The Lindy effect in psychology: Network dynamics reinforce depression over time

Eren Asena

University of Amsterdam

Abstract

This paper studies the factors that sustain mental disorders by taking a network approach. The network theory suggests that mental disorders are networks of symptoms that causally interact (Borsboom, 2017). Symptom networks share certain dynamics with other complex systems: abrupt transitions between stable states, critical slowing down and hysteresis (Cramer et al., 2016). These findings suggest that symptom networks that have transitioned to a pathological state tend to remain that state. We argue that this tendency leads to the Lindy effect in symptom networks. The Lindy effect means that the conditional probability of surviving beyond a time point, given survival until that time point, increases over time (Taleb, 2014). In other words, time benefits future survival. A symptom network is considered to have survived until a time point if it has remained in a pathological state until that point. We first show how the Lindy effect is formalised by examining the stopping time distribution of Brownian motion with an absorbing barrier (Cook, 2012; Taleb, 2018). Specifically, we describe the hazard function of the stopping time distribution and make a distinction between "strong Lindy" and "weak Lindy". Strong Lindy is a monotonically decreasing hazard function whereas weak Lindy means a $\cap$ shaped hazard function. Then, major depressive disorder (MDD) networks were simulated, manipulating the level of symptom connectivity. As before, the presence of the Lindy effect in these networks were tested using hazard functions, and in addition, survival probabilities conditioned on time. Afterwards, we fit a distribution to the network lifetimes. The lifetime distribution of strongly connected networks were heavy tailed and showed the Lindy effect; the longer a network had been depressed, the more likely it was to remain depressed. The lifetime distribution of weakly connected networks were light tailed and did not show the Lindy effect. After discussing caveats and alternative explanations of the findings, we conclude that network dynamics and the resulting Lindy effect can explain several findings in psychology such as the chronicity of depression (Swaminath, 2009) and the frequency distribution of remission times (Simon, 2000; Patten et al., 2010).

The Lindy effect in psychology: Network dynamics reinforce depression over time

The World Health Organization (WHO) considers depression as the largest cause of global disability, accounting for 7.5% of all years lived with disability (WHO, 2017). Current treatments of depression are inadequate to bear such a high burden. This is apparent from the fact that many individuals with depression do not respond to treatment and for many of them the disorder takes a chronic course (Kelley et al., 2018; Swaminath, 2009). Understanding how depression becomes chronic and why it is hard to treat can help develop better treatment plans. This paper aims to contribute to this understanding. Specifically, we build on the network theory of mental disorders (Borsboom, 2017) to investigate what sustains major depressive disorder (MDD) over time.

Mental disorders as symptom networks

The network theory suggests that mental disorders are networks of symptoms (Borsboom, 2017). That is, each symptom is a node in a network and causally interacts with the other symptoms (Borsboom & Cramer, 2013). An external stressor may activate one of the symptoms in the network, which can activate others. As a result, an entire network of symptoms can be activated, forming the disorder. For example, losing one's spouse can trigger depressed mood. The depressed person may feel guilty about the last fight they had with their partner. The feeling of guilt can make them ruminate at night, causing sleep problems. Consequently, the person may feel fatigued, which can diminish their interest in normally pleasant activities such as spending time with friends. In turn, withdrawing from social interactions can remove many positive experiences, which can worsen depressed mood. In other words, not only can the activation of a symptom spread through the network and create a whole syndrome, but the symptom-symptom interactions that cause this spread can form positive feedback loops, reinforcing each other in vicious cycles.

Modeling symptom networks

Symptom networks can be modelled using the Ising model. (Epskamp, Maris, Waldorp, & Borsboom, 2018). The Ising model is used in statistical physics to describe phase transitions in a field of particles. In symptom networks modelled this way, the nodes (i.e., symptoms) can be in one of two states: on or off. A symptom being ‘on’ means that the symptom is active. Likewise, a symptom is ‘off’ if it is not active. Neighbouring symptoms are connected to each other by weight parameters (Epskamp et al., 2018). A positive weight means that two symptoms tend to be in the same state whereas a negative weight means that the symptoms tend to be in different states. The absolute value of the weight determines the strength of this relationship. The connectivity of the nodes (i.e., how strongly connected the symptoms are) can be manipulated by multiplying the weight parameters of neighbouring nodes with a connectivity parameter.

Symptom networks as complex dynamical systems

Symptom networks show interesting behaviours that are common among complex dynamical systems. A system is complex and dynamical if its behaviour at the macro level changes depending on the interactions between its components (Mitchell, 2011). For example, the behaviour of the brain evolves as a result of the interactions among many neurones, neurotransmitters, etc. Complex dynamical systems with strongly connected components have tipping points, at which the system abruptly transitions from one stable state to another (Scheffer et al., 2012). For example, a seemingly healthy banking system can suddenly collapse. Initially, the system will recover from small perturbations (e.g., a bank having difficulties can recover with subsidies). However, over time the banking system will reach a tipping point at which the failure of a few large banks will spread throughout the network and cause sudden collapse of the entire system.

The cusp catastrophe model provides a useful way to model these abrupt shifts between distinct states (Thom, 1973, as cited in Grasman, van der Maas, & Wagenmakers,

2009; Zeeman, 1971, as cited in Stewart & Peregoy, 1983). In cusp catastrophes, the state of a system is determined by two variables called the control variables (Flay, 1978). The first control variable is called the normal factor and the second is called the splitting factor. For low values of the splitting variable, increases in the normal factor causes continuous and gradual changes in the system. However, if the splitting variable goes up, a small change in the normal variable near the tipping point will make the system abruptly transition to another stable state. Consequently, the system will have discontinuous states; it will not evolve gradually but jump from one state to another.

In a simulation study, Cramer et al. (2016) used the cusp catastrophe model to show that MDD networks can have discontinuous states. Specifically, external stress was used as the normal variable and connectivity as the splitting variable. The researchers stressed the symptom network in increasing amounts over time. When a tipping point was reached, highly connected networks abruptly transitioned from a non-depressed state to a depressed state. In contrast, symptom activation did not spread as much in weakly connected networks. Therefore, they managed to recover quickly and did not transition to a depressed state. In line with these findings, Hosenfeld et al. (2015) observed that most MDD patients in their sample experienced depression as two distinct states. This was indicated by the fact that the frequency distribution of the symptoms was bimodal for the majority of the sample. To put it simply, the participants either experienced many symptoms together or almost no symptoms at all. Also, the probability of remaining in either the depressed or the non-depressed state was high, suggesting that both states were stable.

Complex dynamical systems give warning signals before reaching tipping points. One of these signals is critical slowing down (Scheffer et al., 2012). Critical slowing down happens when the system takes an increasingly longer period of time to recover from small perturbations. In MDD, this may translate to remaining sad for a long time after a minor dispute with a loved one (Cramer et al., 2016). Critical slowing down is reflected in increased autocorrelations between the states of the system over subsequent time points.

Since the system takes some time to recover from a perturbation, its state at subsequent time points resemble each other. Cramer et al.'s (2016) MDD networks showed increased autocorrelations before transitioning to a depressed state. Finally, complex dynamical systems show hysteresis (Stewart & Peregoy, 1983). Hysteresis means that if a system has transitioned to another state, the amount of change in the normal variable that is necessary to bring the system back to the first state is larger than the amount that caused the initial transition. In Cramer et al.'s (2016) MDD networks, the amount of stress reduction needed to bring a depressed network back to the non-depressed state was larger than the amount of stress that initially triggered the depression. Importantly, the hysteresis effect was stronger for highly connected networks. It seems that when highly connected networks transition to a pathological state, they tend to remain in that pathological state until stress is greatly reduced. To an extent, they become self-sustaining – the person remains depressed even after the stressful life event that triggered the depression is gone.

It will be beneficial to give an interim summary at this point. We mentioned that mental disorders such as MDD can become chronic and resistant to change. We also argued that the network theory of psychopathology can help us understand why. The explanatory power of the network theory resides in the characteristics which symptom networks share with other complex dynamical systems: discontinuous jumps between pathological and non-pathological states, increased autocorrelations indicating a tipping point, and hysteresis. These findings imply that once a network transitions to a depressed state, it tends to remain in that state. Furthermore, these characteristics are observed because symptoms causally interact with each other, forming feedback loops. These feedback loops reinforce the depressed state over time. The present study gives a proof of concept that as a consequence of this reinforcing tendency, highly connected networks show a phenomenon called the Lindy effect (Mandelbrot, 1982; 2004; Taleb, 2014).

The Lindy effect

The Lindy effect refers to the idea that the remaining life expectancy of a nonperishable thing (i.e., it does not have an organic limit to its lifetime) is directly proportional to its age (Taleb, 2014). That is, the longer it has survived, the longer it is expected to survive in the future. For example, the Lindy effect implies that the Bible is more likely to be around in the next century than a recently published novel, because it has already survived for thousands of years. It has stood the test of time. The Lindy effect implies *antifragility* through time (Taleb, 2014). Whereas fragile things are harmed by stress / volatility, *antifragile* things benefit from it. Time being a stressor, an artefact that obeys the Lindy effect benefits from time. We suggest that the Lindy effect sheds light on the evolution of mental disorders over time, including how they become chronic. Specifically, we argue that the Lindy effect is implied by network dynamics, explain how it can elucidate the progression of mental disorders, describe ways of detecting it in data, and test its presence in MDD networks.

