Modeling Change in Networks

Haslbeck, J.M.B.; Ryan, O.; van der Maas, H.L.J.; Waldorp, L.J.

DOI
10.4324/9781003111238-14

Publication date
2022

Document Version
Final published version

Published in
Network Psychometrics with R

Chapter 11

Modeling Change in Networks

Jonas M. B. Haslbeck[1], Oisín Ryan[2], Han L. J. van der Maas[1], & Lourens J. Waldorp[1]

1. University of Amsterdam, Department of Psychology
2. Utrecht University, Department of Methodology & Statistics

11.1 Introduction

In the previous two chapters, we considered models for the relationships between variables over time (temporal effects) and within the same measurement occasion (contemporaneous effects). A key assumption of these models is that their parameters do not change across time. If only the means change across time, this problem can in some situations be mitigated by detrending the time series (see also Chapter 10). However, in many situations changes in parameters over time might be the very thing we are interested in. In this chapter we introduce time-varying network models, which allow us to explicitly model such changes.

The difference between parameters that stay the same across time and time-varying parameters is illustrated in Figure 11.1, which displays the (lag-0) partial correlations of a hypothetical time-varying Gaussian graphical model (GGM) capturing the relationships between the variables 'anxious,' 'worried,' and 'irritated' at the same time point in a time series over 100 days. The top panel displays the partial correlations as a function of time. We see that the partial correlation between 'anxious' and 'worried' is 0.4 at the beginning of the time series and monotonically decreases to zero at the end of the time series. In contrast, the partial correlations between 'worried' and 'irritated,' and 'anxious' and 'irritated,' are constant (0.2 and 0) across the entire time series.

Cite this chapter as:

Haslbeck, J. M. B., Ryan, O., van der Maas, H. L. J., & Waldorp, L. J. (2022). Chapter 11. Modeling change in networks. In Isvoranu, A. M., Epskamp, S., Waldorp, L. J., & Borsboom, D. (Eds.). *Network psychometrics with R: A guide for behavioral and social scientists*. Routledge, Taylor & Francis Group.

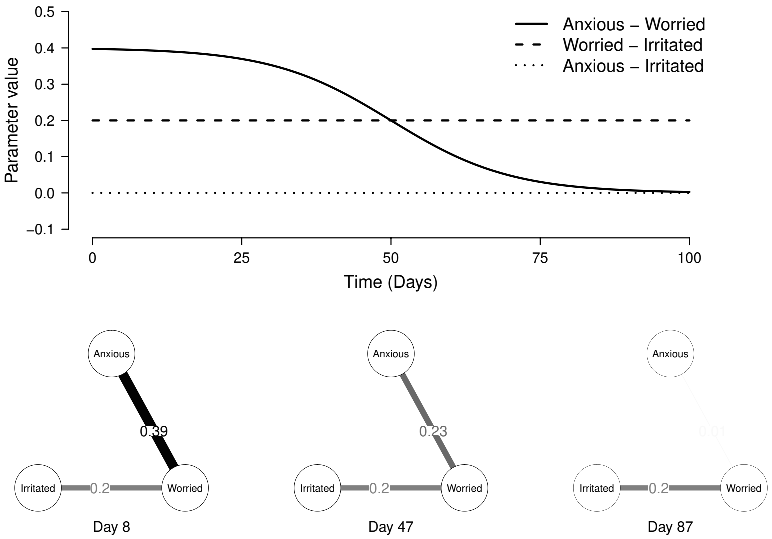

Figure 11.1. Illustration of a time-varying network model over a time period of 100 days. Top panel: The partial correlations between the three variables 'anxious,' 'worried,' and 'irritated' as a function of time. Lower panel: The corresponding partial correlation networks at days 8, 47, and 87.

The lower panel of Figure 11.1 displays the corresponding partial correlation networks on day 8, 47, and 87. Consistent with the top panel, the partial correlation between 'anxious' and 'worried' decreases over the course of the time series, while the other two partial correlations stay the same. This shows that not every parameter in a time-varying model needs to change across time. If the parameters of a model do not change across time, the model is often referred to as a stationary model.[1] Note that in Figure 11.1 we only displayed the three partial correlations as a function of time, and no other parameters of the GGM.

Time-varying analyses are relevant when we are interested in change within an individual. For example, take the transition of an individual into a state in which they are diagnosed with a mental disorder. From the network perspective of psychopathology mental disorders arise from direct interactions between symptoms and vulnerability is closely connected to the nature of these interactions (Borsboom, 2017). For example, it has been suggested that more dense symptom networks are associated with depression diagnoses (e.g., see Cramer et al., 2016; Pe et al., 2015; van Borkulo et al., 2015, but also see Schweren et al., 2018). If this was true, we would expect that the density (or global

[1]Technically, a process is stationary if its moments are the same in all parts of the time series. There are some cases in which the parameters of a model do not change across time, but the moments of the distribution are not stationary. An example is a vector auto-regressive (VAR) process with eigenvalues outside the limit circle. However, for the cases discussed in this chapter, non-time-varying parameters and stationarity can be used interchangeably.

strength) of individual symptom networks tends to increase when transitioning from a healthy state to a state with a diagnosis. This prediction could be tested with time-varying models. Time-varying models can also be used to study the impact of unplanned or planned events, such as negative life events or interventions. In addition, time-varying models are central to early warning signals (EWS) of critical transitions, which may allow one to anticipate transitions in alternative states (van de Leemput et al., 2014, however, also see Dablander et al., 2021). The methods to detect EWS currently used in the literature are typically based on univariate measures such as variances and autocorrelations. However, most systems of interest consist of several variables. Time-varying models are the natural multivariate extension of such EWS detection procedures.

