Running head:

Complexity, Chaos and Catastrophe: Modeling Psychopathology as a Dynamic System

Jolanda J. Kossakowski
Angélique O. J. Cramer

Word count: 10116

Keywords

Psychopathology; mental disorders; network perspectives; mean field approximation; complex dynamical systems; cusp catastrophe

J. J. Kossakowski • A. O. J. Cramer
Department of Psychology, University of Amsterdam, Nieuwe Achtergracht 129-B, 1001 NK Amsterdam, the Netherlands
E-mail: J.J.Kossakowski@uva.nl
Website: http://psychosystems.org; http://jolandakossakowski.eu; http://aojcramer.com

1. Introduction

In recent years, the notion that correlations between questionnaire items (e.g., insomnia and fatigue; hereafter called: variables) are the result of direct interactions between these variables (e.g., insomnia → fatigue) has grown in popularity. The rise of this network perspective started in 2006, when van der Maas et al. (2006) demonstrated that the *positive manifold*, the phenomenon that scores on cognitive tasks (e.g., verbal comprehension and working memory) are positively correlated, can be explained by means of a *network approach*. In this groundbreaking study, van der Maas et al. argued that, instead of explaining the positive manifold by means of a latent (i.e., unobserved) variable called *g* (i.e., general intelligence), the positive manifold can be explained by means of a *mutualism model*, in which variables have mutual, reinforcing, relations (e.g., verbal comprehension → working memory). Figure 1 shows an example of both the latent variable and the mutualism model (Data used for this example are freely available in the statistical software program R; R Core Team, 2016). Here, we have six variables (denoted by squares) that represent six test batteries of an intelligence test: Cattell's culture-fair test (general), vocabulary (vocab), reading comprehension (reading), maze (maze), block design (blocks) and a picture completion test (picture). In the left panel, the latent variable example, the latent variable *g* is the cause of the six variables, whereas in the mutualism model (right panel of Figure 1), each of the six variables has direct and mutual relations with every other variable in the model. Note that *k1 ... k6* represent environmental resources (such as the parents' level of education or a child's age) that may influence individual variables. Van der Maas and colleagues showed that, under certain circumstances, the latent variable model and the mutualism model were statistically equivalent, even though they are conceptually very different. That is, these results told us that you do not always need a latent variable model to explain the existence of a positive manifold between psychological variables. The primary aim of the present chapter is to elaborate on this network perspective as a theory whilst also providing the reader with a set of methods (including practical examples) that can be easily used for one's own data.

#FIGURE 1 ABOUT HERE#

The positive manifold is not unique to the topic of intelligence. For example, symptoms of psychological disorders (e.g., symptoms of major depression) are also consistently positively correlated with one another (i.e., positive manifold; e.g., Hartman et al., 2001); the same manifold was observed for variables that pertain to dimensions of normal

personality (e.g., neuroticism; Dolan, Oort, Stoel, & Wicherts, 2009). As such, van der Maas and colleagues (2006) paved the way for exploring the feasibility of a network perspective in other areas of psychological research: e.g., psychopathology (e.g., Borsboom, Cramer, Schmittmann, Epskamp, & Waldorp, 2011; Cramer, Waldorp, van der Maas, & Borsboom, 2010; McNally et al., 2014), personality (Cramer et al., 2012), health-related quality of life (Kossakowski, Epskamp, et al., 2016) attitudes (Dalege et al., 2016) and the International Classification of Diseases and Related Health Problems (ICD) and the Diagnostic and Statistical Manual of Mental Disorders (DSM; Tio, Epskamp, Noordhof, & Borsboom, 2016). In particular, the network perspective on psychopathology (e.g., psychological disorders) has not only resulted in a more mature conceptualization of disorders (and comorbidity between them) as networks of directly interacting symptoms (e.g., Cramer et al., 2010), it has also inspired the development of a host of new methods with which one is able to estimate networks for psychopathological data (see Fried et al., 2016 for an extensive review of the empirical psychopathological network literature).

As previously stated, many psychological constructs can be studied from a network perspective. In a network, individual variables are represented as nodes (circles); the mutual relation between two variables is visualized as a line between these two variables (hereafter called: edge; M. Newman, 2010). Figure 2 depicts two ways of constructing a network, using the same intelligence data as was used for Figure 1. The left network in Figure 2 shows a binary network where an edge can be either present ("1") or absent ("0"). As such, all edges have the same weight (i.e., "1") and thus the same thickness and color. The right network features green edges with different widths. Green edges represent positive relations between nodes, whereas red edges, not currently shown in Figure 2, represent negative relations between variables/nodes (Costantini et al., 2015). Green edges indicate that two variables reinforce each other: as one variable increases/decreases, the other increases/decreases as well. In contrast, red edges indicate that, as one variable increases/decreases, the other variable decreases/increases. The width and color saturation of the edge denotes the strength of the relation: the more saturated and thicker the edge, the stronger the relation between two nodes (Epskamp, Cramer, Waldorp, Schmittmann, & Borsboom, 2012). Edges do not necessarily need to connect two nodes: it is also possible for a node to reinforce itself (e.g., insomnia results in more insomnia). A so-called *self-loop* is then present in the network: an edge that starts and ends at the same node. In Figure 2 for example, the nodes that denote the vocabulary test and the block design test have a strong self-loop.

#FIGURE 2 ABOUT HERE#

In networks such as the one shown in the right panel of Figure 2, mutual relations between variables – that is, the *connectivity* of a network – can differ. This means that some variables have a stronger or weaker relation than others. For example, in Figure 2, the vocabulary test (vocab) and the block design test (blocks) have a strong relation in comparison to the other relations in the network, such as the relation between the variables reading test (reading) and the maze test (maze). This means that they have, relatively speaking, a strong influence on each other: when individuals score high on the vocabulary test, they also tend to score high on the block design test. The stronger the relations between the variables, the stronger the variables influence each other. *High connectivity* refers to a network with generally strong relations between variables, whereas *low connectivity* refers to a network with generally weak relations between variables. In other words, the higher the connectivity in a network, the stronger the relations between the variables in the network are. Considering intelligence, more strongly connected networks can be beneficial: in a strongly connected network, different aspects of an intelligence test reinforce each other more, which may lead to higher scores on an intelligence test. Thus: stronger connections may result in higher intelligence. Considering the development of psychopathology, this characteristic may not be so beneficial. For example, van Borkulo et al. (2015) showed that depressed participants, who still had the diagnosis of major depressive disorder (MDD) after two years, had a more strongly connected network than depressed participants who were in remission after two years. In this context, strong network connectivity can be seen as a measure of *vulnerability*: individuals who are more vulnerable for developing an episode of MDD tend to have a more strongly connected network than individuals who do not have this vulnerability for MDD (Cramer et al., 2016).

In this chapter, we aim at providing the reader with the basics of network modeling and applying these modeling principles to studying vulnerability in psychopathological constructs. We will start by providing an overview of various methods for estimating individual networks from empirical data (section 2), so that we can investigate vulnerability in psychological constructs at the level of the individual. Next, in section 3, we will elaborate on catastrophe theory, which is a central theory pertaining to studying vulnerability from a network perspective. Following this, we explain the Cramer model (Cramer et al., 2016) in section 4, accompanied by an application of this model to symptoms of mania (R. Kessler et al., 2014). We will proceed with the explanation of the empirical mean field approximation

(Kossakowski, Gordijn, Riese, & Waldorp, 2016; Waldorp & Kossakowski, 2016) in section 5, which is also accompanied by an application of the model to various positive and negative affects. We will conclude this chapter by taking a critical perspective: the network approach currently faces several challenges that have not been met (yet) and are important for advancing network modeling in psychology in general, and psychopathology in particular, in the next few years.