The Lindy effect and mental disorders

The lifetimes of nonperishable things are not constrained by organic limits (N.B., throughout this paper, the terms ‘lifetime’, ‘failure time’, ‘death time’, ‘event time’, ‘stopping time’ and ‘hitting time’ have the same meaning in different contexts. They all refer to the total time that has passed from the starting point of a process until the final point of the process.). Without a natural upper bound, the distribution of an event time “...is constrained only by fragility” and follows a power law (Taleb, 2014, p. 317; see Newman, 2005 or the Wikipedia page “Power law,” n.d. for an overview of power law distributions). This makes the remaining life expectancy proportional to past survival. For example, the life expectancy of a 100-year-old book is expected to be ten times as long as the life expectancy of a 10-year-old book (Taleb, 2014). However, life expectancy is a “probabilistically derived average” (Taleb, 2014, p. 319), meaning the predictions will be

far from perfect.

Psychological disorders are technically perishable. The most obvious way to see this is that they will go away when the person dies. Therefore, it is likely that their lifetimes are not solely determined by fragility and their time distribution is not really a power law. However, the Lindy effect does not necessarily require a power law distribution and symptom networks can be *antifragile* within limited time frames. This is because we can observe increasing additional life expectancy in various heavy-tailed distributions. For example, Xie (2020) found that the viewing durations of online classes followed a lognormal distribution with the Lindy effect in its tail. The more the students viewed a class, the more they wanted to learn – at least after they had viewed it for a while.

The point is that given sufficient time, the dynamics of a system can improve its chances of survival. When we say a mental disorder has ‘survived’, we mean that it is still present. This means that the Lindy effect is less likely to apply to the symptoms of a 90-year old than to the symptoms of a 20-year-old (assuming these are the ages of depression onset). Although it is not clear at what point time starts to benefit survival, the chronicity of certain mental disorders such as MDD and the network dynamics we discussed above all indicate that the Lindy effect may be present in mental disorders. Our argument is as follows: the symptom interactions which allow a highly connected network to become self-sustaining will continue after the stress is gone. As a result, the symptoms will keep reinforcing each other, holding the network in the depressed state long after the tipping point. Consequently, the probability of node failure will decrease over time and the probability of remaining in the depressed state will increase. Thus, the disorder will show the Lindy effect.

Finding the Lindy effect

To test if the Lindy effect is present in symptom networks, we need a way of detecting it in general. First, the Lindy effect conceptually means that the chances of

survival increases with age. In the case of symptom networks, it means that the probability of remaining in the pathological state will increase with the duration of the disorder. This definition is important: a symptom network is considered to have ‘survived’ until a time point if it has remained in the pathological state until that time point, and it is considered to have died or failed if it has transitioned back to the non-depressed state.

Having defined the Lindy effect conceptually, let us look at its properties more technically. A random variable obeys the Lindy effect if the conditional expectation of its remaining lifetime after a time point, given that it has survived until that point, increases (Eliazar, 2017). Whether or not this is the case can be tested by means of survival analysis (Aalen & Gjessing, 2001). Survival analysis examines the time between a starting point and an event such as death. Our goal is to analyse the probability distribution of this time, using hazard functions and conditional survival probabilities (explained below). In a recent work, Taleb (2021) used the same approach to formalise the Lindy effect.

If the conditional probability of survival increases, then the conditional probability of death within a time interval should decrease. Dividing this probability by the length of the time interval, we get the hazard function. The hazard function is also called the “force of mortality”, or the instantaneous failure rate. Let X be the time random variable. Writing this concept formally:

$$\lim_{\Delta t \rightarrow 0} \frac{P(t < X \leq t + \Delta t | X > t)}{\Delta t}. \quad (1)$$

Note that since this is a rate and not a probability, it can take values larger than 1. Hazard rates can be computed from the expression

$$h(t) = \frac{f(t)}{S(t)}, \quad (2)$$

where $f(t)$ is the density function and $S(t)$ is the survival function. The survival function gives the probability that the random variable will be larger than a given value, $P(T > t)$. It is defined as $S(t) = 1 - F(t)$, where $F(t)$ is the cumulative distribution function (CDF).

The CDF gives the probability that the lifetime will take a value less than or equal to a given time point. Thus, the survival function is the complement of the CDF. Figure 1 visualises different hazard functions associated with different probability distributions. Alternatively, and perhaps more intuitively, the Lindy effect can be formalised as increasing conditional survival probabilities. The conditional probability of surviving further t time points given past survival of s time points (Hieke et al., 2015) is given by:

$$CS(t|s) = \frac{S(s+t)}{S(s)} \quad (3)$$

To see how these patterns would appear in data and how they can be detected, it is helpful to work with a stochastic process that can generate a distribution of lifetimes. Taleb (2018; 2021) used Brownian motion as this process and we followed suit.

Brownian motion: Defining strong vs. weak Lindy

In 1827, the botanist Robert Brown noted that the motion of pollen grains in water looked random (Brownian Motion, n.d.). Einstein showed in 1905 that the the behaviour of particles in a fluid was indeed random and independent of past motion (Dahl, 2010; Brownian motion, n.d.). Although Brownian motion was initially celebrated for proving the existence of atoms and molecules, the random fluctuations it describes can be seen in various phenomena. For example, financial analysts use Brownian motion to model stock prices. Brownian motion can take three main forms called standard Brownian motion, arithmetic Brownian motion (i.e., BM with drift) (ABM) and geometric Brownian motion (GBM) (Pacati, 2010). These are visualized and described in Figure 2.

A real valued stochastic process W_t is a standard Brownian motion (also called a Wiener process) if it has the following properties (Dahl, 2010):

1. The process starts at $W_0 = 0$.
2. The process has independent increments: for all $t > 0$, the increment $W_{t+s} - W_t$ is an independent random variable.

3. The process has normal increments: for all $t \geq 0$ and $s \geq 0$, the increment $W_{t+s} - W_t$ is normally distributed with mean 0 and variance s , $W_{t+s} - W_t \sim N(0, s)$.
4. The process is continuous: W_t is almost surely continuous in t .

A stochastic process S_t is an Arithmetic Brownian motion if it is described by the stochastic differential equation (SDE)

$$dS_t = \mu dt + \sigma dW_t, \quad (4)$$

with initial value $S_0 = s_0$. This gives (see Pacati (2010), p. 4-6 for the derivations):

$$S_t = s_0 + \mu t + \sigma W_t, \quad (5)$$

where μ is the drift term and is negative in this paper (Taleb, 2021). Having negative drift means that the process goes down on average, despite taking random steps (Asena, 2021). With positive drift it would go up. In contrast, standard Brownian motion does not tend towards any direction. It is called *arithmetic* Brownian motion because the drift term scales only with the time increment, but not the current value. Therefore, it has an additive effect on the process. Geometric Brownian motion is the solution to the following SDE:

$$dS_t = \mu S_t dt + \sigma S_t dW_t, \quad (6)$$

which is

$$S_t = S_0 e^{\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W_t}. \quad (7)$$

Unlike in ABM, the GBM sample paths have a curve because GBM is a multiplicative process. Whereas the drift term is constant and is added to the current value in ABM, it is linear and is multiplied by the current value in GBM (Wu & Buyya, 2015). Also, standard BM starts from 0 by definition, but ABM and GBM do not have to.

Brownian motion is relevant for this paper because it arises naturally when modelling death / failure times. The reason is that death is a discrete event and the time of death

can be represented by the stopping time of Brownian motion (Taleb, 2018). The stopping time of Brownian motion is the first time the process reaches a threshold value B (MIT, 2013). For $W_0 > B$ (i.e., BM approaches the threshold from above), this is defined as:

$$\tau_B = \inf \{t : W_t \leq B\}. \quad (8)$$

The threshold B is called the absorbing barrier. Standard BM and ABM with absorbing barriers are visualized in Figure 3.

The stopping time distribution of Brownian motion with a constant absorbing barrier is well known for both the standard and the drifted cases. Therefore, we do not re-derive them and instead point the reader to sources that do. We derive the hazard functions of some other important distributions in the Appendix and they will be discussed in the main text as well. To begin, the stopping time of standard Brownian motion follows a Levy distribution (see Siegrist (2020) for the proof) with a location parameter μ and scale parameter c ("Levy distribution", n.d.):

$$f(x; \mu, c) = \sqrt{\frac{c}{2\pi}} \frac{e^{-\frac{c}{2(x-\mu)}}}{(x-\mu)^{3/2}}, \quad \mu \leq x < \infty. \quad (9)$$

The scale parameter is related to the absorbing barrier as $c = B^2$. The asymptotic behaviour of the density is to decay as a power law (see Appendix B), which means that the Levy distribution is fat-tailed. The CDF of the Levy distribution is ("Levy distribution", n.d.):

$$F(x; \mu, c) = \text{erfc} \left(\sqrt{\frac{c}{2(x-\mu)}} \right) \quad (10)$$

where $\text{erfc}(x) = 1 - \text{erf}(x)$ is the complementary error function and the error function is given by:

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt \quad (11)$$

The survival function does not have a closed form, so we plot the hazard function from the formula:

$$h(x) = \frac{f(x)}{S(x)} = \frac{f(x)}{1 - F(x)} \frac{\sqrt{\frac{c}{2\pi}} \frac{e^{-\frac{c}{2(x-\mu)}}}{(x-\mu)^{3/2}}}{1 - \operatorname{erfc}\left(\sqrt{\frac{c}{2(x-\mu)}}\right)} \quad (12)$$

We visualize the density and hazard function in Figure 4.