Other instructive examples come from the field of developmental psychology. It is possible that a developmental transition, say from non-conserver to conserver as measured by Piaget's conservation task, is solely a change in state and not in the underlying network structure, but developmental theory usually expects that such changes are accompanied by structural changes in the relations between variables. In Dolan and van der Maas (1998) and Schmittmann et al. (2005) it has been shown, with the fit of mixture factor and latent Markov models, that the correlation structure in answers to conservation anticipation items is qualitatively different between non-conservers and conservers. Similar changes in the correlations can be expected for other Piagetian tasks such as the balance scale task (van der Maas & Jansen, 2003). Another example are dyadic interactions such as between parent and child. Such interactions are sometimes time varying (Bringmann et al., 2018). It is interesting to note that the mutualism network model of general intelligence (van der Maas et al., 2006) is a developmental model in which the relations between nodes (cognitive functions) are not varying over time. The main reason for this was that time-invariant relations were sufficient to explain the positive manifold, the main phenomenon to be explained in intelligence research. A newly proposed network model of intelligence is time-varying as it models intelligence as growing networks (Savi et al., 2019).

On a more abstract level, time-varying models allow one to capture the variability of parameters across time. This is important even if the time-varying parameters are initially not related to any events or covariates, because variability is an empirical fact that needs to be explained by a formal theory explaining the phenomenon at hand (Borsboom et al., 2021; Haslbeck, Ryan, et al., 2021). For example, if the relation between 'anxious' and 'worried' fluctuates significantly in depressed individuals, a theory about anxiety and worrying in depressed individuals would need to account for these fluctuations.

11.2 Time-varying network models

A time-varying model is a model that includes parameters that vary across time. This means that we can turn any of the stationary models discussed in the previous chapters such as GGMs, Ising models, mixed graphical models (MGMs) or VAR models into a time-varying model by making its parameters a function of time. For example, for the case of the GGM used above, we would give both the mean vector $\boldsymbol{\mu}_t$ and the partial correlation matrix $\boldsymbol{\Omega}_t$ an index for time, where t is an ordered variable capturing the time from the start until the end of the time series. That way, we obtain a time-varying

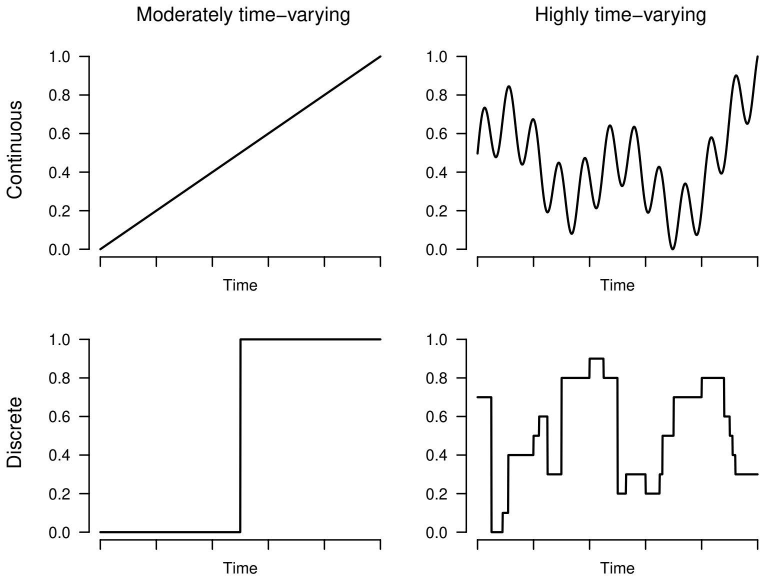

Figure 11.2. Examples of time-varying parameter functions. Top left: moderately time-varying linear function; top right: highly time-varying continuous function; bottom right: highly time-varying step function; bottom left: moderately time-varying function of discrete jumps (with multiple states). The value of the time-varying parameter is on the y-axis.

Gaussian distribution. In such a time-varying network model, each parameter can depend on time in a different way. Some parameters may be constant across time, as we have seen in the example in Figure 11.1, and others may vary in various ways across time. While a time-varying network model always consists of many (time-varying) parameters, we will now focus our discussion on a single parameter. This parameter could be a time-varying partial correlation or mean in a GGM, a time-varying autoregressive or cross-lagged effect in a VAR model, or a time-varying interaction between categorical variables in an MGM. Of course, there is an infinite number of ways in which a parameter can be a function of time. However, in order to develop some intuition for which time-varying models are easier or harder to estimate, we consider four illustrative examples in Figure 11.2 along two dimensions: first, whether the change is continuous (first row) or not (second row), and little change (first column) versus a lot of change (second column) across time.

The top left panel of Figure 11.2 shows a time-varying parameter that is linearly increasing over time, which is the simplest possible time-varying continuous function. But the parameter does not have to follow some simple function. The top right panel displays a parameter that also changes continuously across time, however, this parameter exhibits more change within a particular time interval. Parameters do not have to change continuously, however, but could instead jump between discrete states. The bottom left panel illustrates the simplest case of this type of time-varying model: the parameter stays the same for the first half of the time series, and then jumps to a new value, and remains there for the rest of the time series. Similarly to the continuous case, there can be many more discrete jumps giving rise to a highly time-varying parameter (see bottom right panel). Of course, the true time-varying model could also be a combination of continuous changes and discrete jumps. One's beliefs about whether the changes in parameters across time are best described by continuous or step-functions are important, because they will inform the type of method to be used to estimate the time-varying model. We will return to this question in the next section.

