The uncommon is common: Structural similarities of symptom heterogeneity across mental disorders

Tobias R. Spiller^{1,2,3*}†, Or Duek^{1,2}†, Markus Helmer¹, John Murray¹, Roland von Känel³, Ilan Harpaz-Rotem^{1,2}

¹Department of Psychiatry, Yale University School of Medicine; 333 Cedar St, New Haven, CT 06510, USA.

²National Center for PTSD, Clinical Neurosciences Division, VA Connecticut Healthcare System; 950 Campbell Avenue, West Haven, CT 06516, USA.

³Department of Consultation-Liaison Psychiatry and Psychosomatic Medicine, University Hospital Zurich (USZ), University of Zurich (UZH); Culmannstrasse 8, CH-8091, Zurich, Switzerland.

*Corresponding author. Email: tobias.spiller@yale.edu

†These authors contributed equally to this work

Abstract

The heterogeneity of symptoms among individuals diagnosed with the same mental disorder complicates the identification of biomarkers and the development of targeted treatments. Yet, the characteristics of this heterogeneity remain largely unknown. We investigated the frequency of disorder-specific symptom combinations, a marker of symptom heterogeneity, in five samples, each assessed for symptoms of a specific disorder (posttraumatic stress disorder, depression, anxiety, schizophrenia and burnout). The frequency of symptom combinations was heavily skewed in all samples, with most symptom combinations being reported only by few individuals. Moreover, the distribution of the symptom combination frequency could be approximated by a power-law and a log-normal distribution. This demonstrates similarities in the structure of symptom heterogeneity among mental disorders. Furthermore, we show that studies with sample sizes typical for research in mental health preclude many rare symptom combinations, limiting the validity of the obtained evidence by such studies for individuals with rare symptom combinations.

Background

Mental health problems can manifest in a wide range of symptoms. The systematic study of these symptoms and their co-occurrence, also known as psychopathology, has underpinned efforts to standardize the diagnosis and categorization of mental disorders. In the most commonly used classification systems for mental disorders — the Fifth edition of the Diagnostic and Statistical Manual of Mental Disorders (DSM-5; American Psychiatric Association, 2013) and the Tenth edition of the International Classification of Diseases (ICD-10; World Health Organization, 1992) — a specific set of symptoms defines a psychopathological syndrome that is labeled as a disorder. Most disorders' definitions are polythetic, meaning that many different combinations of symptoms (symptom phenotypes) are sufficient to fulfill a disorder's criteria. Consequently, individuals diagnosed with the same mental disorder can and do display vastly different symptoms (Nandi et al., 2009; Olbert et al., 2014). For example, the DSM-5 definition of posttraumatic stress disorder (PTSD) can theoretically be met in 636,120 different ways (Galatzer-Levy & Bryant, 2013). Further complicating the symptom heterogeneity of mental disorders, individual symptoms of the same mental disorder respond differently to treatment (Eeden et al., 2019), differ in their genetic correlations (Thorp et al., 2020), and differently relate to other disorders (Lux & Kendler, 2010). Researchers have blamed this within-disorder heterogeneity for the ongoing illusiveness of biomarkers (Kapur et al., 2012) and the lack of neuroscientifically-informed therapeutic agents for mental disorders (Cuthbert & Insel, 2013).

Researchers have used several strategies to reduce within-disorder heterogeneity, including the addition of a so-called specifier to the diagnostic criteria which allows to define a more homogenous diagnostic subgroup within a diagnostic category (e.g., melancholic depression in the DSM-5) (American Psychiatric Association, 2013), or the division of one disorder into distinct diagnostic categories (e.g., complex PTSD and PTSD in the upcoming ICD-11) (Brewin et al., 2017). Data-driven approaches have also been used to identify homogenous subgroups within disorders (e.g., latent class analysis) (Klonsky & Olino, 2008), and to identify common factors in symptoms of a specific mental disorder (e.g., exploratory factor analysis) (Asmundson et al., 2000). However, these strategies have neither driven systematic study of the empirical characteristics of phenotype heterogeneity, nor yielded consistent results (Rasmussen et al., 2019; van Loo et al., 2012). The U.S. National Institute of Mental Health responded to this intractable problem by instituting a funding policy change in 2009 that replaced symptom-based diagnostic categories with the Research Domain Criteria (RDoC; *11*). Despite RDoCs influence on mental health research, symptom-based definitions of mental disorders still drive clinical research, FDA approval of new therapeutic drugs, as well as clinical practice and are thus affected by symptom phenotype heterogeneity.