We can see that the hazard function has an inverted U ($\cap$) shape. That is, it initially increases, reaches a maximum and then decreases asymptotically (we derive the asymptotic behaviour of the hazard function in Appendix B). That is, standard Brownian motion shows the Lindy effect only after a certain time point, or as Taleb puts it, there is a "pre-Lindy" period (Taleb, 2021, p. 53). The first author called this "weak Lindy" in an online article (Asena, 2021). This is contrasted with "strong Lindy", which is a non-increasing hazard function. The stopping time of ABM with negative drift and an absorbing barrier $B < S_0$ follows an inverse Gaussian distribution (IG) (Primožic, 2011, p. 9), $\tau \sim IG\left(\frac{B-S_0}{\mu}, \frac{(B-S_0)^2}{\sigma^2}\right)$. Now, if $X(t)$ is a GBM, then $Y(t) = \log X(t)$ is an ABM. Thus, the distribution of the new stopping time ρ of GBM with an absorbing barrier $L < S_0$ is also an IG, but with log-transformed parameters, $\rho \sim IG\left(\frac{\log L - \log S_0}{\mu - \frac{1}{2}\sigma^2}, \frac{(\log L - \log S_0)^2}{\sigma^2}\right)$ (Primožic, 2011, p. 9). The PDF and the CDF of the IG are ("Inverse Gaussian distribution", n.d.):

$$f(x; \mu, \lambda) = \sqrt{\frac{\lambda}{2\pi x^3}} e^{-\frac{\lambda(x-\mu)^2}{2\mu^2 x}}, \quad 0 < x < \infty \quad (13)$$

$$F(x; \mu, \lambda) = \Phi\left(\sqrt{\frac{\lambda}{x}}\left(\frac{x}{\mu} - 1\right)\right) + e^{\frac{2\lambda}{\mu}} \Phi\left(-\sqrt{\frac{\lambda}{x}}\left(\frac{x}{\mu} + 1\right)\right), \quad (14)$$

where μ is the mean, λ is the shape parameter, and $\Phi(x)$ is the standard normal distribution function. Again, we cannot reach a closed form survival function, so we plot the hazard function from the formula:

$$h(x) = \frac{f(x)}{S(x)} = \frac{f(x)}{1 - F(x)} = \frac{\sqrt{\frac{\lambda}{2\pi x^3}} e^{-\frac{\lambda(x-\mu)^2}{2\mu^2 x}}}{1 - \Phi\left(\sqrt{\frac{\lambda}{x}}\left(\frac{x}{\mu} - 1\right)\right) + e^{\frac{2\lambda}{\mu}} \Phi\left(-\sqrt{\frac{\lambda}{x}}\left(\frac{x}{\mu} + 1\right)\right)} \quad (15)$$

Figure 5 shows the density and the hazard functions of the IG distribution. We see that as is the case with the Levy distribution, the hazard function of the IG distribution has a $\cap$ shape, increasing until it reaches a peak and then decreasing (Lemeshko et al., 2010). In conclusion, Brownian motion does not seem to produce strong Lindy. Then what does? We later show that the Pareto distribution (Appendix C) and, under certain conditions, the Weibull (Appendix D) and the Gamma (Appendix E) distributions, give strong Lindy. Pareto is a power law distribution. Although the Weibull and the Gamma distributions can have heavy tails, they are not as heavy as in the case of Pareto. This fact is important in combination with the results we obtain from Brownian motion; the Levy case shows that fat tailed distributions do not necessarily give strong Lindy, and more importantly, the Weibull and the Gamma cases show that strong Lindy does not need power laws. Finally, we want to note that the lognormal distribution, which is not a power law distribution but has tails heavy enough to be confused with Pareto, gives weak Lindy (see Appendix F). The lognormal may or may not prove to be important for the role of the Lindy effect in psychology. "Lognormal arises when many random variables cooperate multiplicatively" (Boccara, 2003, p.324), which is exactly the type of dynamic we are hypothesising in this paper. Furthermore, Xie (2020) gives evidence that the time distribution of behaviours can be lognormal, with the Lindy effect in the tail. To sum up this section, Figure 6 gives an overview of the relationship between Brownian motion and the Lindy effect.

Current study

Aims

The overarching goal of this study was to study the Lindy effect in symptom networks. Under this general goal, we had two aims: a) to show how the Lindy effect can

be detected in data and b) to test the presence of the Lindy effect in MDD symptom networks. We showed with the stopping time distribution of Brownian motion and the appendix how the Lindy effect can be demonstrated using hazard functions. In the remainder of this paper, we first simulate networks to generate a distribution of depression durations. Then, we estimate the hazard rates and the conditional survival probabilities directly from the generated sample, using non-parametric methods. Finally, we approximate the empirical distribution of the network lifetimes with a theoretical distribution (i.e., fit a distribution to data). As mentioned, hazard functions are one of the several functions that characterise probability distributions. In a way, the most important question concerning the relationship between the Lindy effect and depression is the true distribution of depression duration. After all, if the true distribution has a decreasing hazard function, then depression must obey the Lindy effect. However, identifying a true distribution from sample data is not a trivial task. This is particularly the case because as we will further discuss, we had a moderate sample size due to computational limitations and used a particular network model that was highly sensitive to parameter choices. For these reasons, we hope to reach more confident conclusions by combining non-parametric estimation of hazards directly from the sample and fitting a distribution to the data.

Hypotheses

This study was largely exploratory. We experimented with many parameter values in the simulations. The choices we made are justified and compared to the alternative choices we could have made in the results and the discussion sections. The only confirmatory testing was on the hypothesis that strongly connected networks would show the Lindy effect, whereas the weakly connected ones would not. This would be demonstrated by a decreasing hazard function characterising the lifetime distribution of strongly connected networks.

Methods

Network simulations

We simulated 1000 MDD symptom networks, each for 10,000 time points. Each additional network costs lots of computational resources with current algorithms. Therefore, we could not simulate more than 1000. Since we expected the Lindy effect to be a natural consequence of the symptom-symptom interactions and the hysteresis effect observed by Cramer et al. (2016), we based our simulations on theirs. The networks were estimated based on cross-sectional data from the Virginia Adult Twin Study of Psychiatric and Substance Use Disorders (VATSPUD), which is a longitudinal study of 8973 Caucasian twins from the Mid-Atlantic Twin Registry (Prescott, Aggen, & Kendler, 2000; Kendler & Prescott, 2006, both as cited in Cramer et al., 2016). The data includes the presence or absence of 14 depression symptoms based on the DSM-III-R, lasting at least 5 days during the previous year (American Psychological Association, 1987). Actually, the DSM includes nine depression symptoms but some of them are collapsed. For example, weight loss and weight gain are defined as a single symptom in the DSM but are separate symptoms in the data set. This is why the data set have more symptoms. The presence or the absence of a symptom was binary (i.e., coded as 1 if present, 0 if absent). The empirical weight and threshold parameters were estimated from this data set using the Ising model described earlier in the introduction. The model was implemented using the R-package **IsingFit** (van Borkulo et al., 2014; van Borkulo, Epskamp, & Robitzsch, 2016). Each symptom at each time point was represented as an activation function:

$$A_i^t = \sum_{j=1}^J cW_{ij}X_j^{t-1} \quad (16)$$

The activation of a symptom i at time t is the weighted summation of the neighbouring symptoms X (a matrix of 1s and 0s) at time $t - 1$, multiplied by the connectivity parameter c . That is, the connectivity parameter scales the weights. For each network, we

randomly sampled the connectivity parameter from a range of values (e.g., 1.15 – 1.30, increased by 0.00015, leading to 1001 values to sample from). The choice of the connectivity parameter deserves its own discussion, so we will save it to the discussion section. The activation of a symptom went into the sigmoid function to output a number between 0 and 1. This was the probability of the symptom being on or off (i.e., 1 or 0).

$$P(X_i^t = 1) = \frac{1}{1 + e^{(b_i - A_i^t)}}, \quad (17)$$

where the term b_i is the threshold of the symptom i . This is the general tendency of a symptom to be present or absent. The smaller the threshold and the larger the activation, the more likely it is that the symptom will be on. A symptom was 1 / on, if its probability was larger than a random uniformly distributed number between 0 and 1.

A network was considered depressed if at least five out of nine symptoms were active at a given time point (American Psychological Association, 2013). The most up to date manual, the DSM-V, say that at least five symptoms should be active during the same 14-day period. However, we ignored the additional time specification because the data did not involve responses given to questions that were posed in terms of momentary assessment. Also, the time period is somewhat arbitrary and there is no substantive reason to expect a precisely 14-day period to make a difference in the patterns we observe. Finally, a network was considered to have transitioned back to a non-depressed state when the total number of active symptoms decreased to below five.