When estimating stationary models, for example a VAR model, a common strategy to obtain more stable estimates is to collect data for a longer period of time, for example three instead of two weeks. However, this strategy does not work for a time-varying model. If we are interested in how the parameters of the VAR model evolve during the first two weeks, then adding a third week does not help. Instead, we need to make sure that enough measurements are available *within* the first two weeks, that is, that the measurement frequency is high enough. How high the sampling frequency needs to be in the interval of interest depends on the form of the time-varying parameters. The simpler time-varying functions on the left of Figure 11.2 require fewer observations than estimating the more complex ones on the right. The exact way to determine how many observations are necessary to estimate the time-varying model with a given accuracy depends on the estimation method at hand. However, across methods the measurements need to be frequent enough to closely follow how the parameter changes across time. For example, it makes sense intuitively that no estimation method can capture the highly time-varying functions in the right panels of Figure 11.2 if there are only a handful of measurements available.

11.3 Estimating time-varying network models

To estimate time-varying network models, we need to make assumptions about how parameters vary across time. This typically involves assuming either that parameters change as a continuous function of time, or that the parameters are piece-wise constant functions across time (see Figure 11.2). If one had infinite data, this choice would not matter too much, because both classes of functions could approximate any function arbitrarily well. However, when estimating time-varying models from empirical data, the choice of assumption makes a big difference. For example, piece-wise constant functions will do poorly in approximating linear relationships, and continuous functions will be a poor approximation of discrete jumps. The choice of assumption therefore has to be carefully considered and should be based on knowledge about the studied phenomenon, either from previous analyses or from theoretical expectations. For example, body weight is probably changing continuously, while the type of errors in a developmental psychology task might jump from one state to another, when a child advances to a new strategy. In the remainder of this section we discuss how a time-varying model could be estimated under these two assumptions.

Methods assuming continuous change

Perhaps the most straightforward method to estimate a time-varying network model with continuously changing parameters is to use a moving window approach. These types of methods move a window with a certain width through the time series and estimate a stationary model in each of the windows. This way one obtains a number of 'local' models throughout the time series which together constitute a time-varying model. Typically, however, one does not use a fixed window, but rather uses all data points for the estimation of all local models and weighs them differently depending on how far away they are from the current 'local' model (e.g., Kolar et al., 2010; Zhou et al., 2010). The key parameter in this method is the width (or bandwidth) of the kernel-function that determines the weights for each data point for a given 'local' model. Similar to the window width in the moving window approach, if the bandwidth is high, many observations around to the time point at which we estimate the model receive a high weight. This has the upside that we use more data and the estimates are less affected by sampling variance. However, it has the downside that we have less sensitivity to detect variations in parameters over time, since we aggregate larger parts of the time series together. Conversely, a low bandwidth offers a high sensitivity to detecting variations in parameters over time, however, the estimates will be affected more by sampling variance. It is thus important to strike a good balance between the two by selecting an appropriate bandwidth parameter for a given time series. Note that selecting an extremely large bandwidth will not allow any variation of parameters across time. Therefore, determining an appropriate bandwidth can also be seen as determining how time-varying (if at all) the parameters of the data generating model are. For a detailed explanation of this method we refer the reader to Haslbeck and Waldorp (2020) and Haslbeck, Bringmann, et al. (2021). Selecting the bandwidth is a model selection problem, which can be tackled with information criteria or data-driven approaches such as cross-validation (see also Chapter 7).

Another approach based on the assumption of continuity is to use a parametric model to explicitly model the time-varying parameters as a function of time. In order to allow for sufficient flexibility, this can be done using the generalized additive modeling (GAM) framework (Wood, 2006). This approach is easiest to imagine for a univariate model consisting only of a single intercept (here the mean of a variable). The GAM-based approach then finds the combination of functions that best fits the time-varying mean of this univariate time series. This idea can be extended to network models including several variables and interaction parameters. Bringmann et al. (2017) used this approach to estimate VAR models. Similar to the bandwidth in the kernel-method described above, the GAM-based method also requires a model selection procedure to determine the extent to which the parameters are time-varying. In this method, this is achieved by a penalized likelihood approach (Wood, 2006), in which increased complexity of the time-varying functions leads to a higher penalty. The penalty parameter in this procedure is selected using generalized cross-validation (Golub et al., 1979). Haslbeck, Bringmann, et al. (2021) provide an overview of the kernel-method and the GAM-based method, including a simulation study evaluating their performance in estimating time-varying VAR models in situations that are typical for psychological research.

Yet another way of estimating time-varying parameters using the assumption of continuity is by using the fused lasso (Hastie et al., 2015), which puts a penalty on differences in parameters in subsequent time points. This approach is implemented for the GGM in the SINGLE algorithm (Monti et al., 2014; Monti, 2014), and Gibbert (2017) provides an implementation of the (group) fused-lasso based method as presented by Gibberd and Nelson (2017).

If all parameters in the true time-varying model are continuous functions of time, and if we have enough measurements in the time interval of interest, we can estimate the time-varying model consistently (Robinson, 1989). However, if the time-varying parameters in the true model exhibit discrete jumps, the above methods will estimate the model incorrectly around those jumps. Thus, if we expect that the parameters are largely changing in discrete jumps across time, we should use this as our assumption to estimate the time-varying model.

Methods assuming discrete change

Many methods that model time-varying parameters which change in a discrete way over time do so using a discrete-valued 'state' or 'regime' variable. The simplest case is when we have an a-priori expectation that changes in parameter values will coincide with changes in the value of some observed (time-varying) variable. For instance, Bringmann et al. (2013) use a multilevel VAR model to analyze experience sampling data collected both before and after participants take part in a treatment program. To test whether the parameters of the VAR network change after treatment, they created a dummy state variable indicating whether observations were pre- or post-treatment, and include this as a moderator in their VAR model. If the regression coefficients for any of the dummy product terms are non-zero, then we have modeled a discrete change in those parameter values. Of course this general procedure can be extended to include multiple different states (for example, each with its own dummy variable).