Although the vast number of theoretical possible symptoms combinations of mental disorders due to their polythetic definition in the DSM-5 has emphasized (Galatzer-Levy & Bryant, 2013; Young et al., 2014), symptom heterogeneity has not been investigated systematically in real world samples (Nunes et al., 2020). In this present study, we aimed to systematically evaluate the heterogeneous presentation of symptoms across several mental disorders using five large samples. Participants' symptom ratings, which were originally obtained on Likert scales, were binarized to resemble the requirements of the DSM-5 for the analysis of four samples. In a first complimentary analysis, the same analysis routine was conducted in a subset of the fifth sample containing only participants with the same overall symptom burden. In a second complimentary analysis, symptoms were not binarized to assess the effect of binarization. In all analyses except for the second complimentary analysis, symptom heterogeneity was assessed with two measures: a) phenotype similarity among individuals and b) the distribution of phenotype frequency within each sample. In a preliminary analysis, the plots of relative phenotype frequencies showed a heavily skewed, heavy tailed distribution (Foss et al., 2013), which appeared linear in log-log coordinates – a hallmark feature of a power-law distribution (Mitzenmacher, 2004). We then used methods introduced by Clauset et al. to test whether the observed distributions obey a power-law or a log-normal distribution (Clauset et al., 2009).

Methods

Procedure

All analyses conducted in this study were secondary analyses of five samples gathered for different purposes. These datasets were obtained from different sources. One sample was obtained from the National Institute of Mental Health Data Archive (NDA). An additional two samples were created using electronic health care record data of the U.S. Department of Veteran Affairs. One sample is accessible from an interactive website aimed to educate the public about personality and one sample was obtained from a privately run web-based health application. The details of the samples are outlined below. In total, we included disorder-specific symptom assessments of 286,645 individuals.

Sample 1: NDA - PANSS

Study design & Setting. We reanalyzed data extracted from the National Institute of Mental Health Archives Data Archive (NDA). We used only “PANSS” as a search query. No additional search criteria were added to maximize the potential sample size. Information about the specific studies that have generated this dataset can be found in the NDA website (doi: 10.15154/1522637) The data of all these studies were stacked within the NDA platform into a combined dataset containing 8532 individual participants.

Measures. Positive and negative symptoms were assessed using the PNASS (Positive and Negative Syndrome Scale). The PNASS is rated on a Likert scale from 1 (*absent*) to 7 (*extreme*). Our analysis was restricted to the two subscales assessing positive and negative symptoms, each consisting of 7 items. We also created a total score by summing the ratings of both subscales which ranged from 14 to 98). Cronbach’s alpha for the total severity score was 0.83 in this sample.

Participants. The compiled original dataset included 23,864 rows of data, after excluding duplicates (by using GUID) we were left with 4779 subjects. Of these, 163 (3.4%) had missing data, yielding a total of 4,616 participants for the analysis. Median age of the sample was 39 (Q1, Q3 = 29, 55) with the majority of participants being male (2963, 64.2%). The median PNASS total score (of positive and negative items) was 33 (Q1, Q3 = 26, 40).

Sample 2: VA – PCL-5

Study design. The data stem from the electronic health care records of the Department of Veteran Affairs (VA) and were initially obtained for clinical purposes.

Setting. We extracted data of all veterans assessed for a possible diagnosis of posttraumatic stress disorder (PTSD) in the VA health care system between October 1, 2018, and September 30, 2019. The data extraction and use for this study were approved by the Institutional Review Board of the VA.

Measures. Symptoms of PTSD were assessed with the PTSD Checklist-5 (PCL-5). The PCL-5 contains 20 items each assessing the severity of a given symptom in the past month. Individual items are rated on a 5-point Likert scale ranging from 0 (*not at all*) to 4 (*extremely*) (Blevins et al., 2015). A total severity score, ranging from 0 to 80, was obtained. Cronbach’s alpha for the total severity score was 0.93 in this sample.

Participants. We obtained data of 165,397 veterans. 113,137 (68.4%) of the assessed veterans had at least one missing PCL-5 item and were therefore excluded from further analysis. Hence, a total of 52,260 participants were included in the final analysis. Median age was 45 (Q1, Q3 = 35, 57) years and the majority of the sample were males (n = 45,193, 86.5 %). The PCL-5 sum score ranged from 0 to 80 points with a median score of 59 points (Q1, Q3 = 50, 67).