2. Constructing Networks

Various methods exist for estimating the connectivity of a network (i.e., the strength of the relations between variables in a network). Networks can be estimated for an entire group, or for individuals. When one estimates a network for a group of individuals, one can draw conclusions, for example, about specific relations between variables, that may be generalizable to the population of which we drew a sample. For this type of network, we only need one measurement per individual (cross-sectional data). It is also possible to estimate an intra-individual network, based upon which one can draw conclusions that only pertain to that specific individual. For this type of network, we need several measurements of one individual (time-series data). The applicability of each estimation method described in the remainder of this section heavily relies on the type of empirical data (categorical, continuous, or a mixture of the two), the number of participants that are being measured, and the number of measurements that have been collected per participant. The methods that we elaborate on in this section can be used to estimate intra-individual networks. In this section, we will discuss three methods: 1) *IsingFit* for binary data; 2) *Graphical Vector Auto Regression* (VAR) for continuous data, and 3) *Mixed Graphical Model* (MGM) for data with both binary and continuous variables.

A general disadvantage of network estimation is that networks can have many edges: in a network with six nodes, we can estimate up to 15 unique edges. However, in a network with 25 nodes, we can estimate as much as 300 unique edges. Many of these edges can be spurious. If we look at Figure 2 for example, an edge is present between the reading test (reading) and the block design test (blocks), while these two tests might in fact not be directly related at all. It could very well be that, when one scores high on the reading test, one also scores high on the vocabulary test (vocab), and as a result, also scores high on the block design test. Thus, it is possible that the relation between the reading test and the block design test is ‘caused’ by the vocabulary test.

All three methods control for such potentially spurious edges by using the *least absolute shrinkage and selection operator* (LASSO; Tibshirani, 1996). The LASSO imposes a parameter *lambda* (λ) that controls the level of sparsity (a sparse network is one with relatively few connections): a high λ will result in many weak edges that will be removed from the network, while a low λ will result in fewer weak edges that will be removed from the network. Choosing this parameter is therefore not trivial as it directly influences the structure of the resulting network. All methods that will be discussed estimate the network using different values for λ (Zhao & Yu, 2006). Then, the most optimal λ parameter is selected using the *Extended Bayesian Information Criterion* (EBIC; Chen & Chen, 2008; Foygel & Drton, 2010), where the network model with the lowest EBIC is chosen. The EBIC uses a parameter of its own *gamma* (γ), that controls whether or not the EBIC prefers simpler models, where simpler indicates a sparse network with fewer edges. A high γ parameter ensures that EBIC prefers the sparse network model with a high λ parameter and fewer edges. The hyperparameter γ ranges between *0* and *1.5* and needs to be manually set. For more information about this method for controlling sparsity, see Epskamp and Fried (2016) and Epskamp (2016).

Each method will be demonstrated with the same empirical dataset. These data are time-intensive measurements of one 57-year old male with a history of depression. Over the course of 239 days, this participant monitored his daily life experiences and affects several times a day, resulting in 1479 time points (for details about these data: Kossakowski, Groot, Haslbeck, Borsboom, & Wichers, 2016; Wichers, Groot, Psychosystems, ESM Group, & ESW Group, 2016). Table 1 shows the items, their meanings, their intended ranges and the response ranges.

#TABLE 1 ABOUT HERE#

2.1 Binary Data: IsingFit

The IsingFit method is based on the Ising model (Ising, 1925). This model originated in statistical physics and was developed to model ferromagnetic materials, such as the configuration of atoms and their corresponding spins. Nodes in an Ising model can only be in one out of two possible states (i.e., ‘active’ or ‘inactive’). Furthermore, only pairwise interactions are allowed in an Ising model, which means that only interactions between two nodes are allowed, and not, for example, three-way interactions (for example, the interaction

between ‘vocab’ and ‘blocks’ in Figure 2 cannot be positive when ‘reading’ is active, and negative when ‘reading’ is inactive). The Ising model consists of two node-specific parameters: an *interaction parameter*, which denotes the strength of the interaction between two nodes and thus the edge weight, and a *threshold parameter* that represents the preference of a node to be either active or inactive regardless of the activity of its neighboring nodes (neighbor defined as all nodes that a particular node has a connection with; van Borkulo et al., 2014). A positive threshold corresponds to a preference of being active, whereas a negative threshold corresponds to a preference of being inactive. A threshold parameter of exactly 0 corresponds to having no preference.

Essentially, the IsingFit method regresses one node on all other nodes in an iterative manner, using the optimal penalty parameter. This means that twice as many regressions will be performed as there are nodes in the network. For example, when we consider Figure 2, the vocab node will be regressed on reading, maze, blocks, picture and general, reading will be regressed on vocab, maze, blocks, picture and general, and so on. After this node-wise regression, we have an interaction parameter from vocab to reading, and from reading to vocab. In order to calculate the final edge weight, the so-called AND-rule is used: only when both interaction parameters are nonzero will an edge be drawn between the two nodes, with an edge weight that is the mean of these two interaction parameters. It is also possible to use the OR-rule, which states that at least one of the two parameters must be nonzero in order for an edge to be drawn between the two nodes. This edge will then have the weight that is associated with the nonzero interaction parameter. For clarity of presentation, all networks in this chapter will be estimated using the AND-rule.

Figure 3 demonstrates the IsingFit method. To improve visual comparison, the layout of the networks is fixed. In the left panel, a network with a γ parameter of 0 is estimated, while in the right panel, a network with a γ parameter of 1 is estimated. It can be seen that item 12 (I feel strong) has many negative interactions with the other variables. This means that, as item 12 increases, items such as item 1(I feel relaxed), 16 (I can concentrate well) and 26 (I have a headache) decrease, and vice versa. The two networks depicted in Figure 3 demonstrate why setting the γ parameter is not trivial: the network in the left panel (estimated with $\gamma = 0$) clearly has more edges than the network in the right panel (estimated with $\gamma = 1$). More specifically, in the network in the left panel, approximately 39% of the edges are nonzero, whereas in the network in the right panel, approximately 19% of the edges are

nonzero. This difference in networks estimated with different values for γ will decrease as the sample size increases.

#FIGURE 3 ABOUT HERE#

The IsingFit method is initially intended for cross-sectional data. However, the procedure is also suitable for time-series data if we assume that all measurements are independent of one another. This means that the scores on measurement *t* do not rely on the scores that were obtained at measurement *t-1*. To illustrate, suppose an individual completes the same questionnaire twice. The assumption then states that the answers given the second time are not (partially) determined by the answers given the first time. It is debatable whether this assumption holds in time-series data, but it is beyond the scope of this chapter to go into further detail.

2.2 Continuous Data: Graphical Vector Auto Regression

When dealing with continuous time-series data, we can estimate two types of networks: a contemporaneous network that shows the relations between variables that occur within the same time point, and a temporal network that shows the relations between variables that occur across time points (Epskamp, Waldorp, Mottus, & Borsboom, 2016). For example, Figure 4 shows hypothetical examples of these two types of networks for affect variables: feeling distressed, active, interested, and alert. The left panel presents a contemporaneous network, which shows that feeling interested and feeling active have a negative relation. This means that, as one feels more interested, one will feel less active *within the same time point*, and vice versa. In a contemporaneous network, relationships are mutualistic (i.e., no arrows) because these relationships unfold within the same time point. The right panel of Figure 4 presents a temporal network. This network shows that feeling alert and feeling distressed have a positive influence on each other: when one is feeling more alert, one will feel more distressed, and when one feels more distressed, one will feel more alert *at the next time point*.