In the subsequent survival analysis, the status of the networks that transitioned to a non-depressed state were coded as 1 (i.e., the MDD network ‘failed’) and the time point at which the transition happened was the failure time, or the ‘lifetime’ of the network. On the other hand, the status of the networks that never transitioned back to a non-depressed state were coded as 0 (i.e., censored at the last time point) and the last time point was the censoring time. We hit each network with external stress at its second time point. So, the activation function at the second time point was:

$$A_i^t = \sum_{j=1}^J cW_{ij}X_j^{t-1} + S_i^t \quad (18)$$

However, all the remaining values were generated without stress. The purpose of adding stress at the second time point was to force all the networks to transition to a depressed state from the beginning. This made sense because we were not interested in what made the networks depressed. Rather, we were interested in how long the networks remained in a depressed state once they transitioned. We started counting time at the second point, after most networks (approximately 85%) became depressed.

For each network, we randomly picked an integer between 1 and 15 and used it as the stress parameter. We specifically picked this range because it was just large enough to make the majority of the networks (approximately 85%) transition to a depressed state but not large enough to activate all of the functions at once. As a result, about half of the networks started the depressed state with 6-7 symptoms and about half started with 10-14 symptoms. Those networks that did not transition to a depressed state when hit by stress were coded as dead from the beginning. This naturally increased the number of events at the first time point. Excluding these networks from the analysis instead of considering them as events did not substantially change the results, since they constituted a small fraction of all the networks to begin with. Finally, we computed the hazard rates and the conditional survival probabilities of the depressed state for each time point. We also estimated the probability distribution of the transition times using various methods that are explained below.

Data analysis

Computing the hazard rates and the conditional survival probabilities.

The **survival** package (Therneau, 2020) was used to obtain Kaplan-Meier survival estimates. Using these estimates, the **bshazard** package (Rebora, Salim, & Reilly, 2018) carried out non-parametric smoothing of the hazard rates based on B-splines. This reduces

the problem of estimating the hazards to estimating coefficients in a Generalised Linear Model framework (Rebora, Salim, & Reilly, 2018). The raw hazards are calculated as $\frac{y(t_i)}{P(t_i)}$, where $y(t_i)$ is the number of events in a given time interval and $P(t_i)$ is the total number of networks at risk within that time interval (Rebora, Salim, & Reilly, 2018, p. 116). The number of events is modelled as a Poisson random variable, with the dispersion parameter ϕ . The smoothing parameter is directly estimated from the data as $\lambda = \frac{\phi}{\sigma_\nu^2}$, where σ_ν^2 is the variance of the normally distributed B-splines with ν being the vector of coefficients that determine the amount of inflection at the knots (Rebora, Salim, & Reilly, 2018, p. 116). Accordingly, larger λ means more smoothing. Underfitting is prevented by using a large number of equally spaced knots and overfitting is prevented with a penalty term. The hazard functions were plotted using the **survminer** package (Kassambara, Kosinski, Biecek, & Fabian, 2020).

Conditional survival probabilities were computed using the **dynpred** package (Putter, 2015). The package returns conditional death probabilities:

$$P(T \leq t + w | T > t) \quad (19)$$

for a fixed width w (Putter, 2015, p. 12) from which we derived conditional survival probabilities as $1 - P(T \leq t + w | T > t)$. Specifically, the package uses the survival estimates previously obtained from the **survival** package (Therneau, 2020). These estimates are used to calculate cumulative hazard rates as $\hat{H}(t) = -\log \hat{S}(t)$ for each time window (i.e., the hazard accumulated between t and $t + w$, given survival until t), based on the Nelson-Aalen estimator (Nelson 1972; Aalen, 1978). The conditional survival probabilities can be derived from (22) in Appendix A and (20) becomes the complement.

Distribution fitting. We first used visual aids to get a sense of the distribution. Specifically, we plotted the frequency distribution, the exponential Q-Q plot and the Zipf / log-log plot of the network lifetimes. The exponential Q-Q plot plots the theoretical quantiles of the exponential distribution against the empirical quantiles of the data. A

linear relationship indicates that the distribution is likely exponential. A concave relationship indicates that the sample distribution has heavy tails – but not necessarily fat tails (Cirillo, 2013). The Zipf plot plots the logarithm of the empirical survival function against the logarithm of the sample values. A perfectly linear relationship – rarely observed – points to the Pareto distribution. On the other hand, a concave relationship indicates that the survival function decreases faster than that of a power law, meaning the tail is lighter. We used this plot to see if we can rule out the Pareto distribution, as we did not expect to observe it in psychological variables. The Zipf plot was generated using Cirillo’s (2013) code. This part of the estimation procedure was exploratory, as it informed our choices in other methods. Specifically, the distribution of strongly connected networks seemed heavy tailed but not fat tailed, and the distribution of weakly connected networks had a much thinner tail. Therefore, we do not report these findings in the results section. Instead, the graphs are placed in Appendix G.

As explained in the discussion of the Lindy effect, the Weibull, Lognormal and Gamma distributions can have heavy tails and can demonstrate both weak (lognormal) and strong Lindy (Weibull and Gamma depending on the parameter values). Therefore, we used the **fitdistrplus** package (Delignette-Muller & Dutang, 2015) to obtain maximum likelihood estimations for these three candidate distributions. We also fit an exponential distribution to the weakly connected networks because of the thinner tail and the relatively well fit in the exponential Q-Q plot (see Appendix G). Importantly, the MLE method was adjusted to accommodate right-censoring. A data point is right-censored if it has not failed within a set time period, and therefore there is no information about its status after the set time point. In this case, all we can say about a network is that it has survived for $T \geq t$ time points. In our case, the censoring point and the maximum lifetime is $10000 - 1 = 9999$, since we started counting time at the second time point, after the networks were stressed. The likelihood function then becomes (Delignette-Muller & Dutang, 2015; "8.4.1.2. Maximum likelihood estimation", n.d.):

$$L(\theta) = \left(\prod_{i=1}^r f(t_i|\theta) \right) [1 - F(T|\theta)]^{n-r} = \left(\prod_{i=1}^r f(t_i|\theta) \right) [S(T|\theta)]^{n-r}, \quad (20)$$

where r is the number of failures, n is the total number of lifetimes (therefore $n - r$ is the number of censored lifetimes) and $F(T)$ and $S(t)$ are the CDF and the survival functions of the candidate parametric distribution evaluated at the time of censoring (also the maximum time point in our case). 95% confidence intervals (CI) were constructed by bootstrapping for $N = 1000$ iterations, both around the point estimates and around the CDF. We report the former in text. For the latter, we plotted the theoretical CDF against the empirical CDF with CIs around the theoretical CDF.

We used the Anderson-Darling (AD) test (Anderson & Darling, 1952) as our exact fit measure. The AD test is a test of whether an empirical distribution comes from a specified theoretical distribution. The test statistic is a weighted sum of the squared distances between the data points and a fitted CDF:

$$A_{AD,n}^2 = n \sum_{i=1}^n \frac{[i/n - F_0(X_i)]^2}{F_0(X_i)(1 - F_0(X_i))} \quad (21)$$

The AD test is more sensitive than the Kolmogorov-Smirnov (KS) test and gives more weight to the tails of the distribution. This is good because the tail behaviour is crucial for the Lindy effect. Just like the KS test, the AD test assumes that the distribution is fully specified. So, the test statistic needs to be adjusted if the parameters are not known in advance but are estimated from the sample. Further adjustment is needed in estimating the CDF to accommodate right-censoring. The adjusted AD test statistic was computed using the **parmsurvfit** package (Jacobson, Wilson, & Pileggi, 2018). The package follows the Minitab procedure to compute the AD statistic for MLE (Minitab, 2019). The CDF is estimated using the adjusted median rank (details can be found in "Reliability Engineering Interactive Training"). To compare different models, we used the Akaike Information Criterion (AIC) (Akaike, 1973) and the Bayesian Information Criterion (BIC) (Schwarz, 1978). These are given as:

$$\begin{aligned}
AIC &= 2k - 2\ln(\hat{L}) \\
BIC &= k\ln(n) - 2\ln(\hat{L})
\end{aligned}
\tag{22}$$

Results

The hazard functions

As expected, we observed the Lindy effect in strongly connected networks. After becoming depressed, the conditional probability that a strongly connected network would transition back to a non-depressed state, given that it has remained in the depressed state for a period of time, decreased over time. This was shown in a decreasing hazard function and increasing conditional survival probabilities (Figure 7). The pattern was particularly clear between the first 0-2000 time points. However, some important findings question the validity of these findings. The most important of these findings is that we had an extremely high overdispersion ($\phi = 15.99$) in the estimated hazard rates, meaning that the sample variance was much greater than assumed by the Poisson model. This indicates that we did not use the correct model to describe the observed hazard rates. One reason for this is that the Poisson distribution has the same parameter λ for both the mean and the variance. Consequently, the variance grows with the event rate, which is an issue when there are many events. A negative binomial model may be a better choice, since it can estimate the variance independently of the mean. Second, we display conditional survival and death probabilities that were computed within 1000 point time-intervals, which is quite large considering that the entire period was 10000 points. The reason we chose this value is that for smaller values the variance in each interval grows very large and masks the overall pattern. This is especially the case because the sample size was small ($N = 1000$ networks), and particularly clear in the right tail because the number of failures decreases substantially towards the tail. The trend in the conditional survival probabilities does not change but becomes less smooth than what we show here. That being said, regardless of

the parameters or the window width we chose, the Lindy effect was robust for the first 2000 time points. Overall however, these results mean that our sample of strongly connected networks only roughly demonstrate the Lindy effect over the complete time period, and the variance of the hazard rates is too large to have confidence in with the Poisson model used by the **bshazard** package (Rebora et al., 2018).

