.36
$x_1 \times x_1 \times x_2 \times x_3$	<.01	<.01	0.70	<.01	<.01	0.66	<.01	<.01	0.43	<.01	<.01	0.22
$x_1 \times x_1 \times x_2 \times x_4$	<.01	<.01	0.33	<.01	<.01	0.76	<.01	<.01	0.84	<.01	<.01	0.81
$x_1 \times x_1 \times x_3 \times x_3$	<.01	<.01	0.37	<.01	<.01	0.44	<.01	<.01	0.19	<.01	<.01	0.17
$x_1 \times x_1 \times x_3 \times x_4$	<.01	<.01	0.27	<.01	<.01	0.16	<.01	<.01	0.36	<.01	<.01	0.19
$x_1 \times x_1 \times x_4 \times x_4$	<.01	<.01	0.47	<.01	<.01	0.12	<.01	<.01	0.23	<.01	<.01	0.17
$x_1 \times x_2 \times x_2 \times x_2$	<.01	<.01	0.35	<.01	<.01	0.51	<.01	<.01	0.43	<.01	<.01	0.52
$x_1 \times x_2 \times x_2 \times x_3$	<.01	<.01	0.96	<.01	<.01	0.92	<.01	<.01	0.97	<.01	<.01	0.80
$x_1 \times x_2 \times x_2 \times x_4$	<.01	<.01	0.82	<.01	<.01	0.35	<.01	<.01	0.28	<.01	<.01	0.44
$x_1 \times x_2 \times x_3 \times x_3$	<.01	<.01	0.28	<.01	<.01	0.18	<.01	<.01	0.10	<.01	<.01	0.04
$x_1 \times x_2 \times x_3 \times x_4$	<.01	<.01	0.39	<.01	<.01	0.16	<.01	<.01	0.27	<.01	<.01	0.12
$x_1 \times x_2 \times x_4 \times x_4$	<.01	<.01	0.13	<.01	<.01	0.05	<.01	<.01	0.07	<.01	<.01	0.06
$x_1 \times x_3 \times x_3 \times x_3$	<.01	<.01	0.76	<.01	<.01	0.90	<.01	<.01	0.99	<.01	<.01	0.98
$x_1 \times x_3 \times x_3 \times x_4$	<.01	<.01	0.96	<.01	<.01	0.99	<.01	<.01	0.66	<.01	<.01	0.80
$x_1 \times x_3 \times x_4 \times x_4$	<.01	<.01	0.82	<.01	<.01	0.94	<.01	<.01	0.74	<.01	<.01	0.88
$x_1 \times x_4 \times x_4 \times x_4$	<.01	<.01	0.75	<.01	<.01	0.53	<.01	<.01	0.52	<.01	<.01	0.59
$x_2 \times x_2 \times x_2 \times x_2$	<.01	<.01	0.52	<.01	<.01	0.57	<.01	<.01	0.56	<.01	<.01	0.50
$x_2 \times x_2 \times x_2 \times x_3$	<.01	<.01	0.76	<.01	<.01	0.60	<.01	<.01	0.64	<.01	<.01	0.49
$x_2 \times x_2 \times x_2 \times x_4$	<.01	<.01	0.71	<.01	<.01	0.68	<.01	<.01	0.63	<.01	<.01	0.69
$x_2 \times x_2 \times x_3 \times x_3$	<.01	<.01	0.47	<.01	<.01	0.32	<.01	<.01	0.30	<.01	<.01	0.44
$x_2 \times x_2 \times x_3 \times x_4$	<.01	<.01	0.57	<.01	<.01	0.26	<.01	<.01	0.18	<.01	<.01	0.29
$x_2 \times x_2 \times x_4 \times x_4$	<.01	<.01	0.41	<.01	<.01	0.43	<.01	<.01	0.37	<.01	<.01	0.36
$x_2 \times x_3 \times x_3 \times x_3$	<.01	<.01	0.64	<.01	<.01	0.80	<.01	<.01	0.89	<.01	<.01	0.31
$x_2 \times x_3 \times x_3 \times x_4$	<.01	<.01	0.28	<.01	<.01	0.48	<.01	<.01	0.76	<.01	<.01	0.36
$x_2 \times x_3 \times x_4 \times x_4$	<.01	<.01	0.55	<.01	<.01	0.77	<.01	<.01	0.74	<.01	<.01	0.77
$x_2 \times x_4 \times x_4 \times x_4$	<.01	<.01	0.68	<.01	<.01	0.95	<.01	<.01	0.86	<.01	<.01	0.91
$x_3 \times x_3 \times x_3 \times x_3$	<.01	<.01	0.70	<.01	<.01	0.99	<.01	<.01	0.95	<.01	<.01	0.64
$x_3 \times x_3 \times x_3 \times x_4$	<.01	<.01	0.44	<.01	<.01	0.83	<.01	<.01	0.96	<.01	<.01	0.83
$x_3 \times x_3 \times x_4 \times x_4$	<.01	<.01	0.79	<.01	<.01	0.96	<.01	<.01	0.77	<.01	<.01	0.72
$x_3 \times x_4 \times x_4 \times x_4$	<.01	<.01	0.75	<.01	<.01	0.76	<.01	<.01	0.48	<.01	<.01	0.54
$x_4 \times x_4 \times x_4 \times x_4$	<.01	<.01	0.30	<.01	<.01	0.33	<.01	<.01	0.18	<.01	<.01	0.23

Table 9: Full parameter estimates, standard errors and p-values for Model G in Section 4.4, for the DE model fit to the emulated ESM data.

F Additional Model Results from ESM Time Series

In this appendix, we provide additional figures to visualize the results of the statistical models fit to the ESM time series in Section 4.

Figure 20: The histograms of the emotion intensity of the four modeled emotions Cheerful, Content, Anxious and Sad, for the ESM data

Figure 21: Panel (a) shows the relationship between Content and Cheerful, two emotions with the same valence, at the same time point. The red line indicates the best fitting regression model, for ESM time series. Similarly, panel (b) shows the relationship between Anxious and Content, two emotions with different valence. Panel (c) displays the correlation matrix as a network, and panel (d) displays the partial correlation matrix as a network.

Figure 22: Time series of the four emotion variables, also shown in panel (a) of Figure 2, with background color indicating whether a given time point is assigned to the first or second component of the mean-switching HMM estimated from the ESM dataset.