Sample 3: Internet - DASS

Study design. The questionnaire was administered online using an interactive website aimed to educate the public about personality tests and to promote open-source assessments and open data.

Setting. The users of the interactive website were informed that their answers would be used for research. At the end of the assessments, the users were asked to confirm whether or not their answers were truthful and appropriate for research. Data of those who did not confirm were not included in the final dataset. Only complete data were provided.

Measures. Respondents were asked to indicate their age and gender (female, male, or other). The Depression Anxiety Stress Scales (DASS) consist of 42 items grouped into three scales (Depression, Anxiety, and Stress). Each item assesses the magnitude of an experienced negative emotional symptom over the past week on a 4-point Likert scale ranging from 1 (*did not apply to me at all*) to 4 (*applied to me very much, or most of the time*) (Lovibond & Lovibond, 1995). For this study, we only considered response to the 13 items of the anxiety subscale. For this subscale, a sum score was calculated ranging from 13 to 52. Cronbach's alpha of the sum score of the anxiety scale was 0.91, demonstrating excellent consistency.

Participants. The obtained dataset consisted of $n = 39,775$ fully completed assessments. After the exclusion of assessments indicating an age older than 95 years ($n = 8$, $< 0.1\%$) and assessments without information on gender ($n = 67$, 0.2%) a total of 39,700 assessments were included in the final analysis. Of these, $n = 30,361$ (76.5%) were provided by women, $n = 8788$ (22.1%) by men and in $n = 551$ (1.4 %) cases, gender was marked as other. Median age was 21 (Q1, Q3 = 18, 25) years. The summed DASS score ranged from 42 to 168 points with a median of 100 points (Q1, Q3 = 77, 123).

Sample 4: Health Application - MBI

Study design. Data were gathered consecutively from respondents to a web-based health application assessing the risk of job burnout.

Setting. The freely accessible application was hosted in Switzerland and aimed to provide a self-assessment of burnout risk to respondents of the survey (for details concerning the application see; von Känel et al., 2016). The web-based application was advertised with a camping on a Swiss news portal for the health care sector (Medinside.ch) in November 2016. The only exclusion criterion for the use of the application was age below 18 years. For this analysis, we additionally excluded respondents older than 70 years, the oldest age of official retirement in Switzerland, and all but the first responses by those participants who completed the questionnaire more than once. Furthermore, the analysis was restricted to data collected between November 2016 and September 2019. Ethical authorization for this study was not required under the Swiss Federal Regulations on Human Research (Swiss Human Research Act). Nonetheless, all data were collected in accordance with the Swiss national ethical guidelines and in compliance with scientific standards specified by Switzerland's Federal regulations. Informed consent to the use of anonymized data for scientific purposes were provided by all respondents.

Measures. Demographic data on age and gender (female or male) were obtained from all participants. Burnout was assessed using the 16-item Maslach Burnout Inventory- General Survey (MBI-GS; Maslach et al., 1996) with each item rated on a 7-point Likert scale ranging from 0 (*never*) to 6 (*daily*). The MBI-GS consists of three dimensions namely *emotional exhaustion* (five items) *depersonalization* (five items) and *professional efficacy* (six items). Overall burnout severity was assessed with a sum score of all 16 items ranging from 0 to 112. The scale demonstrated acceptable internal consistency with Cronbach's alpha of 0.75, 0.91, 0.85, 0.68 for the MBI-GS total score and the three dimensions *emotional exhaustion*, *depersonalization* and *professional efficacy*, respectively.

Participants. A total of 32,258 assessments were completed in the study period. Of these, 179 repeated entries, 124 participants older than 70 years, and 717 participants with missing data regarding their age and 6,915 participants with missing data regarding the MBI-GS were excluded, resulting in a final sample size of 24,323 individuals. The majority of the participants ($n = 17,070$; 70.2%) were female. Median age was 38 (Q1, Q3 = 29, 48) years. Burnout severity ranged from 0 to 91 points with a median of 42 points (Q1, Q3 = 27, 57).

Sample 5: VA - PHQ-9

Study design. Data were initially gathered for clinical purposes at the Department of VA and were extracted out of the electronic health care records.