#FIGURE 4 ABOUT HERE#

Both networks are equally important and interesting: when collecting time-series data, one is often interested in the progression of an individual throughout time. At the same time, important associations between variables may occur within a single time point that are not captured when solely considering temporal relations between variables. Figure 4 shows this: fictitious relations are present between feeling distressed, active, interested, and alert. Here, we see that a positive, mutual relation between feeling alert and feeling active exists in the contemporaneous network (left panel), which means that, as one feels more active, one also feels more alert and vice versa. At the same time, in the temporal network (right panel), feeling active negatively influences feeling alert, which means that, as one feels more active, one will feel less alert at the next time point. Also note that there is a negative mutual relation between feeling distressed and feeling active in the contemporaneous network, and that the temporal network shows that this relation also exists over time: feeling distressed negatively influences feeling active over time, and vice versa. As Aristotle one said: "the sum is more than its parts"; estimating both a temporal and a contemporaneous network will give us a more complete picture of the progression of a construct, such as a psychological disorder or intelligence, and will yield more information than focusing on either one of the two networks (Epskamp, Waldorp, et al., 2016).

We can estimate both these network structures using Graphical Vector Auto Regression (VAR; Epskamp, Waldorp, et al., 2016). This method uses complex matrix algebra to estimate the elements of a matrix containing the temporal relations between the items and the elements of a matrix containing the contemporaneous associations between the items (Wild et al., 2010). It is beyond the scope of this chapter to go into depth with respect to the estimation of the temporal and contemporaneous relations. For a more detailed and technical explanation, see Epskamp et al. (2016) and Wild et al. (2010). Figure 5 visualizes the essence of the Graphical VAR method: For each node, we regress it on itself at the following time point *t+1*, and we regress it on all other items at the following time point *t+1*. In the example shown in Figure 5, node S is regressed on itself and node C, both at the previous time point *t-1*. The procedure is repeated for node C, which is regressed on itself and node S at the previous time point *t-1*.

#FIGURE 5 ABOUT HERE#

Figure 6 shows network structures for data of the 57-year old male with depression estimated using Graphical VAR. For explanatory purposes, we assumed all variables to be

continuous. The left panel shows the contemporaneous network, and the right panel the temporal network, estimated with $\gamma = 0$. We left the network structures for $\gamma = 1$ out for clarity of presentation. In the contemporaneous network (left panel), variable 5 (I feel lonely) has strong positive relations with items 2 (I feel down) and 6 (I feel anxious). Furthermore, a strong negative relation is present between items 16 (I can concentrate well) and item 28 (I am sleepy). This means that, as the participant can concentrate more, he is less sleepy at the same time point. In the temporal network (right panel), we see strong positive self-loops for items 11 (I doubt myself), 22 (I am tired) and 27 (I have a headache): when the participant doubts himself at one time point the feedback reinforces this feeling, and he will doubt himself more at the next time point. Interestingly, the contemporaneous network has many more edges than the temporal network: in the contemporaneous network, approximately 37% of the edges are nonzero, while in the temporal network, approximately 4.5% of the edges are nonzero. This shows the importance of looking both at contemporaneous and temporal network structures to uncover the dynamics of an individual throughout time.

#FIGURE 6 ABOUT HERE#

2.3 Mixed Data: Mixed Graphical Model

In the previous two sections, we discussed methods that estimated one network structure. When we estimate one network structure, we assume that this structure does not change over time, and is thus *stationary*. When dealing with time-series data, we are interested in how an individual potentially changes over time, and therefore we are also interested in how the associated network might change over time. Where the IsingFit and the Graphical VAR method both estimate one (contemporaneous) network, the Mixed Graphical Model (MGM) method estimates multiple (more than 1) network structures that together demonstrate how the relations between variables in a network of an individual may change over time (Haslbeck & Waldorp, 2016a, 2016b). For example, with the MGM method, we can estimate two, five, ten, or more network structures (how many is up to the user to decide) that, when shown sequentially, display potential changes in the network structure throughout time.

In order to estimate time-varying networks using the MGM method, we start out by identifying the type of variables in the time-series: variables can be continuous (e.g., age) or categorical (e.g., sex). The time-series data is then cut into *z* pieces (the number of pieces is

chosen by the user), and for each piece, the network structure is estimated, taking the types of variables (i.e., continuous or categorical) into account (see Haslbeck & Waldorp, 2016b for a more detailed description). Similar to the IsingFit method, each node is regressed on all other nodes. For example, when we, again, consider Figure 2, the vocab node will be regressed on reading, maze, blocks, picture and general, reading will be regressed on vocab, maze, blocks, picture and general, and so on. After this node-wise regression, we obtain two edge weights for the edge between vocab and reading, for example. When both of these edge weights are nonzero, the mean of these edge weights will be used to visualize the relation between, in this example vocab and reading. When either of the two edge weights is zero, no edge will be drawn between the two variables.

Because one is estimating multiple network structures that together show the potential changes in a network structure throughout time, one needs to decide how many network structures are desired, which is an arbitrary choice. Also, the MGM method assumes that the network structure does not change within each piece of time for which the network structure is estimated. The MGM method assigns each time point within one piece a certain weight, that indicates how important that time point is for the estimation of the edges. Time points that are seen as more important will have a stronger influence on the edge estimation process than time points that are seen as less important. For this chapter, we use the default assignment of weights to time points, as suggested by Haslbeck and Waldorp (2016a). For a more detailed and technical explanation of the MGM method, see Haslbeck and Waldorp (2016a) and Haslbeck and Waldorp (2016b).

#FIGURE 7 ABOUT HERE#

Figure 7 shows network structures for the single patient data estimated using MGM. In order to estimate an MGM, we dichotomized items 2 (I feel down), 5 (I feel lonely), 6 (I feel anxious) and 10 (I feel guilty), because their scale ranges from -3 to +3 and the MGM method cannot handle negative scores. Furthermore, we assumed that items with seven response categories were continuous, resulting in three continuous items and 25 categorical items. We decided (arbitrarily) to divide the time-series up into five pieces and estimate five networks accordingly. We estimated the five network structures both with $\gamma = 0$. We left the network structures estimated with $\gamma = 1$ out for clarity of presentation. Starting at the top network in Figure 7 and then going down, the changes in the network structure become

visible. For example, the edge between items 5 (I feel lonely) and 10 (I feel guilty) is not present in the first network, but slowly appears in the fourth and fifth network. This means that the relation between feelings of loneliness and feelings of guilt developed over time. Furthermore, it can be seen that the edge between items 2 (I feel down) and 6 (I feel anxious) appears in the second network, but disappears again in the fifth network. This indicates that, at the beginning and at the end of the data collection period, feeling down and feeling anxious were not related, but this relation developed into a strong positive relation, and then disappeared again. This shows the importance of estimating not one network structure, but several network structures, as Figure 7 shows that we gain more insight into the dynamics of an individual throughout time compared to only one network.

3. Catastrophe Theory

One of the problems with networks (also called *systems*) is that they easily become complex. Even a relatively simple model with, say, 10 variables already entails more than 50 parameters to be estimated. For a slightly more complex model with 14 variables the number of parameters to be estimated is larger than 100. As such, if one is interested in understanding and/or predicting the behavior of a system (network) as a whole then this becomes more difficult as the number of system variables increases and with it the number of parameters. Therefore, in other fields, one reduces the complexity of a network model in order to understand its general dynamics by means of using a *cusp catastrophe model*. With this mathematical model, one reduces complexity by considering two orthogonal control variables - instead of the more than 100 parameters when a system is comprised of 15 variables - and their impact on the behavior of a system.

#FIGURE 8 ABOUT HERE#

Figure 8 shows a hypothetical cusp catastrophe model with two control variables, stress (x-axis) and connectivity (y-axis; Cramer et al., 2016) that have their own way of influencing the behavior of the system (i.e., a psychological disorder in our example; z-axis). Panel A of this figure shows the 3D representation while Panel B shows the 2D representation with stress on the x-axis and the behavior of the system on the y-axis for two situations: weak connectivity (green line, top figure of Panel B) and strong connectivity (red line, bottom figure of Panel B). Stress acts as the so-called *normal* variable while connectivity acts as the so-called *splitting* variable.