The more general case is when we are unsure when or if a discrete change in parameters occurs over our window of observation. Although these methods are a little bit more complex than above, they typically also involve modeling a time-varying discrete 'state.' One popular class of models used for this purpose are hidden Markov models (HMM, e.g., Zucchini et al., 2017). HMMs model a system which changes between multiple states (also known as *components*), each with its own set of parameters, but where the state variable is unobserved and must be inferred from the data. As well as obtaining the parameters which describe the K different states, HMMs yield a simple so-called *Markov* model of how the state-switching behavior evolves over time in the form of transition probabilities, specifying how probable it is to switch from one state to the other (or remain at the current state) at the next time point. This allows one to obtain a time-varying model of the form displayed in the bottom row of Figure 11.2. In the classic form of the model, the HMM components represent Gaussian distributions, with each component having a potentially different (and so, time-varying) mean and covariance matrix. The HMM has also been extended such that the components themselves are time series models, often known as *Markov-Switching* time series models (Chow et al., 2018; Hamaker et al., 2010, 2016; Lu et al., 2019).

In the time series literature, models whose parameters change in a discrete fashion over time are often referred to as *regime-switching models* (Hamilton, 1989; Kim & Nelson, 1999) and this class of models includes many different ways of modeling the unobserved state/regime variable. One simple but popular method is to model the state of the system as determined by a threshold variable: the parameters of the time series model change in a discrete manner depending on whether the threshold variable falls above or below a threshold value at a given point in time (Tong & Lim, 1980). The threshold variable itself can be anything but must be specified by the researcher, while the threshold value is estimated from the data. Haslbeck, Ryan, et al. (2021) illustrate the application of threshold-VAR models to estimate time-varying VAR networks, where the threshold variable is one of the four variables included in the VAR model itself. See also Hamaker et al. (2010).

The appropriateness of these models depends on the assumptions we are willing to make about the mechanisms by which the parameters of our network model change over time. For example, threshold-based regime-switching models may be most appropriate when the researcher has some knowledge about which variable(s) determine the state of the system, or when sufficient data are available to reliably test multiple alternatives. HMMs and Markov-Switching models may be more appropriate when we believe the state-switching behavior to be governed by an unobserved variable, or by some unknown function of the observed variables. For an extended discussion and illustrative comparison of univariate HMM and threshold-based models fit to psychological time series data, we refer readers to Hamaker et al. (2010, 2016). Although many of the approaches mentioned here were originally developed for single-subject data, many multiple-subjects and multilevel extensions to these models also exist (Asparouhov et al., 2017; Chow et al., 2018; De Haan-Rietdijk et al., 2016; Ou et al., 2019).

Testing stationary vs. time-varying models

All method discussed above have some sort of mechanism built in that selects an optimal level of variation of parameters across time. The kernel-based approach accomplishes that by using a bandwidth that controls to what extent parameters can vary across time, and selects this hypertuningparameter with a cross-validation scheme. The GAM-based approach uses penalty for more complex models (i.e., more time-varying parameters), and selects the hypertuningparameter for this penalty using generalized cross-validation. And in the case of HMMs we perform model selection between different numbers of components, for example, using an information criterion.

In some situations, one might be interested in performing a formal test of whether a parameter or an entire model is stationary or time-varying. In the GAM-based method, it is possible to perform a significance test on individual parameters, however, there is no global test. For the HMMs, one selects the number of components, often with an information criterion like the Bayesian Information Criterion (BIC, Schwarz et al., 1978). If a single component is selected, we conclude that the time series is stationary, if there are two or more components, we conclude that it is time-varying. For all methods, including the kernel-based method, one can perform a hypothesis test on single parameters or the entire model with the null hypothesis that the population model is stationary. The sampling distribution under the null hypothesis can be constructed by repeatedly reshuffling the ordering of the data and estimate time-varying networks on them.

11.4 Estimating time-varying GGMs from time series of mood measurements

We now apply the ideas from above to estimate a time-varying network model on a time series of a single individual. We use a time series that contains a large variety of measurements of psychological, contextual, and symptom variables on 1,476 occasions on 238 consecutive days with up to ten measurements on each day. The measurements were taken from an individual diagnosed with major depression during a double-blind medication reduction period, which took place in two steps (day 42 and 98). To keep the present illustration simple, we focus on the six mood-related variables 'anxious,' 'lonely,' 'ashamed,' 'guilty,' 'worried,' and 'irritated,' which were measured on a 7-point Likert scale. For further details on this time series, see Kossakowski et al. (2017).

The goal of the analysis is to explore how the dependencies between variables change over the course of the time series. Similar to the example in the introduction, we are interested in the pairwise relations between variables at the same time point and how these relations may change over time. In addition, we assume that those relations are linear and that treating the ordered categorical variables as continuous will not distort the relationships between variables too much. Consequently, we can use the multivariate Gaussian distribution as our network model and therefore need a method that allows one to estimate a time-varying Gaussian distribution. In the interest of keeping the example simple, we model the time-varying lag-0 relationships between variables with a time-varying Gaussian distributions. However, for the reasons outlined in Chapter 9, it may often be advantageous to decompose the lag-0 relationships in temporal and contemporaneous relationships, for example with the graphical VAR (GVAR) model and the method presented here also works for the GVAR model.