Setting. We extracted demographic data and all available Patient Health Questionnaire-9 assessments from veterans who consulted with the VA between January 1, 2019, and December 31, 2019. The Institutional Review Board of the VA approved the data extraction and the use of the data for this study.

Measures. The Patient Health Questionnaire (PHQ-9) assesses nine depressive symptoms according to the Fifth edition of the Diagnostic and Statistical Manual of Mental Disorders (DSM-5) in the past fourteen days. Participants rated each symptom on a 4-point Likert scale ranging from 0 (*not at all*) to 3 (*nearly every day*) (Kroenke et al., 2001). All items were added up to a sum score ranging from 0 to 27. The sum scores Cronbach's alpha was 0.88 in this sample.

Participants. A total of 166,469 assessments were extracted. Of these, 1072 (0.6 %) were excluded due to missing PHQ-9 items. Thus, a total of 165,397 participants were included in the final analysis. The median age of the sample was 49 (Q1, Q3 = 37, 63) years with a minority of females ($n = 29,702$, 18.0%). The median overall PHQ-9 score was 14 (Q1, Q3 = 8, 19) with a range from 0 to 36.

Statistical analysis

The basic analytic procedure consisted of a main analysis and two complementary analyses. The main analysis included four steps: (a) binarization (i.e., presence or absence) of individual items, (b) calculation of the unique phenotypes and their characteristics, (c) assessment of the frequency distribution, and (d) assessment of symptom coherence. The first complementary analysis investigated the distribution of symptom phenotypes in subgroups of patients with each subgroup consisting of individuals with identical overall symptom burden. The procedure included an adapted version of all four steps of the main analysis. In the second complementary analysis, we investigated the distribution of symptoms of non-binarized symptom phenotypes based on steps (b) and (c) of the main analytic procedure. The main analysis was carried out using data of studies 1 to 4, the complementary analyses were computed using a subset of data from sample 5. Only complete records without missing items were included in the analysis. All analyses were carried out using *R* version 4.0.3 and 4.1.0 Important packages used were among others *tidyverse* (Wickham et al., 2019) and *powerLaw* (Gillespie, 2015). The analytic code and detailed information about the used packages are available on GitHub: https://github.com/orduek/Symptoms_Heterogeneity_Psych.

Main analyses

Binarization. The classification systems often require symptoms to be present (over a minimal time period or at least a certain number of times during a time period) or not (respectively, fewer or briefer than required). Thus, for four of the five datasets, symptoms are assessed as a binary state, namely as being either present or absent. Moreover, designating items as present or absent facilitates the characterization of symptom phenotypes, which are defined as the individual combinations of present and absent symptoms. This procedure limits the number of possible phenotypes of a given disorder to 2 to the power of the number of items (2^X), whereas the use of the original Likert scale would lead to a much higher number of possible phenotypes, e.g., a disorder with 15 symptoms leads to a staggering $5^{15} = 30,517,578,125$ versus $2^{15} = 32,768$ individual phenotypes. Therefore, we binarized the ratings in four datasets based on cut-offs outlined in the instructions of the questionnaires or on the basis of commonly used conventions. In sample 1 (PANSS), a cut-off of 3 reflecting mild psychopathology was chosen. Thus, items ratings of 3 or below were considered to represent the absence of a symptom, and ratings of 4 or higher (i.e., 4, 5 and 6), indicating at least moderate psychopathology, were considered to represent the presence of a symptom (Kay et al., 1987). In sample 2 (PCL-5), items rated 1 or lower were considered as an indicator of symptom absence and ratings of 2 or higher (i.e., 2, 3 and 4) as an

indicator of symptom presence. This cut-off was suggested by the developers of the PCL-5, outlined in their validation study (Blevins et al., 2015, p.). In sample 3 (DASS), item ratings of 2 or lower were taken as an indicator of symptom absence and ratings higher than 2 (i.e., 3 and 4) were considered to represent symptom presence. This was based on the median of the Likert-scale. In sample 4 (MBI), the cut-off was made at 3, with item ratings 3 or lower resulting in a coding of the symptom as absent and with a rating higher than 3 (i.e., 4, 5 and 6) being coded as symptom presence. In sample 5 (PHQ-9), when using a binarized phenotype, the cut-off of 1 was based on the median of the rating scale. Hence, item ratings 1 or lower and higher than 1 (i.e., 2 and 3) were binarized to symptom absence and presence, respectively. In an additional analysis, the symptoms were not binarized, but left with their original four level (zero to three) coding (see Additional analyses).