For sufficiently low values of the splitting (i.e., connectivity) variable, the behavior of the system responds to increasing values of the normal (i.e., stress) variable in a smooth and continuous fashion (see green line in top figure of *Panel B* of Figure 8): that is, when connectivity of a, say, depression system is relatively weak, then more stress results in the system being more depressed in a continuous fashion. Cramer et al. called such systems *invulnerabl*e or resilient networks. For higher values of the splitting variable (i.e., strong connectivity) the outcome surface splits and bifurcates from smooth, continuous changes to sudden, discontinuous jumps for increasing values of the normal variable (i.e., stress; see red line in bottom figure of Panel B of Figure 8): that is, when connectivity of a psychopathological system is relatively strong, then at a certain point (i.e., the tipping points in bottom figure of Panel B of Figure 8) the system either jumps from a non-disordered to a disordered state, or the other way around, which is called a *critical transition*. Cramer et al. called such systems *vulnerable* networks. In between these tipping points lies a *forbidden* zone (in bottom figure of Panel B of Figure 8: the dashed part of the red line): within this zone, the behavior of the system (i.e., the disorder) is unstable to such an extent that even a tiny amount of stress will already kick the system into a stable (non-)disordered state (i.e., the solid parts of the red line). A final phenomenon that is present in networks with strong connectivity is *hysteresis*. Figure 9 shows a hypothetical example of hysteresis for a symptom network that is strongly connected. The *x*-axis represents stress while the *y*-axis represents the state of the system (from a healthy - none or few symptoms active - to a disordered state - most or all symptoms active). The green line (UP) shows what happens if stress is increasing: as we already saw in Figure 8 increasing stress results in a sudden jump from a healthy state to a disordered state. But what happens when, once disordered, we start lowering the amount of stress that is influencing the system? This is shown with the red line (DOWN): the amount of stress needed to get the system back into a healthy, non-disordered, state (jump from "health" to "disorder" in the green line) exceeds the amount of stress that tipped the system into a disordered state in the first place (jump from "disorder" to "health" in the red line). That is, hysteresis is the gap, marked with an arrow in Figure 9, between the green and the red line.

#FIGURE 9 ABOUT HERE#

4. The Cramer Model

In an earlier section we elaborated on the theoretical principles of a cusp catastrophe model. We sketched one possibility of such a model in which a disorder is the result of an interaction between two control parameters, the connectivity of the system (i.e., the strength of relations in a symptom network of a disorder) and the amount of stress (i.e., external perturbation) that is put on that system (see Figure 8). This may sound appealing but we first need to establish that, indeed, vulnerability of a system to becoming disordered has something to do with the connectivity of that system. In this section, we first provide some information on mania data that we used as an example for showing the workings of the Cramer model. Next, we outline the specifics of this model and show that, indeed, vulnerable networks/systems are those with strong connections between its symptoms. We conclude this section by showing how one might detect critical transitions with this model.

4.1 Mania Data

Data for our example came from the National Comorbidity Survey Replication (NCS-R). This is a nationally representative household survey of English speakers 18 years and older in the United States (see R. C. Kessler et al., 2004). The NCS-R survey schedule is the version of the World Health Organization (WHO) Composite International Diagnostic Interview that was developed for the WHO World Mental Health Survey Initiative (WMH-CIDI; R. C. Kessler & Ustun, 2004). The interviews were conducted between February 2001 and April 2003. A total of 9282 respondents participated in Part 1 of the interview (core diagnostic assessment), the data of which we used for this chapter. Specifically, we used data for the 10 symptoms of mania: elevated mood, restlessness, pressure of speech, flight of ideas, loss of social inhibitions, decreased need for sleep, inflated self-esteem, distractibility, reckless behavior, and marked sexual energy. For each of these criteria, respondents had to indicate whether or not they had suffered from that symptom in the past year. As such, the data were dichotomous. Missing data were imputed with zeroes.

We estimated network parameters for the 10 symptoms of mania in the NCS-R dataset with the IsingFit method as explicated earlier. Figure 10 shows the resulting network. Note that the skip structure of the NCS-R schedule is responsible for the fact that elevated mood (*mod* in Figure 10) is such a central symptom in this network: a ‘yes’ response is needed to proceed to the remainder of the questionnaire while a ‘no’ results in skipping the other nine mania criteria. Other relatively strong connections seem intuitively plausible: for example, a strong connection between loss of social inhibitions and marked sexual energy;

and between distractibility and flight of ideas. The empirical parameters of this network are used in the next section as input for the formal dynamic Cramer model.

#FIGURE 10 ABOUT HERE#

4.2 The Cramer Model

The model assumes the following. First, symptoms X_i can be 'on' (1; active) or 'off' (0; inactive). Second, symptom activation takes place over time t such that, for example, elevated mood at time t may cause marked sexual energy at time $t + 1$. Third, a symptom i receives input from symptoms with which it is connected in the empirical mania symptom network based on the NCS-R data. We call these symptoms with which symptom i is connected *neighbor* symptoms. The empirical weight parameters are collected in a matrix **W** for the $J = 10$ mania symptoms: entry W_{ij} thus represents the logistic regression weight between symptoms i and j as estimated from the NCS-R data. Next, the total amount of activation a symptom i receives at time t is the weighted (by **W**) summation of all the neighboring symptoms **X** - that is, the vector that contains the "0" (inactive) and "1" (active) values - at time $t - 1$. Cramer et al. (2016) call this the *total activation function* (boldfaced parameters are estimated from the NCS-R mania data):

$$A_i^t = \sum_{j=1}^{J} \boldsymbol{W_{ij}} X_j^{t-1} \tag{1}$$

An example: suppose we wish to compute the total activation function for marked sexual energy, which is connected to elevated mood with a NCS-R weight parameter of 2 and to reckless behavior with a NCS-R weight parameter of 0.80. Also assume that at time $t - 1$ elevated mood is activated while reckless behavior is not (encoded in the vector **X**). In this simple example, the total activation at time t for the symptom marked sexual energy is $2 \cdot 1 + 0.80 \cdot 0 = 2$.

In the next step, we develop a function for computing the probability of symptom i becoming active at time t. This logistic probability function states the following: the probability of symptom i becoming active at time t depends on the difference between the total activation of its neighboring symptoms and the threshold (estimated with the NCS-R

mania data) of symptom *i* (in formula 2: $\boldsymbol{b_i} - A_i^t$). The more the total activation exceeds the threshold of symptom *i* at time *t*, the higher the probability that symptom *i* becomes active (in formula 2 $P(X_i^t = 1)$) at time *t*. Cramer et al. (2016) call this the *probability function* (boldfaced parameters are estimated from the NCS-R mania data):

$$P(X_i^t = 1) = \frac{1}{1 + e^{\boldsymbol{b_i} - A_i^t}} \quad (2)$$

Please note that the parameter $\boldsymbol{b_i}$ denotes the absolute value of the threshold of symptom *i* as estimated from the NCS-R data. So let us return to our simple example above in which marked sexual energy was only connected to elevated mood and reckless behavior. Suppose that the threshold for marked sexual energy is 1 and we already know that the total activation at time *t* equals 2. Then the probability of having marked sexual energy at time *t* equals $\frac{1}{1+e^{1-2}} = 0.73$. This probability becomes much lower when the threshold exceeds the total activation, for example when the threshold would be 4. In that case, the probability becomes $\frac{1}{1+e^{4-2}} = 0.11$. A special case arises when A_i^t is equal to $\boldsymbol{b_i}$, that is, when the amount of activation of the neighbors of symptom *i* is exactly equal to the threshold of symptom *i*. In that case, the probability of symptom *i* becoming active is exactly 1/2 (e.g., $\frac{1}{1+e^{2-2}} = 0.50$).