As discussed in Section 11.3, the way in which parameters change over time can take many forms, and the estimation method should match our expectations about this form. In the present exploratory analysis, however, we have no expectations about the functional form and we choose a method that allows us to estimate parameters as continuous functions of time. Specifically, we choose the kernel-based method described in Section 11.3 which obtains a time-varying model by estimating a series of 'local' models across the time series. As discussed above, this method requires us to specify a bandwidth which we here set to 0.1. In addition, we provide the time stamps of all measurements, which allows us to take unequal intervals of measurements into account. We estimate the model using an unregularized nodewise regression scheme that combines estimates by averaging them, and we standardize all variables such that the final estimates can be interpreted as partial correlations (for details see Haslbeck & Waldorp, 2020). Finally we need to specify the number of local models to be estimated across the time series. The larger the number of local models, the more similar adjacent local models will be, and the only downside of choosing more is that estimation takes longer. Here, we choose 30 local models. Consequently, the estimated time-varying Gaussian distribution consists of $6(6-1)/2 \times 30 = 450$ partial correlations and 6×30 intercepts. The large numbers of parameters in time-varying models make it challenging to report all of them, which means that in practice one only reports the parameters that are of central interest. However, for the purpose of illustration, we consider several different visualizations. The Tutorial Box 11.1 shows how to estimate this model using the R-package *mgm*.

The time-varying Gaussian graphical model (GGM) was estimated with the function `tvmgm()` from the R-package *mgm*:

```
library("mgm")
output <- tvmgm(data = data,
                type = rep("g", 8),
                level = rep(1, 8),
                timepoints = timestamps,
                estpoints = seq(0, 1, length=30),
                bandwidth = 0.1,
                lambdaSeq = 0,
                threshold="none")
```

Next to the data, `tvmgm()` requires specification of the types of variables (here we create a vector with eight times `"g"` for our eight Gaussian variables) and the number of categories for each variable in case they are categorical. For continuous variables the number of categories is defined by convention as 1. This specification is necessary, because `tvmgm()` is also able to estimate time-varying models involving mixed variables.

Next, we specify the time stamps of each measurement in the data matrix via the `timepoints` argument. If this argument is not provided, the estimation algorithm assumes equidistant measurements. This is used by the weighting function to make sure that the time scale of the measurements corresponds to the time scale of the time-varying model. This can be a problem, for example, if a time series spans over a few weeks, but in one week measurements are missing for a few days. Next, we specify the estimation points with the `estpoints` argument as 30 points from the beginning (0) to the end (1) of the time series.

The argument `bandwidth` specifies the bandwidth parameter that should be used for estimation. Here, we specify the parameter to 0.1, however, in practice it should be selected using model selection. This can be done with the cross-validation scheme implemented in the function `bwSelect()`, which takes the same input as `tvmgm()` plus a sequence of candidate bandwidth values and specifications for the cross-validation scheme.

Finally, `lambdaSeq = 0` specifies that the estimation should be unregularized, and that no additional thresholding should be performed on the estimates (`threshold="none"`).

Tutorial Box 11.1. Estimating time-varying GGM using the kernel-weighting method assuming continuously time-varying parameters.

A natural choice of presenting the estimates in the network context would be to visualize the parameters of the time-varying Gaussian distribution by displaying its partial correlation matrix at different points in the time series. Figure 11.3 displays those networks at days 32, 66, 151, and 194 of the time series spanning 238 days.

Inspecting the partial correlations at those four time points shows us that the networks are changing across time and that different parameters seem to be changing more than others. For example, the partial correlation between 'anxious' and 'guilty' is relatively stable at around 0.40 and the partial correlation between 'lonely' and 'irritated' is relatively stable around 0. On the other hand, the partial correlation between 'anxious' and 'ashamed' is

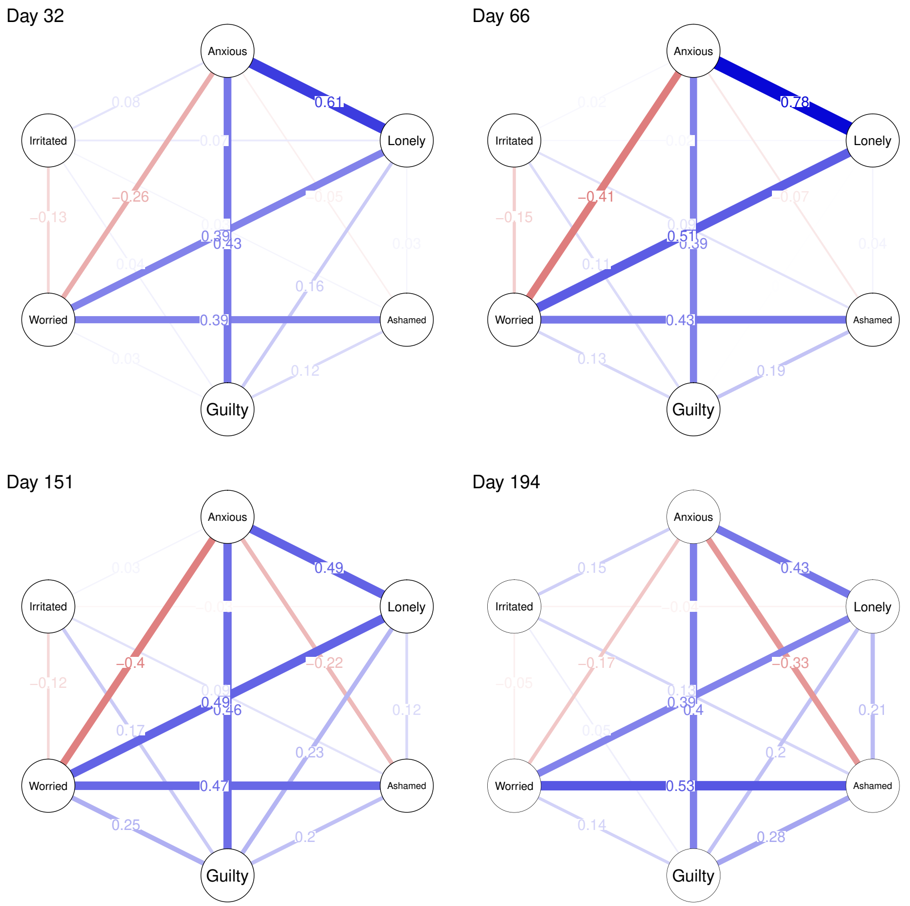

Figure 11.3. The partial correlation matrix of the time-varying Gaussian distribution at days 32, 66, 151 and 194 in the time series spanning 238 days. Blue edges indicate positive partial correlations, red edges indicate negative partial correlations, and the width of edges is a function of the absolute value of the partial correlation.