Phenotype characteristics. Phenotypes were defined as the individual combination of symptom presence and absence of all of the assessed symptoms in a single sample. Thus, two individuals were considered to have the same symptom phenotype, when their individual combination of binarized symptoms was exactly the same. In one additional analysis, a phenotype was defined as the combination of individual ratings of each symptom. Thus, the phenotypes of two individuals were only considered equivalent in case both participants rated every assessed symptom on their original scale exactly the same. For each analysis, the frequency of each phenotype was calculated. Subsequently, several characteristics of the phenotypes were assessed. These included the times the most abundant phenotype was reported and the details of the combination of symptoms of the most abundant phenotypes.

Distribution of frequencies of phenotypes. We aimed to formally describe the distribution of the frequencies of phenotypes in the individual samples. As we suspected that the data obey a power-law distribution, we used the methods for the investigation of power-law distributions in empirical data outlined by Clauset et. al (2009) (Clauset et al., 2009). Quoting Clauset and colleagues (2009), “a quantity x obeys a power law if it is drawn from a probability distribution

$$p(x) \propto x^{-\alpha},$$

where α is a constant parameter of the distribution known as the exponent or scaling parameter...[which]...typically lies in the range $2 < \alpha < 3$ ”. However, empirical data often do not obey a power law for all values of x , but only for values greater than some minimum value of x , termed $x_{\min}$. In principle, the analysis of empirical data suspected to obey a power law follows three steps: First, the parameters $x_{\min}$ and α are estimated, using a combination of a maximum-likelihood fitting methods (for technical details see; Clauset et al., 2009). Second, the goodness-of-fit between the estimated power-law and the empirical data are tested, yielding a p -value. A p -value lower than the chosen *alpha* level (i.e., 0.05 in this case) refutes the null hypothesis stating that the empirical data do fit a power-law distribution with the estimated parameters $x_{\min}$ and α . Third, the fit of an alternative distribution, namely the log-normal distribution, is tested and compared to the fit of the power-law distribution using a likelihood ratio test, again implemented to yield a p -value. In this implementation, a p -value below 0.05 is indicative for a better fit of one of the competing distributions.

Symptom coherence. We used the Jaccard index to quantify symptom coherence. The Jaccard index is a prominent measure to assess the similarity between two sets of objects (Fletcher & Islam, 2018; Jaccard, 1901). With regard to this study, the Jaccard index is defined as follows: Given two symptom phenotypes p_i and p_j , the Jaccard index J_{ij} is defined as the number of symptoms occurring in both phenotypes p_i and p_j divided by the number of symptoms that occur in either of the two phenotypes p_i and p_j .

$$J_{ij} = \frac{\text{number of symptoms occurring in both } p_i \text{ and } p_j}{\text{number of symptoms occurring in either } p_i \text{ and } p_j}$$

In general, the Jaccard index is defined as the ratio of the size of the intersection of two sets (e.g., A, B) and the union of these two sets. Thus, it ranges from 0 to 1, with 1 indicating absolute coherence of the two sets (Fletcher & Islam, 2018; Olbert et al., 2014).

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}$$

To obtain the J_{ij} in our study, we used the Jaccard distance D_{ij} , which is equivalent to $1 - J_{ij}$. The Jaccard distance D_{ij} or the dissimilarity of all pairs of phenotypes within a sample was calculated using the *dist()* function of the R *Base* package. To obtain the Jaccard index, the resulting values were subtracted from 1. The symptom coherence across individuals within a sample was then assessed by calculating the median J_{ij} and the interquartile range of all pairs of individuals within a sample. As the PHQ-9 data was extremely large, which ended up with a 10.1gb size vector, in order to calculate the summary statistics of the Jaccard index in this sample we have sampled 100,000 samples (with return) and summarized it. To test for stability, we repeated the procedure 1000 times. All Q1, Q3 and median results ended up the same in all samples.

Complementary analyses

In our first additional, we investigated the effect of an expanded phenotype definition. Therefore, we did not binarize the symptoms. Thus, the number of possible phenotypes increases rapidly, depending on the number of levels of the symptom assessment with the respective Likert scale. In the case of our analysis, we investigated symptoms of depression assessed with the PHQ-9, resulting in $4^9 = 262,144$ possible phenotypes. In contrast, binarizing the symptoms would result in $2^9 = 512$ possible phenotypes underscoring the effect of increasing the base when keeping the exponent constant. Due to the non-binary nature of the symptoms, calculation of the Jaccard index was not possible. The remaining analytic procedure was the same as for the other four studies.