In contrast, weakly connected networks showed much weaker indications of the Lindy effect. Specifically, the hazard function and the conditional survival probabilities of the weakly connected networks lacked a clear pattern. Figure 8 shows that the hazards increase for the first 10 time points, after which they have a vaguely discernible tendency to decrease. Unlike in strongly connected networks, there is relatively little overdispersion ($\phi = 1.44$). By 70 time points, pretty much all networks are dead, which is why we estimated the conditional probabilities within 10 point windows. In part due to this small time window, the CIs are so large and the estimates so haphazard that the plots are very difficult to interpret. Like the hazard rates, the conditional survival probabilities do not show a clear pattern; they increase for the first 30 time points, after which the pattern disappears. Overall, we can confidently conclude that weakly connected networks do not show the Lindy effect in any meaningful sense.

Distribution fit

The ML estimates for the Weibull distribution were $k = 0.36$, 95% CI [0.34, 0.38] and $\lambda = 1632.26$, 95% CI [1349.15, 1994.72]. The ML estimates for the lognormal distribution were $\mu = 5.97$, 95% CI [5.74, 6.20] and $\sigma = 3.53$, 95% CI [3.37, 3.69]. Finally, the ML estimates for the gamma distribution were $\alpha = 0.27$, 95% CI [0.25, 0.28] and $\beta \approx 0.00006$, 95% CI [0.00004, 0.00007]. These results match what the hazard function shows.

Specifically, the shape parameters of the Weibull and the Gamma distributions are both smaller than 1 and the standard deviation of the lognormal is rather large (we consider > 2 to be large because then the hazard starts to decrease quickly enough to be confused for a

monotonically decreasing function (Sweet, 1990; Kurniasari, Widayarni, Warsono, & Antonio, 2019)).

Among these three models, the lognormal distribution fits the data the best ($AIC = 13265.53$, $BIC = 13275.34$), followed by the Weibull distribution ($AIC = 13281.05$, $BIC = 13290.86$), and lastly the gamma distribution ($AIC = 13349.54$, $BIC = 13359.36$).

However, a closer comparison of the models using ECDF and Q-Q plots tells a more nuanced story. Panel (a) of Figure 9 plots the theoretical CDF against the sample values and Panel (b) plots the empirical quantiles against the theoretical quantiles. Both plots show that, even though the lognormal distribution matches the body of the distribution more closely, the Weibull distribution describes the tails better. This indicates that the tails are not as heavy as those of a lognormal distribution.

On the other hand, neither the lognormal nor the Weibull distributions fit the data based on the AD test. That is, we rejected the null hypothesis that the data came from the theoretical distribution for both the lognormal ($A^*2 = 526.53, p < .001$) and the Weibull distributions ($A^*2 = 521.42, p < .001$)¹. We could not compute a test statistic for the gamma distribution, but the previous model comparisons make it extremely unlikely that the gamma distribution would fit the data. It is noteworthy that the p-values are difficult to compute for the adjusted AD statistic and the package does not provide them (see the footnote to see how we arrived at the reported p-values). Therefore, the results of the AD test may be more informative for relative fit than absolute fit, especially considering that the CDF plot shows decent fit with high confidence for the body of the distribution. The tails are less adequately described by either distribution – although better by Weibull than lognormal – and as we noted, the AD tests is very sensitive at the tails. More importantly, we showed in our analysis of the Lindy effect that it is more important to identify the class

¹ The p-values could not be computed from the package and are not available for these values of the test statistic. We estimated the range as $p < .001$ because the values are far above the critical values computed for the adjusted AD statistic in D’Agostino and Stephens (1986), which is the standard text on the matter.

to which a distribution belongs in terms of its tail behaviour than to identify the exact distribution. We also have strong evidence that the lifetime distribution of strongly connected networks is heavy tailed, as shown in both Appendix G, the parameter estimates reported above, the close fit of these models in the CDF plot, and the finding that the tail is between that of a heavy tailed Weibull and a heavier tailed lognormal.

As for the lifetime distribution of the weakly connected network, the ML estimates were $k = 1.36$, 95% CI [1.27, 1.47] and $\lambda = 7.84$, 95% CI [7.48, 8.22] for Weibull, $\mu = 1.66$, 95% CI [1.61, 1.71] and $\sigma = 0.87$, 95% CI [0.82, 0.90] for lognormal, $\alpha = 1.76$, 95% CI [1.62, 1.94], $\beta \approx 0.25$, 95% CI [0.22, 0.27] for Gamma, and $\lambda = 0.14$, 95% CI [0.13, 0.15] for the exponential distribution. As was the case with the strongly connected networks, these parameters make sense in light of the hazard function of the weakly connected networks; Weibull and gamma have thinner tails for shape > 1 and the standard deviation of the lognormal is much smaller. Comparing the models with AIC and BIC show that the Weibull (AIC = 5789.90, BIC = 5799.71) and the gamma distributions (AIC = 5774.06, BIC = 5783.87) are close to each other in terms of goodness-of-fit, with the gamma distribution being the better model ($AIC_{Weibull} - AIC_{Gamma} = 5789.90 - 5774.06 = 15.84$ and $BIC_{Weibull} - BIC_{Gamma} = 5799.71 - 5783.87 = 15.84$). In comparison, the lognormal (AIC = 5865.91, BIC = 5875.73) and especially the exponential distribution (AIC = 5933.686, BIC = 5938.59) have very poor fit. On the other hand, the AD test favoured the Weibull model ($A^*2 = 15.40, p < .001$) over the gamma model ($A^*2 = 17.86, p < .001$). The lognormal ($A^*2 = 38.51, p < .001$) and the exponential distributions ($A^*2 = 46.18, p < .001$) fared poorly in comparison. These values indicate poor exact fit for all the models, but less so compared to the lifetime distribution of the strongly connected networks. Due to the small difference between the Weibull and the gamma models, the contradiction between the relative fit indices and the AD test, and the poor fit of the lognormal and the exponential models across the board, we plotted the CDFs and the quantiles only to compare the Weibull model to the gamma model. This let us avoid cluttering in the plots. As with the

strongly connected networks, Panel (a) of Figure 10 plots the theoretical CDF against the sample values and Panel (b) plots the empirical quantiles against the theoretical quantiles. The CDF plot suggests that both models fit the data well but we cannot distinguish between the two. A closer look at the Q-Q plot indicates that both models fit the body of the distribution nicely but neither describes the tails very well. As was the case with the strongly connected networks, this can explain the sensitivity of the AD test.

Discussion

The results suggest that highly connected networks can in principle show the Lindy effect. Namely, the longer a highly connected network has been depressed, the more likely it is to remain depressed in the future. This pattern was reflected in a decreasing hazard function and increasing conditional survival probabilities. This means that after transitioning to a depressed state, only highly connected networks remain in a depressed state. The others go back to a non-depressed state quickly. So, if a network has been depressed for a while, it is likely to remain depressed because it is more likely that it is a highly connected network that got stuck in the depressed state. In other words, weakly connected networks fail early, whereas strongly connected networks remain depressed. The Lindy effect was most salient between the first 0 – 2000 time points. The conditional probability of transitioning back to a non-depressed state fell sharply in that period and levelled off thereafter. The lifetime distribution of strongly connected networks is heavy tailed, with most failing early and a few remaining depressed for a long time. Relative measures favour the lognormal distribution as the best fit, but the right tail is better described by a Weibull distribution. That is, the tail is not as heavy as that of a lognormal. Reducing connectivity cuts the tail – networks with small connectivity levels cannot remain in the depressed state for a long time. The resulting distribution clearly has a short tail and the survival rates do not show a clear pattern. Consequently, the Lindy effect was less apparent in weakly connected networks. Specifically, the conditional probability of

remaining depressed had a very general tendency to increase for the first 30 time points. After that, it decreased. This makes sense, as most networks stopped being depressed in the first few time points, the vast majority had gone out of depression by time point 30, and all of the networks had stopped being depressed by the 70th time point. One might say there is weak Lindy, but the CIs were so large that any interpretation of the pattern is dubious.