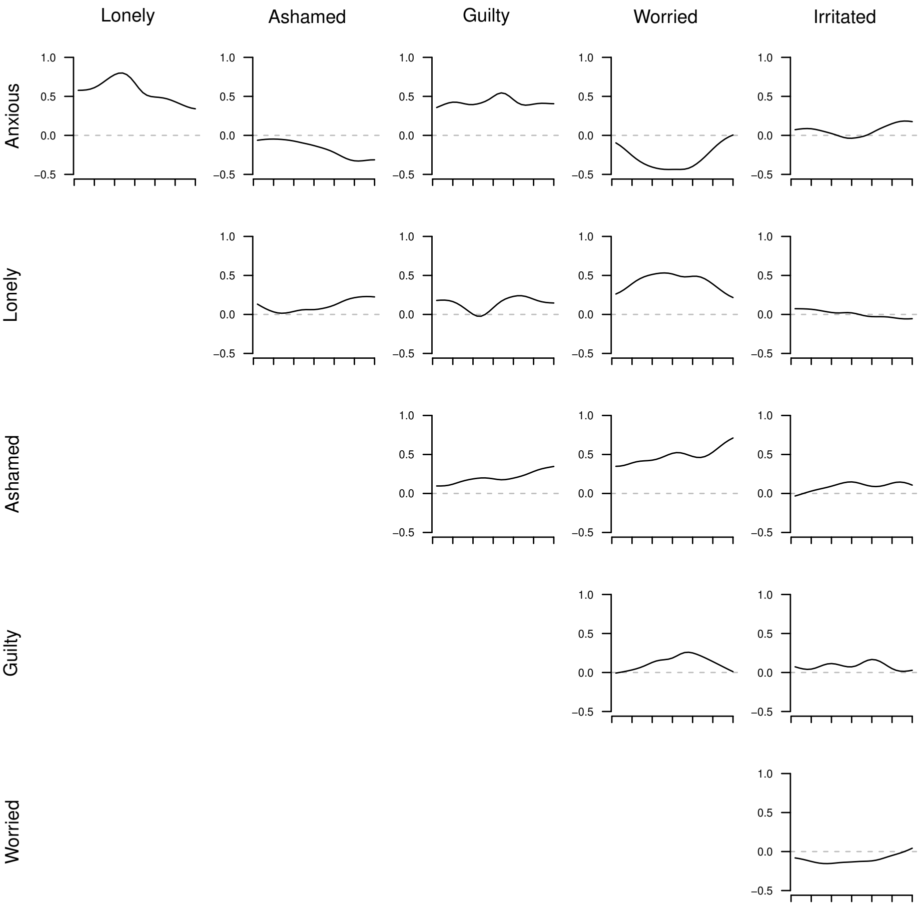

Figure 11.4. The 15 partial correlations in the time-varying Gaussian distributed as a function of time.

close to zero at day 32 but this negative effect increases throughout the rest of the time series. Another example is the partial correlation between 'anxious' and 'worried' which is −0.26 on day 32, but then becomes stronger on day 66 and 151 (−0.41 and −0.4) and then becomes weaker again on day 194 (−0.33).

An alternative visualization that allows one to better inspect the time-varying nature of single parameters is to plot the partial correlations as a function of time. Figure 11.4 displays all partial correlations as a function of time. This visualization provides a more complete view of the time-varying model since it displays all its parameters except the time-varying intercepts. Consistent with the visualization in Figure 11.3 we see that the partial correlation between 'anxious' and 'guilty' stays largely the same, however, we now also see that the partial correlation increase up to 0.50, which we might have missed when only inspecting the four network snapshots. Similarly, we get a better picture of the U-shaped function in which the partial correlation between 'anxious' and 'worried' changes across time.

Finally, Figure 11.5 displays the six time-varying intercepts of the model. We see that the intercepts of 'anxious,' 'lonely,' 'ashamed,' and 'guilty' stay largely the same across the time series. The intercepts of 'worried' and 'irritated,' however, seem to be increasing.

How could these types of results be interesting for empirical research? First of all, the time-varying analysis established that there is some variation in parameters across time, and that this variation differs across parameters. In a second step, we can go about explaining this variation. This could be done by relating changes in parameters to additional observational variables, for example capturing changes in the social or work environment of the individual. Another possible way to explain changes in parameters are interventions, as in the present case with the double-blind medication reduction scheme.

In this application we kept the analysis as simple as possible to focus on the estimation, visualization and interpretation of time-varying parameters. In this process we ignored several crucial questions, such as whether a partial correlation at a given point in time is actually different from zero in the population or only different from zero due to sampling variation. As in non-time-varying analyses, this problem can be addressed by performing significance tests or by using regularization (for details see Haslbeck & Waldorp, 2020). Another question we ignored is how to select the bandwidth parameter which controls how time-varying parameters can be. One way to select the bandwidth is to use a cross-validation scheme in which time-varying models are fit on training sets and their prediction errors are evaluated on hold out sets. Related to the selection of the bandwidth parameter is the question of whether the entire model or individual parameters