To summarize, the Cramer model is a process model that develops over time. The probability of a symptom becoming active at a particular point in time depends on both its threshold and the amount of activation it receives from its neighboring symptoms at that same point in time. The more activation a symptom *i* receives from its neighboring symptoms and the lower its threshold, the higher the probability of symptom *i* becoming active. In the next few sections, we will use this model to study the behavior of a strongly connected mania system when put under stress. Will this behavior conform to what we would expect from the cusp catastrophe model we elaborated on earlier (i.e., the red line in Panel A and bottom graph of Panel B of Figure 8)?

4.3 Putting the Cramer model under stress: critical transitions?

For the purposes of this section we are particularly interested in the potential behavior of a strongly connected system that is put under stress. Would we find, as a cusp catastrophe model predicts, that putting stress on a strongly connected system results in sudden discontinuous jumps from one state to the other? More specifically, would there be a *critical*

transition, the change that occurs when the condition of a person exceeds the tipping point (dots in bottom graph of Figure 8), which transfers the person from one state (e.g., no mania) to another (e.g., mania) in response to a relatively small external force: i.e., in the bottom graph of Panel B of Figure 8 only a relatively small amount of stress is needed to get the system past the tipping point into an alternative stable state (i.e., one of the solid parts of the red line in bottom graph of Panel B in Figure 8; Olde Rikkert et al., 2016). Cramer et al. (2016) showed that for a strongly connected depression network, it was indeed the case that the behavior of the depression system was discontinuous, with sudden jumps from a non-depressed state to a depressed state with a forbidden zone in between both stable states. Here, we sought to replicate these findings for the NCS-R mania data.

First, in order for the mania system to be sufficiently strongly connected, we multiplied the weights matrix **W** that was estimated with IsingFit on the NCS-R mania data with $c = 1.5$. Second, we extended the total activation formula in 1 with a stress parameter S_i^t a number that was added to the total activation of the neighbors of symptom *i* at time *t*: the higher S_i^t - that is, the more stress - the higher the total activation function, and thus the higher the probability (formula 2) that symptom *i* will become active at time *t*. This resulted in the following modified total activation function:

$$A_i^t = \sum_{j=1}^{J} c\boldsymbol{W_{ij}} X_j^{t-1} + S_i^t \tag{3}$$

The probability function remained the same as stated in formula 2. We simulated 10000 time points starting with all symptoms being "off" (i.e., vector **X** with only zeroes). At each time point, we computed total activation and the resulting probability of a symptom becoming active. Next, symptom values (either "0" or "1", denoting inactive and active, respectively) were sampled using these probabilities. Meanwhile, over the course of the 10000 points S_i^t was repeatedly gradually increased from -15 to 15 and then decreased from 15 to -15 with small steps of 0.01. The impact of the stress parameter on the behavior of the mania system was quantified by computing the average state *M* of the system, the average number of symptoms active at a certain time point *t*. Specifically, since all the stress parameter values were used multiple times during the simulation - because of the repeated increasing and decreasing of the stress parameter during the course of the simulation - we averaged states within 0.20 ranges of these stress parameter values.

#FIGURE 11 ABOUT HERE#

Figure 11 shows the main results of the simulation: the *x*-axis represents stress while the *y*-axis represents the average state of the system *M* (note: the higher on the state variable, the more mania symptoms are active, and thus the more 'manic' the system). The grey line (and points) represents the average number of active symptoms when stress was increasing; and the black line (and points) represents the average number of active symptoms when stress was decreasing. The results are consistent with what we predicted from the cusp catastrophe model. First, the behavior of the system was discontinuous with sudden jumps from a non-manic to a more manic state and vice versa. That is, a small increase (decrease) in stress can lead to a disproportional reaction, resulting in a more manic (non-manic) state with more (less) symptoms active. Second, Figure 11 clearly shows that during the transition from healthier to more manic states, and vice versa, a 'forbidden zone' (from around 1 to 7 symptoms) was crossed that does not seem to function as a stable state (i.e., no data points in that area, see black box in Figure 11). Third, and final, the results show clear hysteresis: the amount of stress reduction needed to get the system into a non-disordered state (i.e., only a few symptoms active or none at all) exceeds the amount of stress that initially tipped the system into a more manic state.

In sum, when put under stress, strongly connected systems that behave over time according to the equations of the Cramer model (Cramer et al., 2016) show critical transitions from non-disordered states to disordered states and vice versa. These critical transitions take place at tipping points and between these tipping points states are unstable therefore 'forcing' the system to either the upper branch of disorder (solid upper red line in bottom graph of Panel B in Figure 8) or the lower branch of relative health (solid bottom red line in bottom graph of Panel B of Figure 8). Perhaps more important than knowing that there are critical transitions, is there any way with which we can detect such upcoming critical transitions? Detecting upcoming critical transitions in real data of individuals might prove beneficial, for example because it would allow for early intervention, before a critical transition to a disordered state takes place, that might help in preventing a diagnosis.

4.4 Detecting critical transitions: critical slowing down

There is evidence for the hypothesis that all catastrophic systems, from financial systems to the climate, display *early warning signals* that that system is approaching a

tipping point (Carpenter & Brock, 2006; Dakos et al., 2008; Fort, Mazzeo, Scheffer, & van Nes, 2010; Scheffer et al., 2014). One such early warning signal is called *critical slowing down*: right before a tipping point, the system becomes increasingly slower in recovering from small perturbations. Pertaining to mania, for instance, one might see that someone has more difficulty than usual to recuperate from a relatively minor daily hassle such as an argument with a spouse over which restaurant to go to.

Numerically, this slowing down can be traced by inspecting autocorrelations: that is, the correlations between values of the same variable at multiple time points (e.g., the correlation between 60 measurements of elevated mood over time). These autocorrelations increase when the system slows down: slowing down means that at each time point, the system much resembles the system as it was at the previous time point, meaning that the autocorrelation is relatively high. The question now is: if we compute the autocorrelations between the state variable *M* of the strongly connected mania system, would we find an increase in these autocorrelations?

Again, we used the NCS-R mania data and specifically inspected the autocorrelations between the state variable *M* (i.e., the total number of active mania symptoms) over time for the network with strong connectivity. The setup of the simulation was identical to the one reported in the previous section about putting a network/system under stress. The results are presented in Figure 12. As expected, when stress was increasing, the autocorrelations between the states of the strongly connected mania network increased (dashed line increasing, starting at roughly the -1 stress point) markedly right before the network abruptly switched from a relatively non-manic to a manic state (thicker dashed line jumping from 4 to 10 symptoms right after the -1 stress point). However, Figure 12 also showed an increase in active symptoms from 0 to 4 symptoms before the sharp increase in autocorrelations so it appears that for this data example, the early warning was not early enough. Since Cramer and colleagues (2016) did show a clear early warning signal for major depression data it remains an open question to what extent the present results are a clue to the unsuitability of autocorrelations as window into predicting critical transitions.

#FIGURE 12 ABOUT HERE#

5. The Empirical Mean Field Approximation

5.1 Description

The empirical MFA is based on the work of Waldorp and Kossakowski (2016). As previously stated, networks can quickly become quite complex. By assuming that variables behave in a similar manner and have the same number of connections to other variables, we can reduce such a complex and multivariate system to a single equation (Kossakowski, Gordijn, et al., 2016; Waldorp & Kossakowski, 2016). We can then use this equation to investigate the behavior of the network as a whole over time.