Taken at face value, these findings support our hypothesis. However, there are important caveats. Most importantly, we had a very large overdispersion for the highly connected networks. This means that the Poisson model used by the **bshazard** package (Rebora et al., 2018) is not appropriate for our data. We suggested that a negative binomial model would be better, since it can estimate the variance independently of the mean. As mentioned, the lognormal distribution is a good fit for the time distribution of certain behavioural data but shows weak Lindy. However, the initial rise in the hazard function can be quite short-lived if the standard deviation is large (see Figure F1. in Appendix F). Therefore, it is possible that we could not capture a subtle initial rise in the estimated hazard function. So, the hazard rate may be non-monotonic, but the possibility of a true lognormal distribution is still low because the tail behaviour is better described by the Weibull distribution. One could say there is also the fact that the AD test rejected the possibility of a lognormal, but we should note that we did not expect to have good exact fit in the first place. That is, we expected to reject the null hypothesis of no difference between the empirical and the theoretical distributions. We did not pre-register this expectation as a confirmatory hypothesis but the reasons for it are important to discuss. First, our sample size was only $N = 1000$ networks because each additional network took a great amount of computational resources. Although this is not a terrible sample size for an empirical experiment, it is not great to estimate a true distribution. Second, we used a very particular network model that generated a particular distribution. This model was highly sensitive to the changes in the connectivity parameter, so it was unlikely that the generated distribution would match any theoretical distribution. Also, it seems to be

common practice in the literature to compare the AD test statistics to look at relative fit, rather than to test exact fit – at least with censored data. Indeed, the p -values do not seem to be reliable for censored data. Likewise, we hope to have demonstrated that the *class* of distributions with which one is dealing and its tail behaviour are more important than the exact distribution, at least when it comes to the Lindy effect. That is why we used the AIC and BIC, which indicated that the lognormal was overall a better fit to the network lifetimes. But we also showed that at a closer look, the Weibull distribution describes the tail behaviour better, whereas the lognormal is a better approximation for the rest. In any case, we can confidently reach the following conclusions: the lifetime distribution of highly connected networks is a) heavy tailed but not fat tailed, b) not as heavy tailed as the lognormal distribution, and c) not a Gamma distribution. The main point is that future work in the field should study the true distribution of depression duration. After all, that is what determines if depression shows the Lindy effect in reality. But regardless of whether or not this is the case, we hope to have shown that the time distribution of a disorder can be very informative about its behaviour, and it is strange that it has not received much attention before.

Also, there are alternative ways of running the network simulations and analysing the data. For example, one could simulate networks with the same connectivity parameter instead of drawing from a sequence of values. The pattern does not change when we do this, but it becomes much less informative. Suppose all the strongly connected networks had a connectivity parameter of 1.3. Then, very few networks ever get out of the depressed state. What we observe is a sharply declining hazard rate, which then remains constant until the end because it is already almost 0. Instead, when we sample the connectivity value, we are able to observe networks that failed somewhere between the two extremes. This is also why the delta of the connectivity vector from which we sampled the values was 0.00015. This allows sufficiently small changes in the connectivity, so that the effect of those small changes can be observed. As a result, we have a more granular pattern in the

transition times. Second, if a network did not become depressed after being stressed, we counted it as failed. This group of networks constitute the proportion that stopped being depressed right at the first time point. That is, they were considered to have stopped being depressed, despite not having been depressed to begin with. The proportion of all networks that were in this situation was 15.6%. Alternatively, we could have excluded these networks before estimating the hazard functions and the conditional survival probabilities. We tried this, and it did not change the patterns we observed. What happened was that the slope of the curves became less steep at the beginning.

Let us now discuss this study in the context of the broader literature. Our results are compatible with several findings in the current psychological literature. For example, Simon (2000) reports that the proportion of patients with depression fell by 90% in 6 months and remained stable for the rest of the follow up period. More importantly, remission at 3 months was the best predictor of remaining in remission in the long term. After controlling for early remission, baseline severity did not have any effect. However, the duration of the index episode still had an effect. Further, only 3% of the patients were still depressed for the majority of the follow ups. Likewise, Patten et al. (2010) found that most depression patients remitted by week 7 and less than 10% of them were depressed at 52 weeks. These results suggest that: a) most patients remit quickly, b) a small proportion of patients remain depressed for a very long time and c) early remission and the duration of the index episode are more important predictors of long-term prognosis than baseline severity. This is precisely what we found in highly connected networks: most networks stopped being depressed quickly, a small proportion remained depressed for a very long time, and longer past depression was associated with a bigger likelihood of remaining in depression. Therefore, network dynamics can explain the findings of Simon and Patten et al. (2010) because both the shape of the distribution of depression times and the positive relationship between past and future depression can arise as a direct result of the reinforcing interactions between symptoms.

Another study important for this paper is by Penninx et al. (2011), who examined the two-year course of depression and anxiety as part of the Netherlands Study of Depression and Anxiety (NESDA). The researchers studied the time to first remission using survival analysis. The findings indicated that longer baseline symptom duration predicted lower risk of remission, meaning that individuals with longer baseline symptoms were more likely to remain anxious or depressed. The researchers summarized: "... once the disorder has become established, not the factors predicting the onset, but the clinical features of the disorder itself are most important for prognosis". Our simulations not only replicate the pattern in these findings, but also show that the level of connectivity between interacting symptoms can be the feature of the disorder that can determine its prognosis to an important extent.

Moreover, Kelley et al. (2018) found that there are groups of subjects with similar treatment response patterns and change over time, after controlling for baseline severity. The researchers identified five groups: early remission, delayed remission, gradual remission, minimal response and no response. A visual inspection of the response patterns (Kelley et al., 2018, Figure 1, p. 14) reveals similarities to geometric Brownian motion sample paths that show the Lindy effect at the tail. Indeed, those who have not remitted by a given time point seem to move further away from the remission cutoff, which was defined as a sum score of ≤ 7 on the Hamilton Rating Scale for Depression (HRSD). This is no proof of the Lindy effect, but it is an important similarity that merits attention. This is because the network dynamics and the resulting Lindy effect can explain this pattern. Specifically, the difference between these groups of responders may be the level of connectivity in their symptom networks. This level can vary between people, just as any other variable. Individuals whose symptom networks have weak connections may remit early. In contrast, those with strong connections may be more robust to treatment and remain depressed.

On the other hand, not all of the findings Kelley et al. (2018) point to the Lindy effect. For example, an indirect implication of the Lindy effect would be a higher

probability of being in the minimal and no response groups conditional on having chronic depression (current episode ≥ 2 years). This is because if mental disorders are really symptom networks and if the level of symptom connectivity determines chronicity to an important degree, then someone who has chronic depression is likely to have a strongly connected network that locked in their depression. This would also imply that their symptom network is more robust to treatment. However, a quick re-calculation of the reported data show that non-responders and minimal responders are not overrepresented among patients with chronic depression (see Kelley et al., 2018, Table 2 and Table 3; also note that this was a simple hand calculation without taking into account the possible effect of dropouts).

Furthermore, not every psychological disorder is likely to show the Lindy effect. For example, Heyman (2013) reports that the probability of quitting drugs is independent of the duration of past addiction. This makes sense at an intuitive level for a few reasons. First, one cannot quit depression the same way one can quit addiction. Indeed, Heyman (2013) argues that the primary reason most people quit is pressure from loved ones. This is not to suggest that social support is irrelevant to depression, but it cannot force a person to suddenly feel happy. Consequently, in depression the features of the disorder may predict long-term outcomes better than factors the patient can control, whereas in addiction factors under the patient's control may be more important. In other words, disorders such as depression and anxiety can be reinforced by network dynamics to a degree that is harder to restrain compared to substance use disorders.

The present study contributes to the current psychological literature in three ways. First, we test the conditional probability of going out of a depressed state as a function of time. Second, we link this question to network dynamics, thereby offering an explanatory mechanism for many of the findings discussed above. Specifically, we show that if the symptoms are sufficiently connected and there is nothing preventing them from being present for a long time (e.g., one cannot lose weight forever), then the disorder can in

principle show the Lindy effect. This offers an explanation for the findings that most people remit early and a few become chronically depressed, that early response to treatment is the best predictor of long-term prognosis, and that the length of the baseline symptoms and the index episode are related to less chance of remission. Finally, the Lindy effect follows from the implications of the network theory. This is important because despite that this study aims to present a proof-of-principle, if empirical data shows that psychopathology indeed shows the Lindy effect, the network theory is the only theory that can provide a satisfactory explanation. In contrast, if psychopathology is memoryless, or if it ages not in reverse but like humans do, then mental disorders may not be networks at all. This study is the first step in testing such implications.

In this study, we experimented with many parameters. Our goal was to explore under what conditions the Lindy effect would arise, and how it would be affected when certain parameters changed. We discussed those changes and our rationale for sticking to certain choices above. As such, we do not make any confirmatory claims except that the symptom-symptom interactions in highly connected networks will lead to the Lindy effect, which remains tentative due to overdispersion. Also, this study gives a proof of principle, not an empirical proof. The presence of the Lindy effect still needs to be tested empirically, ideally in a solid time series study. Unfortunately, there is not much data fit for this purpose (we used cross sectional data to simulate intra-individual networks). Clinical psychology can benefit from prospective time series studies that analyse the progression of symptoms. Furthermore, it would be highly interesting to study the evolution of not just MDD but anxiety disorders as well. This is because anxiety can also be chronic and anxiety symptoms are likely to have strong positive feedback loops (Heeren & McNally, 2016). Finally, we tried to develop a formal hypothesis test to detect the Lindy effect in data, based on resampling methods. However, trying to obtain a single value to confirm its presence did not work very well. We decided that estimating empirical hazard functions and conditional probabilities was a more informative method to study the Lindy effect.