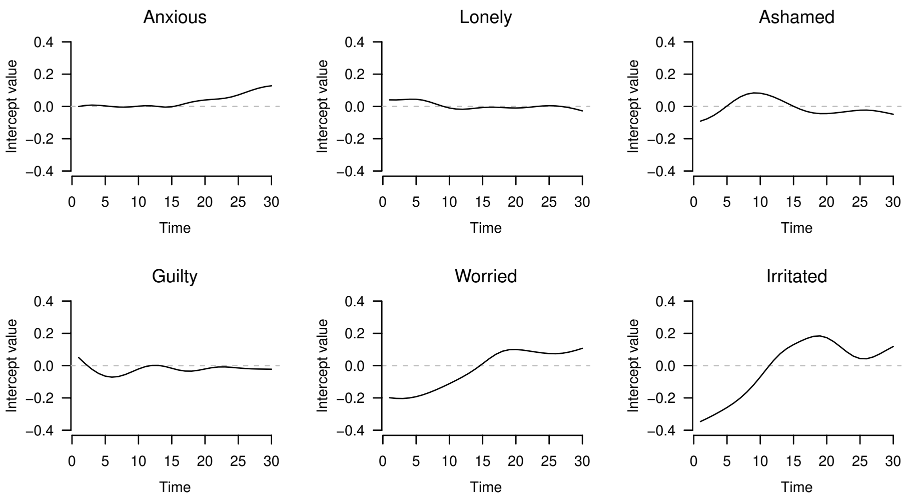

Figure 11.5. The six intercepts in the time-varying Gaussian as a function of time.

are actually time-varying in the population, or whether the observed time-varyingness is merely a function of sampling variation. For the method used here, these questions can be answered by a hypothesis test that simulates data under the null hypothesis that the model is stationary.

11.5 Conclusion

In this chapter we considered time-varying network models, which are models in which parameters can change across time. These models are always relevant when we are interested in how individuals change across time based on intensive longitudinal data. We discussed the difficulties involved in estimating time-varying models and presented a number of estimation methods that make different assumptions about the way in which the parameters change across time. Finally, we illustrated the use of time-varying network models by estimating a time-varying GGM to a $N = 1$ time series of mood measurements.

11.6 Exercises

Conceptual

11.1. I am fitting a stationary VAR model to experience sampling method (ESM) data collected over three weeks, which requires the assumption that the data generating mechanism is also stationary over those three weeks. Now I add measurements of another two weeks to obtain more stable estimates. How does this affect the assumption of stationary? Does this assumption become stronger, weaker, or is it unaffected by adding time points from subsequent weeks?

11.2. When estimating time series models under the assumption of stationarity, we can increase the length of the time series to estimate parameters more accurately. Why does this not work for time-varying models?

11.3. I estimate a time-varying VAR model with five variables and 50 observations. I obtain a high bandwidth parameter, indicating that the estimated time-varying model is largely stationary. Can I conclude that the true model is largely stationary?

11.4. Give an example for a system that changes its interactions continuously, and one example for a system whose interactions jump from one value to another.

11.5. Assume that the true time-varying model exhibits a discrete jump in the middle of the time series. Now, you estimate the time-varying model with one of the methods assuming continuous change. Describe the error you expect from this estimation procedure.

True or false

11.6. Increasing the window width in the moving-window approach increases the sensitivity to detect time-varying parameters.

11.7. The time-varying GAM method uses a weighting approach to estimate time-varying parameters.

11.8. A HMM with a single component is a stationary model.

11.9. In a time-varying model, all parameters are time-varying.

11.10. The bandwidth parameter in the kernel-weighting method can be selected using cross-validation.

References

Asparouhov, T., Hamaker, E. L., & Muthén, B. (2017). Dynamic latent class analysis. *Structural Equation Modeling: A Multidisciplinary Journal*, *24*(2), 257–269.

Borsboom, D. (2017). A network theory of mental disorders. *World Psychiatry*, *16*(1), 5–13.

Borsboom, D., van der Maas, H. L. J., Dalege, J., Kievit, R. A., & Haig, B. D. (2021). Theory construction methodology: A practical framework for building theories in psychology. *Perspectives on Psychological Science*, *16*(4), 756–766.

Bringmann, L. F., Ferrer, E., Hamaker, E. L., Borsboom, D., & Tuerlinckx, F. (2018). Modeling nonstationary emotion dynamics in dyads using a time-varying vector-autoregressive model. *Multivariate Behavioral Research*, *53*(3), 293–314.

Bringmann, L. F., Hamaker, E. L., Vigo, D. E., Aubert, A., Borsboom, D., & Tuerlinckx, F. (2017). Changing dynamics: Time-varying autoregressive models using generalized additive modeling. *Psychological Methods*, *22*(3), 409–425.

Bringmann, L. F., Vissers, N., Wichers, M., Geschwind, N., Kuppens, P., Peeters, F., Borsboom, D., & Tuerlinckx, F. (2013). A network approach to psychopathology: New insights into clinical longitudinal data. *PLoS One*, *8*(4), e60188.

Chow, S.-M., Ou, L., Ciptadi, A., Prince, E. B., You, D., Hunter, M. D., Rehg, J. M., Rozga, A., & Messinger, 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. J., van Borkulo, C. D., Giltay, E. J., van der Maas, H. L. J., Kendler, K. S., Scheffer, M., & Borsboom, D. (2016). Major depression as a complex dynamic system. *PLoS One*, *11*(12), e0167490.

Dablander, F., Pichler, A., Cika, A., & Bacilieri, A. (2021). Anticipating critical transitions in psychological systems using early warning signals: Theoretical and practical considerations. *Psychological Methods*. https://doi.org/10.1037/met0000450

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.

Dolan, C. V., & van der Maas, H. L. J. (1998). Fitting multivariage normal finite mixtures subject to structural equation modeling. *Psychometrika*, *63*(3), 227–253.

Gibberd, A. J., & Nelson, J. D. B. (2017). Regularized estimation of piecewise constant Gaussian graphical models: The group-fused graphical lasso. *Journal of Computational and Graphical Statistics*, *26*(3), 623–634.