#FIGURE 13 ABOUT HERE#

Suppose we have a network with 9 nodes, such as the one depicted in the left panel of Figure 13. In this network, nodes are either active (1, black) or inactive (0, white). For example, in the left panel of Figure 13, three nodes are active (i.e., black). In the MFA, we check for each node how many of its neighbors (nodes that are directly connected to a specific node) are active. In Figure 13 (left panel), when we focus on the middle node, we see that it has four neighbors, and three of them are active. We use this information on the activity of a node's neighbors and whether or not this is the majority of a node's neighbors to set a probability parameter p that decides whether or not a specific node becomes active itself at the next time point $t + 1$. As a rule, we state that, when at least the majority of a node's neighbors is active, the probability for that specific node to become active will be $1 - p$. When less than the majority of a node's neighbors is active, this probability will be p. Looking again at Figure 13, 3 out of 4 neighbors (i.e., the majority) of the middle node are active at time point t, so the probability for the middle node to become active at time point $t + 1$ is $1 - p$. In the right panel of Figure 13, we see that the middle node indeed became active as a result of its active neighbors and its probability parameter. We can repeat this process for each node within a time point, and for each time point, thereby creating a $t \times n$ matrix that contains 0s and 1s for each time point t and each node n.

For each time point, we can calculate the number of active nodes proportionally to the total number of nodes, called the *density*. These densities can be visualized, as is shown in Figure 14 with time on the x-axis and the densities on the y-axis. In this figure, we see how the density, the proportion of active nodes in a network, changes over time. In this example, we see that the density suddenly jumps from one stable state to the other, which is a critical transition. In this case, even though we spotted the transition, we are too late in investigating

whether or not the individual was vulnerable for the critical transition before it occurred. Fortunately, the empirical MFA enables us to investigate this vulnerability.

#FIGURE 14 ABOUT HERE#

The network in Figure 13 is called a *grid* where each node has the same number of neighbors. In psychology, it is hard to come up with an example in which each variable has the same number of neighbors. The MFA assumes that each node has the same number of neighbors, and Waldorp and Kossakowski (2016) investigated the performance of the MFA when this assumption is no longer satisfied, and nodes can have a different number of neighbors. They also investigated two other types of network structures: a *random graph* and a *small world graph*, both of which are displayed in Figure 15. A random graph is a network in which each edge has a constant probability to be drawn, p_e which is independent of the size of the network (Bollobás, 2001; Durett, 2007). In Figure 15 (left panel), we see a random graph with an edge probability of *0.3*, which means that each edge has a probability of *0.3* to be present in the network. This also means that around *30%* of the edges are present in the network. In the right panel of 15, we see a small world graph, which is a network structure that starts out as a grid (i.e., each node has the same number of neighbors, thus edges), but in which each edge has an constant probability p_w to be rewired to two other, possibly non-neighboring nodes and added to the grid (Newman & Watts, 1999), and which is also independent of the size of the network. In Figure 15 (right panel) we see a small world graph with a rewiring probability of *0.2*, which means that each edge in the original grid has a probability of *0.2* to be rewired and added to the network.

#FIGURE 15 ABOUT HERE#

In Waldorp and Kossakowski (2016), the performance of the MFA was shown by means of a simulation study: they theoretically determined the shape of the bifurcation diagram, and determined how well simulated data matched the theoretical bifurcation diagram under various conditions. In Kossakowski et al. (2016), they showed how the MFA can be used in practice. Using maximum likelihood (ML) estimation, it was shown how the probability parameter p can be estimated from the data. This empirical MFA method works as follows: first, the network structure is estimated using IsingFit. As there currently is no valid test that can distinguish between a random graph and a small world graph, it is assumed

that the network structure is a random graph, with the ratio of edges that is non-zero and the possible number of edges as the edge probability p_e. The edge probability p_e is needed to create the bifurcation diagram. After that, the probability parameter p is estimated from the data using ML estimation. The estimate of p ($\hat{p}$) is then set off against the bifurcation diagram.

Figure 16 shows an example of such a bifurcation diagram. A bifurcation diagram is created by plugging in several values for p (x-axis), and plotting the resulting expected density μ_p (y-axis). A bifurcation diagram consists of two parts: the first is the area where two lines occur (at $x = 0.1$ to $x \approx 0.40$). This is the area where critical transitions can occur; the two lines in this part of the diagram represent the two phases that an individual can be in. The second part is the area where only one line exists (at $x \approx 0.40$ to $x = 1$): there is only one line, and therefore, critical transitions cannot occur. The empirical MFA compares $\hat{p}$ with the point at which the line splits into two lines, called the critical value. When $\hat{p}$ is smaller than the critical value, we conclude that the individual has an increased risk for experiencing a critical transition. When $\hat{p}$ is higher than the critical transition, we conclude that the individual does not have an increased risk for experiencing a critical transition. For a more detailed and technical description, see Kossakowski et al. (2016).

#FIGURE 16 ABOUT HERE#

5.2 Application of the Empirical MFA to general affect data

Data for our example were collected from and by the first author (hereafter called participant JK). For testing purposes, the participant monitored her daily affects using the *experience sampling method* (ESM; Csikszentmihalyi & Larson, 1987). With ESM, the participant's waking hours are divided by s intervals of equal size. Within these intervals, the participants received a beep at a random time (within that interval) and is requested to fill out a questionnaire, often on a palmtop or smartphone. Over the course of seven days, participant JK (a 27-year-old female) completed 45 measurements, and missed 22 measurements; in total 67 measurements were collected. Participant JK completed the *Positive Affect Negative Affect Scale* (PANAS; Watson, Clark, & Tellegen, 1988), a 20-item questionnaire that contains ten positive mood states and ten negative mood states, each measured on a 5-point Likert scale,

729 ranging from ‘very slightly or not at all’ to ‘extremely’. Table 2 shows the twenty different
730 mood states and their node labels.
731 Positive variables were recoded, so that high scores indicate a more negative affect.
732 Missing measurements were replaced by the previous measurement, after which the variables
733 were dichotomized using a median split, which is necessary for using IsingFit. Finally, we
734 removed six variables due to observing either of two response categories less than four times.
735
736 #FIGURE 17 ABOUT HERE#
737 #TABLE 2 ABOUT HERE#
738
739 Figure 17 (upper panel) depicts the network structure of participant JK, estimated
740 with the IsingFit method. For example, item 1 (feeling interested) has a strong connection
741 with items 3 (feeling excited), 5 (feeling strong) and 13 (feeling ashamed). Interestingly, the
742 association between items 3 and 5 is negative, which means that, as participant JK feels more
743 excited, she feels less strong and the other way around. The association between items 2
744 (feeling distressed) and 4 (feeling upset) makes sense: as participant JK feels more upset, she
745 tends to feel more distressed and the other way around.
746 Figure 17 (middle panel) shows the trajectory of the density (the proportion of active
747 nodes) of participant JK over time. It can be seen that sometimes, the density does not
748 change, especially in the middle of the measurement period, but also that, at several
749 occasions, some peaks occur that are indicative of high general negative affect. However,
750 since there is no clear pattern in the densities, we would expect that participant JK is not
751 vulnerable for a critical transition from a positive affective state to a negative affective state.
752 The lower panel of Figure 17 shows the bifurcation diagram for participant JK, and $\hat{p}$ that we
753 estimated from the data. It shows that $\hat{p}$ is much higher than the critical value, which lies
754 around *0.26*. This means that participant JK is not vulnerable for a critical transition.
755
756 **6. Discussion**
757 In this chapter, we have outlined a network perspective on psychological constructs. We
758 have demonstrated various methods for estimating network structures and, subsequently, we
759 have shown – with the Cramer model and the Empirical MFA – how one can use these
760 network structures to investigate vulnerability and critical transitions: what makes certain
761 people vulnerable for developing psychopathology and can we anticipate critical transitions
762 from, say, health to a disordered state? Pertaining to the latter question: yes, we may. In the

Cramer model, we saw that individuals with more strongly connected networks are more vulnerable for the development of psychopathology. Detecting critical transitions by means of autocorrelations, however, was an only partially successful endeavor. In the Empirical MFA, the vulnerability was estimated from the data: the lower the probability parameter, the more vulnerable an individual is for a critical transition. For our general affect examples, we showed that it is possible to see and potentially anticipate critical transitions from one state to another.