Other researchers are welcome to try to come up with a practical, robust test to detect the Lindy effect in empirical data.

Figure 1. Hazard functions of different types of processes / distributions. The green line shows a decreasing hazard function, illustrating the Lindy effect. This means that the conditional rate of death / failure right after a time point, given survival until that time point, decreases over time. The Pareto distribution is an example of a distribution that has a decreasing hazard function. The blue line illustrates a constant hazard function. This means that age has no bearing on future survival. The exponential distribution, which is memoryless, has a constant hazard rate. The red line shows an increasing hazard function. This means that the chance of death increases with age. This is the case for thin tailed distributions such as the normal distribution. Lifetimes that follow a thin tailed distribution are heavily constrained.

(a) Standard Brownian motion. Simulated for 1000 time points, starting at $X_0 = 0$.

(b) Arithmetic Brownian motion. Simulated for 1000 time points, starting at $X_0 = 0$.

(c) Geometric Brownian motion. Simulated for 1000 time points, starting at $X_0 = 1$.

Figure 2. Standard Brownian motion (a), Arithmetic Brownian motion (b) and Geometric Brownian motion (c). For each process, 100 sample paths were simulated for 1000 points.

(a) Standard Brownian motion with an absorbing barrier, simulated for 1000 time points. The process started at $X_0 = 0$ and the absorbing barrier was set at an arbitrary $B < 0$.

(b) Arithmetic Brownian motion with negative drift and an absorbing barrier, simulated for 1000 time points. The process started at $X_0 = 0$ and the absorbing barrier was set at an arbitrary $B < 0$.

Figure 3. Standard Brownian motion and arithmetic Brownian motion with an absorbing barrier. A sample path stops when it takes a value smaller than or equal to the absorbing barrier. Notice that most sample paths die quickly, but some survive for a long time. The further away they move from the barrier, the less likely they are to hit it at the next time point. We get hints of the Lindy effect.

(a) The density of the Levy distribution, plotted for different values of the scale parameter.

(b) The hazard function of the Levy distribution, plotted for different values of the scale parameter. The hazard function increases until it reaches a peak, and then decreases asymptotically.

Figure 4. The density and the hazard function of the Levy distribution, plotted for different values of the scale parameter.

(a) The density of the Inverse Gaussian distribution, plotted for different parameter values.

(b) The hazard function of the Inverse Gaussian distribution, plotted for different parameter values. As is the case with the Levy distribution, the hazard function increases until it reaches a peak, and then decreases.

Figure 5. The density and the hazard function of the Inverse Gaussian distribution, plotted for different values of the scale parameter.

Figure 6. Brownian motion generates a distribution of stopping times with associated hazard functions. We observe the Lindy effect as a decreasing hazard function.

(a) The smoothed hazard estimates of strongly connected networks. The points are the raw hazards. The black line is the smoothed hazard estimates. The red dashed lines are the 95% confidence bands around the estimated hazard rates. The hazard rates are decreasing, particularly in the first 2000 time points but the raw hazards have high variance.

(b) The conditional survival and death probabilities of strongly connected networks. The conditional survival probabilities are increasing (green) whereas the conditional death probabilities are decreasing (red). The black lines are the 95% confidence bands. The probabilities were estimated in 1000-point time intervals.

Figure 7. The hazard rates (a) and the conditional survival and death probabilities (b) of strongly connected networks. Both show the Lindy effect.

(a) The smoothed hazard estimates of weakly connected networks. The points are the raw hazards. The black line is the smoothed hazard estimates. The red dashed lines are the 95% confidence bands around the estimated hazard rates. There is a very general decreasing pattern for the first 30 time points but all networks are dead by 70 time points. The CIs are extremely large.

(b) The conditional survival and death probabilities of weakly connected networks. The conditional survival probabilities are roughly increasing (green) for the first 30 time points but lack any clear pattern. Almost all networks are dead by 70 points. The probabilities were estimated in 10-point time intervals.

Figure 8. The hazard rates (a) and the conditional survival and death probabilities (b) of weakly connected networks. Neither show the Lindy effect.

(a) The CDF plot of the strongly connected networks, Weibull, lognormal and gamma distributions. The x-axis shows the sample values. The y-axis is the CDF. The black line is the ECDF, the red line is the theoretical CDF of the Weibull distribution based on the ML estimates, the green line is the theoretical CDF of the lognormal and the blue line is the theoretical CDF of the gamma distribution. The lognormal fits the body of the distribution well but the Weibull distribution described the tails better. This indicates that the empirical distribution is heavy tailed but lighter tailed than the lognormal. The rectangular region consists of the censored values.

(b) The Q-Q plot of the strongly connected networks, Weibull, lognormal and Gamma distributions. The x-axis shows the theoretical quantiles of the fitted distributions and the y-axis shows the empirical quantiles. Again, we see that the body is better described by the lognormal but the tails are better described by the Weibull. However, neither describes the tails perfectly. The rectangular region consists of the censored values.

Figure 9. The CDF (a) and the Q-Q (b) plots of strongly connected networks and the theoretical distributions. The lognormal distribution fits the body of the distribution well, but the tails are better described by the Weibull distribution. The rectangular regions are the censored values.

(a) The CDF plot of the weakly connected networks, Weibull and gamma distributions. The x-axis shows the sample values. The y-axis is the CDF. The black line is the ECDF, the red line is the theoretical CDF of the Weibull distribution based on the ML estimates and the green line is the theoretical CDF of the gamma distribution. The two models are indistinguishable based on the CDF plot. The rectangular region consists of the censored values.

(b) The Q-Q plot of the weakly connected networks, Weibull and gamma distributions. The x-axis shows the theoretical quantiles of the fitted distributions and the y-axis shows the empirical quantiles. We see that Neither fit the tails well but the body of the distribution is well described by both. However, we still cannot distinguish between the two. The rectangular region consists of the censored values.

Figure 10. The CDF (a) and the Q-Q (b) plots of strongly connected networks and the theoretical distributions (only the Weibull and the gamma due to poor relative fit estimates for the lognormal and the exponential distributions. Neither the Weibull nor the gamma distribution seems to fit the tails but the body is described well by both. The two are indistinguishable based on the plots. The rectangular regions are the censored values.

Appendix A

Deriving hazard functions

We first note a few equalities that will be helpful when deriving hazard functions:

$$h(t) = \frac{f(t)}{S(t)} = \frac{-\frac{dS(t)}{dt}}{S(t)} = -\frac{d}{dt} \log S(t). \quad (23)$$

Integrating the hazard function gives the cumulative hazard function $H(t)$, which can be interpreted as the risk of death accumulated until a given time point. A key relationship between the survival function and cumulative hazard follows:

$$\begin{aligned} \int h(t)dt &= -\log S(t) \\ &= -H(t) = \log S(t) \\ &= S(t) = e^{-H(t)} \end{aligned} \quad (24)$$

Appendix B

The Levy distribution

We noted that the asymptotic behaviour of the Levy density is a power law ("Levy distribution", n.d.). That is:

$$\lim_{x \rightarrow \infty} \sqrt{\frac{c}{2\pi}} \frac{e^{-\frac{c}{2(x-\mu)}}}{(x-\mu)^{3/2}} = \sqrt{\frac{c}{2\pi}} \frac{1}{x^{3/2}} \quad (25)$$

We can integrate this to get the survival function:

$$\begin{aligned} S(x) &= \int_x^\infty \sqrt{\frac{c}{2\pi}} \frac{1}{t^{3/2}} dt \\ &= \sqrt{\frac{c}{2\pi}} \int_x^\infty \frac{1}{t^{3/2}} dt \\ &= \sqrt{\frac{c}{2\pi}} \left(-\frac{2}{\sqrt{t}} \Big|_x^\infty \right) \\ &= \sqrt{\frac{c}{2\pi}} \left(\lim_{t \rightarrow \infty} -\frac{2}{\sqrt{t}} + \frac{2}{\sqrt{x}} \right) \\ &= \sqrt{\frac{c}{2\pi}} \left(0 + \frac{2}{\sqrt{x}} \right) \\ &= \sqrt{\frac{c}{2\pi}} \frac{2}{x^{1/2}} \end{aligned} \quad (26)$$

and we get the hazard:

$$\begin{aligned} h(x) &= \frac{f(x)}{S(x)} = \frac{\frac{c}{2\pi} \frac{1}{x^{3/2}}}{\frac{c}{2\pi} \frac{2}{x^{1/2}}} = \frac{\frac{1}{x^{3/2}}}{\frac{2}{x^{1/2}}} \\ &= \frac{x^{1/2}}{2x^{3/2}} = \frac{1}{2} x^{-1} = \frac{1}{2x} \end{aligned} \quad (27)$$

Appendix C

The Pareto distribution

The Pareto distribution is defined by the following survival function, which follows a power law ("Pareto distribution", n.d.):

$$S(x) = \begin{cases} \left(\frac{m}{x}\right)^\alpha, & \text{if } x \geq m \\ 1, & \text{if } x < m \end{cases}$$

where m is the "minimum possible value of X " ("Pareto distribution", n.d.) and α is the shape parameter. We get the PDF by differentiating the CDF:

$$\begin{aligned} F(x) &= 1 - S(x) = 1 - \left(\frac{m}{x}\right)^\alpha \\ f(x) &= \frac{dF(x)}{dx} = -\alpha \left(\frac{m}{x}\right)^{\alpha-1} \left(-\frac{m}{x^2}\right) \\ &= \frac{\alpha m}{x^2} \left(\frac{m}{x}\right)^{\alpha-1} = \frac{\alpha m}{x^2} \frac{x}{m} \left(\frac{m}{x}\right)^\alpha \\ &= \frac{\alpha}{x} \left(\frac{m}{x}\right)^\alpha = \frac{\alpha m^\alpha}{x x^\alpha} \\ &= \frac{\alpha m^\alpha}{x^{\alpha+1}} \end{aligned} \tag{28}$$

We derive the hazard with the usual formula:

$$\begin{aligned} h(x) &= \frac{f(x)}{S(x)} = \frac{\alpha m^\alpha}{x^{\alpha+1}} \frac{x^\alpha}{m^\alpha} \\ &= \frac{\alpha x^\alpha}{x^{\alpha+1}} = \frac{\alpha x^\alpha}{x^\alpha x} = \frac{\alpha}{x}, \end{aligned} \tag{29}$$

which decreases as $x(\text{time})$ grows.