Gibbert, A. (2017). Graphtime. https://github.com/GlooperLabs/GraphTime

Golub, G. H., Heath, M., & Wahba, G. (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, *21*(2), 215–223.

Hamaker, E. L., Grasman, R. P. P. P., & Kamphuis, J. H. (2010). Regime-switching models to study psychological processes. In P. C. M. Molenaar & K. M. Newell (Eds.), *Individual pathways of change: Statistical models for analyzing learning and development* (pp. 155–168). 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.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society*, *57*(2), 357–384.

Haslbeck, J. M. B., Bringmann, L. F., & Waldorp, L. J. (2021). A tutorial on estimating time-varying vector autoregressive models. *Multivariate Behavioral Research*, *56*(1), 120–149.

Haslbeck, J. M. B., Ryan, O., Robinaugh, D. J., Waldorp, L., & Borsboom, D. (2021). Modeling psychopathology: From data models to formal theories. *Psychological Methods*. https://doi.org/10.1037/met0000303

Haslbeck, J. M. B., & Waldorp, L. J. (2020). mgm: Estimating time-varying mixed graphical models in high-dimensional data. *Journal of Statistical Software*, *93*(8).

Hastie, T., Tibshirani, R., & Wainwright, M. J. (2015). *Statistical learning with sparsity: The lasso and generalizations*. Taylor & Francis.

Kim, C.-J., & Nelson, C. R. (1999). *State-space models with regime switching: Classical and Gibbs-sampling approaches with applications*. The MIT press.

Kolar, M., Song, L., Ahmed, A., Xing, E. P., et al. (2010). Estimating time-varying networks. *Annals of Applied Statistics*, *4*(1), 94–123.

Kossakowski, J., Groot, P., Haslbeck, J. M. B., Borsboom, D., & Wichers, M. (2017). Data from ‘critical slowing down as a personalized early warning signal for depression’. *Journal of Open Psychology Data*, *5*(1), 1.

Lu, Z.-H., Chow, S.-M., Ram, N., & Cole, P. M. (2019). Zero-inflated regime-switching stochastic differential equation models for highly unbalanced multivariate, multi-subject time-series data. *Psychometrika*, *84*(2), 611–645.

Monti, R. P., Hellyer, P., Sharp, D., Leech, R., Anagnostopoulos, C., & Montana, G. (2014). Estimating time-varying brain connectivity networks from functional MRI time series. *Neuroimage*, *103*, 427–443.

Monti, R. (2014). pySINGLE. github.com/piomonti/pySINGLE

Ou, L., Hunter, M. D., & Chow, S.-M. (2019). What’s for dynr: A package for linear and nonlinear dynamic modeling in R. *R Journal*, *11*(1), 91–111.

Pe, M. L., Kircanski, K., Thompson, R. J., Bringmann, L. F., Tuerlinckx, F., Mestdagh, M., Mata, J., Jaeggi, S. M., Buschkuehl, M., Jonides, J., et al. (2015). Emotion-network density in major depressive disorder. *Clinical Psychological Science*, *3*(2), 292–300.

Robinson, P. M. (1989). Nonparametric estimation of time-varying parameters. In H. P. (Ed.), *Statistical analysis and forecasting of economic structural change* (pp. 253–264). Springer.

Savi, A. O., Marsman, M., van der Maas, H. L. J., & Maris, G. K. J. (2019). The wiring of intelligence. *Perspectives on Psychological Science*, *16*(6), 1034–1061.

Schmittmann, V. D., Dolan, C. V., van der Maas, H. L. J., & Neale, M. C. (2005). Discrete latent Markov models for normally distributed response data. *Multivariate Behavioral Research*, *40*(4), 461–488.

Schwarz, G. et al. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*(2), 461–464.

Schweren, L., van Borkulo, C. D., Fried, E. I., & Goodyer, I. M. (2018). Assessment of symptom network density as a prognostic marker of treatment response in adolescent depression. *JAMA Psychiatry*, *75*(1), 98–100.

Tong, H., & Lim, K. (1980). Threshold autoregression, limit cycles and cyclical data. *Journal of the Royal Statistical Society. Series B (Methodological)*, *42*(3), 245–292.

van Borkulo, C. D., Boschloo, L., Borsboom, D., Penninx, B. W. J. H., Waldorp, L. J., & Schoevers, R. A. (2015). Association of symptom network structure with the course of depression. *JAMA Psychiatry*, *72*(12), 1219–1226.

van de Leemput, I. A., Wichers, M., Cramer, A. O. J., Borsboom, D., Tuerlinckx, F., Kuppens, P., van Nes, E. H., Viechtbauer, W., Giltay, E. J., Aggen, S. H., Derom, C., Jacobs, N., Kendler, K. S., van der Maas, H. L. J., Neale, M. C., Peeters, F., Thiery, E., Zachar, P., & Scheffer, M. (2014). Critical slowing down as early warning for the onset and termination of depression. *Proceedings of the National Academy of Sciences*, *111*(1), 87–92.

van der Maas, H. L. J., Dolan, C. V., Grasman, R. P., Wicherts, J. M., Huizenga, H. M., & Raijmakers, M. E. (2006). A dynamical model of general intelligence: The positive manifold of intelligence by mutualism. *Psychological Review*, *113*(4), 842–861.

van der Maas, H. L. J., & Jansen, B. R. J. (2003). What response times tell of children's behavior on the balance scale task. *Journal of Experimental Child Psychology*, *85*(2), 141–177.

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

Zhou, S., Lafferty, J., & Wasserman, L. (2010). Time varying undirected graphs. *Machine Learning*, *80*(2), 295–319.

Zucchini, W., MacDonald, I. L., & Langrock, R. (2017). *Hidden Markov models for time series: An introduction using R*. CRC press.