Throughout this entire chapter, we investigated vulnerability for a critical transition that an individual may or may not have. In the mania example described in section 4 the conceptual meaning of the two possible states and the transition is clear: the system is in either a state of mania or not. However, in the general affect example described in section 5 this may not be as clear. In terms of general affect, what are the two states that an individual can be in, and what does the transition between one and the other mean? A vulnerability analysis such as the ones performed in this chapter might be conceptually meaningless if the system's two states are ill-defined or not defined at all. Therefore, not every (psychological) construct might lend itself for such an analysis, and when applying these models, researchers are best advised to first have a clear view on what the two states are in the system, and what it means if an individual transitions between them. Additionally, some psychopathological constructs might have more than two states. Consider bipolar disorder for example: this disorder is characterized by switching between two states, a depressive state with predominantly depressive symptoms (e.g., feeling blue, insomnia) and a manic state with predominantly manic symptoms (e.g., irritability, sexual promiscuity). In this case it would make sense to hypothesize the existence of not 2 but 3 states: one healthy state (e.g., bipolar patient in remission), a depressive state and a manic state. It remains to be seen if one could model a system with three states at all and if so, if the general principles of for example catastrophe theory (e.g., hysteresis, discontinuous behavior) would apply.

In this chapter, we studied critical transitions from a healthy state to a disordered state. We also saw that a strongly connected network is indicative of a vulnerability for a critical transition: when a strongly connected network is put under stress, the system may jump from one state to the other, whereas a weakly connected network, when put under stress, may gradually transition. When healthy individuals are being studied, having a strongly connected network is a disadvantage, for those individuals may experience a critical transition from a healthy state to a disordered state. However, a critical transition can always occur in two ways: from healthy to disordered and from disordered to healthy. In the case of depressed

patients, a critical transition may occur from a disordered to a healthy state. In this particular case, having a strongly connected network is an advantage: if some intervention would succeed in turning a few symptoms "off" then strong connectivity might result in the triggering of additional symptom deactivation. In future research, it is therefore important to study critical transitions both from a healthy →disorder perspective, and from a disorder → healthy perspective.

The ultimate goal of the models shown in this chapter is to assess an individual's vulnerability for a critical transition, before the transition itself occurs. It is quite possible that, by the time individuals enter therapy, it is too late for us to assess their vulnerability; for a critical transition may already have occurred. In such a case, the goal of therapy also transitions from preventing a critical transition to a disordered state, to helping that individual transition back to a healthy state. If we would solely want to prevent critical transitions, we would need to track virtually every individual from the age of 15 onwards in order to assess their vulnerability for a critical transition from a healthy to a disordered state. This would be an impossible task to carry out. Furthermore, many individuals who may never experience mental problems are in this scenario forced to complete questionnaires whose results will never be used. In studying critical transitions, researchers should therefore not only focus on the prevention of critical transition, but also on how critical transitions can be induced, so that therapists can use our models to help patients return to a healthy state faster.

In the empirical examples of both models, missing data were imputed with either zeroes (Cramer model) or with the previous measurement (Empirical MFA). In any type of data collection, but especially with intensive data collection like ESM, participants may fail to fill out certain questions for numerous reasons. The design of the questionnaire may lead to missing data, as certain questions may only be asked when a specific answer is given to some other question, or participants may skip questions and fail to return to these questions. In the case of ESM data, participants may forget their phone or may not simply be in the mood to complete the questionnaire when it is asked of them several times a day. The effect of missing data on the results of a vulnerability analysis is, while writing this chapter, unclear. The methods for dealing with missing data showed in this chapter led to a decrease in item variance, which may result in removing that specific item from the data that is used for the entire analysis. Future research should therefore focus on investigating the effects of the different types of missing data (missing completely at random, such as missing random measurements; missing at random, such as forgetting the phone with which you completed the questionnaires; or missing not at random, such as not being in the mood to complete a

questionnaire before having a cup of coffee in the morning) on the results, and how much data can be missing before results become unreliable. A way of gauging the reliability of network parameter estimates that is currently available for cross-sectional data and that could potentially be used for time-series data, is implemented in R (Epskamp, Borsboom, & Fried, 2016). This software package enables calculating confidence intervals around edge weights by means of bootstrapping these weights. The resulting confidence intervals indicate how safe one is in concluding that the estimated edge weight is reliable (e.g., potentially less reliable if the confidence interval around an edge weight includes zero).

This chapter described two models for studying vulnerability and critical transitions that are currently still in their developmental phase. Even though there are some serious conceptual and technical challenges that need to be overcome (Fried & Cramer, 2016), both the Cramer model and the Empirical MFA show promising results and have the potential to aid clinical psychologists and psychiatrists in both successfully treating patients and preventing relapses by, for example, tracking individuals during therapy in order to find the window of optimal opportunity for such successful interventions. That is, in our view, network modeling of psychopathological phenomena holds the promise of ultimately leading the field of psychopathology into an era of personalized care.

References

Bollobás, B. (2001). *Random Graphs*. book, Cambridge, UK: Cambridge University Press.

Borsboom, D., Cramer, A. O. J., Schmittmann, V. D., Epskamp, S., & Waldorp, L. J. (2011). The small world of psychopathology. *PloS ONE*, *6*(11), e27407.

Carpenter, S. R., & Brock, W. A. (2006). Rising variance: a leading indicator of ecological transition. *Ecology Letters*, *9*, 311–318.

Chen, J., & Chen, Z. (2008). Extended bayesian information criterioa for model selection with large model spaces. *Biometrika*, *95*, 759–771.

Costantini, G., Epskamp, S., Borsboom, D., Perugini, M., Mottus, R., Waldorp, L. J., & Cramer, A. O. J. (2015). State of the aRt personality research: A tutorial on network analysis of personality data in R. *Journal of Research in Personality*, *54*, 13–29.

Cramer, A. O. J., Sluis, S., Noordhof, A., Wichers, M., Geschwind, N., Aggen, S. H., … Borsboom, D. (2012). Dimensions of normal personality as networks in search of equilibrium: You can't like parties if you don't like people. *European Journal of Personality*, *26*, 414–431. article.

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 dynamical system. *PloS One*.

Cramer, A. O. J., Waldorp, L. J., van der Maas, H. L. J., & Borsboom, D. (2010). Comorbidity: a network perspective. *Behavioral and Brain Sciences*, *33*, 137–193. article.

Csikszentmihalyi, M., & Larson, R. (1987). Validity and reliability of the experience-sampling method. *The Journal of Nervous and Mental Disease*, *175*, 526–536. article.

Dakos, V., Scheffer, M., van Nes, E. H., Brovkin, V., Petoukhov, V., & Held, H. (2008). Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences USA*, *105*, 14308–14312. article.

Dalege, J., Borsboom, D., van Harreveld, F., van den Berg, H., Conner, M., & van der Maas, H. L. J. (2016). Toward a Formalized Account of Attitudes:The Causal Attitude Network (CAN) Model. *Psychological Review*, *123*, 2–22.

Dolan, C. V, Oort, F. J., Stoel, R. D., & Wicherts, J. M. (2009). Testing Measurement Invariance in the Target Rotated Multigroup Exploratory Factor Model. *Structural Equation Modeling: A Multidisciplinary Journal*, *16*(2), 20. article.

Durett, R. (2007). *Random Graph Dynamics*. book, Cambridge, UK: Cambridge University Press.