Appendix D

The Weibull distribution

The Weibull distribution is defined by the following cumulative hazard function (Hogg, Tanis, & Zimmerman, 2013, p. 115).

$$H(x) = \left(\frac{x}{\lambda}\right)^k, \quad (30)$$

where λ is the scale parameter and k is the shape parameter. We show that the density of Weibull can indeed be derived from this cumulative hazard and it will become clear during the process that the corresponding hazard function is monotonically decreasing for $k < 1$.

$$h(x) = H'(x) = k \left(\frac{x}{\lambda}\right)^{k-1} \frac{1}{\lambda} \text{ and recalling (17),} \quad (31)$$

$$S(x) = e^{-(x/\lambda)^k} \quad (32)$$

Since $h(x) = f(x)/S(x)$, $f(x) = h(x)S(x)$. So, we get

$$f(x) = h(x)S(x) = \frac{k}{\lambda} \left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k}, \quad 0 \leq x < \infty \quad (33)$$

which is the density of the Weibull distribution ("Weibull distribution", n.d.). Notice that the hazard function (21) decreases monotonically for $k < 1$, giving strong Lindy. Figure C1 shows the density and the hazard function of the Weibull distribution for different parameter values.

(a) The density of the Weibull distribution, plotted for different shape values.

(b) The hazard function of the Weibull distribution, plotted for different shape values.

Figure D1. The density and the hazard function of the Weibull distribution, plotted for different values of the shape parameter.

Appendix E

The Gamma distribution

The PDF of the Gamma distribution is:

$$f(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}, \quad 0 < x < \infty \quad (34)$$

where

$$\Gamma(\alpha) = \int_0^\infty x^{\alpha-1} e^{-x} dx \quad (35)$$

is the Gamma function. We integrate the PDF to get the survival function:

$$\begin{aligned} S(x) &= \int_x^\infty \frac{\beta^\alpha}{\Gamma(\alpha)} t^{\alpha-1} e^{-\beta t} dt \\ &= \frac{\beta^\alpha}{\Gamma(\alpha)} \int_x^\infty t^{\alpha-1} e^{-\beta t} dt \end{aligned} \quad (36)$$

Now, we change variables to bring this function back to the form of the Gamma function.

Let $y = \beta t$, so $t = \frac{y}{\beta}$, $dy = \beta dt$ and $dt = \frac{1}{\beta} dy$. Then, using the new variables and changing the limits of integration:

$$\begin{aligned} S(x) &= \frac{\beta^\alpha}{\Gamma(\alpha)} \int_{\beta x}^\infty \left(\frac{y}{\beta} \right)^{\alpha-1} e^{-y} \frac{1}{\beta} dy \\ &= \frac{\beta^\alpha}{\Gamma(\alpha)} \int_{\beta x}^\infty y^{\alpha-1} \frac{1}{\beta^{\alpha-1}} \frac{1}{\beta} e^{-y} dy \\ &= \frac{\beta^\alpha}{\Gamma(\alpha)} \frac{1}{\beta^\alpha} \int_{\beta x}^\infty y^{\alpha-1} e^{-y} dy \\ &= \frac{\Gamma(\alpha, \beta x)}{\Gamma(\alpha)}, \end{aligned} \quad (37)$$

where $\Gamma(\alpha, \beta x)$ is the upper incomplete Gamma function. We get the hazard from the now familiar formula:

$$\begin{aligned}
 h(x) &= \frac{f(x)}{S(x)} = \frac{\frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)}}{\frac{\Gamma(\alpha, \beta x)}{\Gamma(\alpha)}} \\
 &= \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)} \frac{\Gamma(\alpha)}{\Gamma(\alpha, \beta x)} \\
 &= \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha, \beta x)},
 \end{aligned} \tag{38}$$

which gives strong Lindy for $\alpha < 1$. Figure D1 shows the density and the hazard functions of the Gamma distribution.

(a) The density of the Gamma distribution, plotted for different shape values.

(b) The hazard function of the Gamma distribution, plotted for different shape values.

Figure E1. The density and the hazard function of the Gamma distribution, plotted for different values of the shape parameter.

Appendix F

The Lognormal distribution

The PDF of the lognormal distribution is:

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{\left(-\frac{(\log x - \mu)^2}{2\sigma^2}\right)}, \quad 0 < x < \infty \quad (39)$$

where μ is the mean and σ is the standard deviation. The CDF is:

$$F(x) = \Phi\left(\frac{\log x - \mu}{\sigma}\right) \quad (40)$$

Since $1 - \Phi(x) = \Phi(-x)$,

$$h(x) = \frac{f(x)}{S(x)} = \frac{f(x)}{1 - F(x)} = \frac{e^{\left(-\frac{(\log x - \mu)^2}{2\sigma^2}\right)}}{x\sigma\sqrt{2\pi}\Phi\left(\frac{\mu - \log x}{\sigma}\right)} \quad (41)$$

As as Figure E1 shows, this function has the $\cap$ shape (Sweet, 1990) and approaches zero asymptotically:

$$\lim_{x \rightarrow \infty} \frac{e^{\left(-\frac{(\log x - \mu)^2}{2\sigma^2}\right)}}{x\sigma\sqrt{2\pi}\Phi\left(\frac{\mu - \log x}{\sigma}\right)} = 0 \quad (42)$$

For a more detailed analysis of the lognormal hazard, interested readers should refer to Sweet (1990), who derives an equation for the location of the maximum – this derivation is outside the scope of this paper.

(a) The density of the lognormal distribution, plotted for varying SDs.

(b) The hazard function of the lognormal distribution, plotted for varying SDs.

Figure F1. The density and the hazard function of the lognormal distribution, plotted for varying SDs.

Appendix G

Exploring the distribution of network lifetimes

(a) The frequency distribution of the lifetimes for strongly connected networks.

(b) The frequency distribution of the lifetimes for weakly connected networks.

Figure G1. The histograms of network lifetimes. The lifetime distribution of strongly connected networks has a longer tail than that of the weakly connected networks. This means that extreme lifetimes are more likely to be observed in strongly connected networks.

(a) The exponential Q-Q plot of the strongly connected networks. The x-axis shows the empirical quantiles of the network lifetimes. The y-axis shows the theoretical quantiles of the exponential distribution. The black lines show 95% confidence intervals. Concavity indicates heavy tails, but not necessarily a power-law.

(b) The exponential Q-Q plot of the weakly connected networks. The x-axis shows the empirical quantiles of the network lifetimes. The y-axis shows the theoretical quantiles of the exponential distribution. The sample quantiles are closer to the theoretical quantiles for the most part, indicating lighter tails. However, there is some evidence of heavy tailedness for larger values.

Figure G2. The exponential Q-Q plots of network lifetimes. The plot is visibly concave for strongly connected networks, meaning their lifetime distribution is heavy tailed. The quantiles of weakly connected networks are closer to the quantiles of the exponential distribution, indicating they have a lighter tail. Although, there is some evidence of heavy tailedness for larger values.

(a) The Zipf plot of the strongly connected networks, including the censored values. The x-axis shows the log of the network lifetimes. The y-axis shows the log of the empirical survival function. Linearity indicates Paretianity, whereas concavity indicates that the survival function decreases faster than that of a Pareto distribution. This means that the empirical distribution is not as heavy tailed as the Pareto distribution. The relationship is concave, with a sharp decline near 10000 time points. This is natural, since the maximum lifetime can be 9999.

(b) The Zipf plot of the strongly connected networks, excluding the censored values. We show the same plot without the censored values to remove the drop-off at the end and clarify the concavity.

Figure G3. The Zipf plot of the strongly connected networks, with and without the censored lifetimes. The x-axis shows the log of the network lifetimes. The y-axis shows the log of the empirical survival function. Linearity points to the Pareto distribution, whereas concavity indicates that the tails of the empirical distribution are not as heavy as the tails of the Pareto distribution. Both plots show a concave relationship. Removing the censored lifetimes in (b) makes the concavity more evident.