884 Epskamp, S. (2016). Regularized Gaussian Psychological Networks: Brief Report on the
885 Performance of Extended BIC Model Selection.
886 Epskamp, S., Borsboom, D., & Fried, E. I. (2016). Estimating psychological networks and
887 their accuracy: a tutorial paper.
888 Epskamp, S., Cramer, A. O. J., Waldorp, L. J., Schmittmann, V. D., & Borsboom, D. (2012).
889 qgraph: Network Visualizations of Relationships in Psychometric Data. *Journal of*
890 *Statistical Software*, *48*, 1–18. article. Retrieved from http://www.jstatsoft.org/v48/i04/
891 Epskamp, S., & Fried, E. I. (2016). A tutorial on Regularized Partial Correlation Networks.
892 Epskamp, S., Waldorp, L. J., Mottus, R., & Borsboom, D. (2016). Discovering Psychological
893 Dynamics: The Gaussian Graphical Model in Cross-sectional and Time-series Data.
894 Fort, H., Mazzeo, N., Scheffer, M., & van Nes, E. H. (2010). Catastrophic shifts in
895 ecosystems: spatial early warnings and management procedures. *Journal of Physics*
896 *Conference Series*, *246*, 1–15. article.
897 Foygel, R., & Drton, M. (2010). Extended Bayesian Information Criteria for Gaussian
898 Graphical Models. In *NIPS* (pp. 604–612). inproceedings.
899 Fried, E. I., & Cramer, A. O. J. (2016). Moving forward: challenges and directions for
900 psychopathological network theory and methodology.
901 Fried, E. I., van Borkulo, C. D., Cramer, A. O. J., Boschloo, L., Schoevers, R. A., &
902 Borsboom, D. (2016). Mental disorders as networks of problems: a review of recent
903 insights.
904 Hartman, C. A., Hox, J., Mellenbergh, G. J., Boyle, M. H., Offord, D. R., Racine, Y., …
905 Sergeant, J. A. (2001). DSM-IV Internal Construct Validity: When a Taxonomy Meets
906 Data. *Journal of Child Psychology and Psychiatry*, *42*, 817–836. article.
907 Haslbeck, J. M. B., & Waldorp, L. J. (2016a). mgm: Structure Estimation for time-varying
908 Mixed Graphical Models in high-dimensional Data. *Submitted for Publication*.
909 Haslbeck, J. M. B., & Waldorp, L. J. (2016b). Structure estimation for mixed graphical
910 models in high-dimensional data. *Submitted for Publication*.
911 Ising, E. (1925). Beitrag zur theorie des ferromagnetismus. *Zeitschrift Für Physik A Hadrons*
912 *and Nuclei*, *31*, 253–258. article.
913 Kessler, R., Adler, L., Berglund, P., Green, J., McLaughlin, K., Fayyad, J., … Zaslavsky, A.
914 (2014). The effects of temporally secondary co-morbid mental disorders on the
915 associations of DSM-IV ADHD with adverse outcomes in the US National Comorbidity
916 Survey Replication Adolescent Supplement (NCS-A). *Psychological Medicine*, *44*,
917 1779–1792.

Kessler, R. C., Berglund, P., Chiu, W. T., Demler, O., Heeringa, S., Hiripi, E., … Zheng, H. (2004). The US National Comorbidity Survey Replication (NCS-R): Design and field procedures. *International Journal of Methods in Psychiatric Research*, *13*, 69–92. article.

Kessler, R. C., & Ustun, T. B. (2004). The World Mental Health (WMH) survey initiative version of the World Health Organization (WHO) Composite International Diagnostic Interview (CIDI). *International Journal of Methods in Psychiatric Research*, *13*, 93–121. article.

Kossakowski, J. J., Epskamp, S., Kieffer, J. M., van Borkulo, C. D., Rhemtulla, M., & Borsboom, D. (2016). The application of a network approach to Health-Related Quality of Life (HRQoL): introducing a new method for assessing HRQoL in healthy adults and cancer patients. *Quality of Life Research*, *25*, 781–792.

Kossakowski, J. J., Gordijn, M. C. M., Riese, H., & Waldorp, L. J. (2016). Mean field dynamics of graphs II: Assessing the Risk for the Development of Phase Transitions in Empirical Data. article.

Kossakowski, J. J., Groot, P. C., Haslbeck, J. M. B., Borsboom, D., & Wichers, M. (2016). Data from "Critical Slowing Down as a Personalized Early Warning Signal for Depression."

McNally, R. J., Robinaugh, D. J., Wu, G. W. Y., Wang, L., Deserno, M. K., & Borsboom, D. (2014). Mental Disorders as Causal Systems: A Network Approach to Posttraumatic Stress Disorder. *Clinical Psychological Science*, *3*, 836–849.

Newman, M. (2010). *Networks: an introduction*. book, Oxford Univ Pr.

Newman, M. E. C., & Watts, D. J. (1999). Renormalization group analysis of the small-world network model. *Physics Letters A*, *263*, 4–6. article.

Olde Rikkert, M. G. M., Dakos, V., Buchman, T. G., de Boer, R., Glass, L., Cramer, A. O. J., … Scheffer, M. (2016). Slowing Down of Recovery as Generic Risk Marker for Acute Severity Transitions in Chronic Diseases. *Critical Care Medicine*, *44*, 601–606. article.

R Core Team. (2016). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.r-project.org/

Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., … Sugihara, G. (2014). Early-warning signals for critical transitions. *Nature*, *46*, 53–59. article.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal*

952 *Statistical Society. Series B (Methodological)*, *58*, 267–288. article.
953 Tio, P., Epskamp, S., Noordhof, A., & Borsboom, D. (2016). Mapping the manuals of
954 madness: Comparing the ICD-10 and DSM-IV-TR using a network approach.
955 *International Journal of Methods in Psychiatric Research*.
956 http://doi.org/10.1002/mpr.1503
957 van Borkulo, C. D., Borsboom, D., Epskamp, S., Blanken, T. F., Boschloo, L., Schoevers, R.
958 A., & Waldorp, L. J. (2014). A New Method for Constructing Networks from Binary
959 Data. *Scientific Reports*, *4*, 1–10. article.
960 van Borkulo, C. D., Boschloo, L., Borsboom, D., Penninx, B. W. J. H., Waldorp, L. J., &
961 Schoevers, R. A. (2015). Association of symptom network structure with the course of
962 depression. *JAMA Psychiatry*, *72*, 1219–1226. article.
963 van der Maas, H. L. J., Dolan, C. V, Grasman, R., Wicherts, J. M., Huizenga, H. M., &
964 Raijmakers, M. E. J. (2006). A dynamical model of general intelligence: The positive
965 manifold of intelligence by mutualism. *Psychological Review*, *113*, 842–861. article.
966 Waldorp, L. J., & Kossakowski, J. J. (2016). Mean Field Dynamics of Graphs I: Evolution of
967 Probabilistic Cellular Automata on Different Types of Graphs. article.
968 Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and Validation of Brief
969 Measures of Positive and Negative Affect: The PANAS Scales. *Journal of Personality*
970 *Social Psychology*, *54*, 1063–1070.
971 Wichers, M., Groot, P. C., Psychosystems, ESM Group, & ESW Group. (2016). Critical
972 Slowing Down as a Personalized Early Warning Signal for Depression. *Psychotherapy*
973 *and Psychosomatics*, *85*, 114–116.
974 Wild, B., Eichler, M., Friederich, H.-C., Hartmann, M., Zipfel, S., & Herzog, W. (2010). A
975 graphical vector autoregressive modelling approach to the analysis of electronic diary
976 data. *BMC Medical Research Methodology*, *10*, 1–13. article.
977 Zhao, P., & Yu, B. (2006). On model selection consistency of lasso. *The Journal of Machine*
978 *Learning Research*, *7*, 2541–2563.
979