Our second additional analysis assessed the phenotypes and their distribution in three subsets of participants in sample 5 (PHQ-9) with the same overall sum score. Therefore, we subset three groups of participants from the sample investigated in sample 5 for a first analysis. All participants in each subgroup had the same sum score, with the three groups having different sum scores. The cut-offs for the subgroups were chosen to reflect the Q1, median, and Q3 of the sum score, such that the three subgroups represent groups of participants with low, medium, and high sum scores indicative of different severity of depressive symptomatology. After the selection of the subsets, the analysis followed the same steps (phenotype characteristics including binarization, distribution of frequencies of phenotypes and Jaccard index) as outlined above.

Results

Main Analysis

The numerical results are summarized in Table 1 (sample 1-4). Each of the four investigated samples contained more than a thousand unique phenotypes. The median frequency of phenotypes ranged from 16 to 2453. The distribution of phenotype frequency was heavily skewed across all four samples, with less than 50 phenotypes being reported relatively often (see Figure 1, Panel A). The parameters of the best fitting power-law models that approximated the empirical phenotype frequency distribution are shown in Table 1 and illustrated in Figure 1, Panel B (sample 1-4). Across all four samples, $x_{\min}$ ranged from 1 to 16 and α ranged from 1.86 to 2.43. The goodness-of-fit comparison between the log-normal and the power-law model revealed a similar fit for all four samples. Phenotype similarity varied substantially across the four samples, ranging from very high (Jaccard index 0.92) in veterans assessed

for PTSD (sample 2), to moderate (Jaccard index 0.61) among individuals assessed for symptoms of burnout (sample 4).

Complementary Analysis

The results of the first complementary analysis are depicted in Table 1 and Figure 1 (in both as sample 5). The median phenotype frequency was 12, and phenotype frequency distribution was heavily skewed. Because symptoms were not binarized, symptom similarity could not be determined with the Jaccard-Index.

Table 2 outlines the results of the second complementary analysis. Across the three investigated subsample, the median phenotype frequency ranged from 5 to 7 and the phenotype frequency distribution was heavily skewed. In the medium symptom burden subgroup, the two-sided but not the one-sided goodness-of-fit comparison rejected the power-law model. Across the three samples, $x_{\min}$ ranged from 14 to 105 and α ranged from 1.77 to 2.62. Symptom similarity varied substantially, ranging from very low (Jaccard index 0.0, low symptom sample), to high (Jaccard index 0.75, high symptom sample).

Discussion

Our findings demonstrate that the frequency distribution of phenotypes is heavy tailed across multiple mental disorders. In other words, most phenotypes occur rarely and only a few phenotypes occur frequently. The findings further support previous investigations and corroborate the common conception that individuals assessed for the same mental disorder can display very different symptom profiles (Fried & Nesse, 2015; Hickling et al., 2013; Olbert et al., 2014). Moreover, this corresponds well to an often-echoed observation by mental health clinicians that patients presenting with textbook-like symptoms are relatively rare, whereas many patients display highly individual combinations of symptoms.

More formally, the observed heavy tailed distributions can be approximated by log-normal and power-law distributions. Both distributions occur in diverse settings and have been used to model both natural and human-made phenomena such as the abundance of species in animal communities (Ulrich et al., 2010), the distribution of wealth, and the distribution of vertices in networks such as the internet (Barabási & Albert, 1999; Mitzenmacher, 2004; Newman, 2005). Thus, both, power-law and log normal distributions are plausible approximations. We were not able to differentiate whether the data obeyed a power-law or a log normal distribution, which is often impossible in empirical contexts (Clauset et al., 2009).

The characteristics of the observed heavy tailed distributions of psychiatric symptom phenotypes have wide-ranging implications for research in mental health and clinical practice. First, the heavy skew of the observed distribution is a property inherent to this distribution and cannot be overcome by studying larger and larger samples. Thus, as long as we continue to refer to the current symptom-based system for the classification of mental disorders, we will be faced with symptom heterogeneity among individuals with the same disorder. This directly compromises the possibility of identifying disorder-specific biomarkers and may help explain why such attempts have largely failed so far (Abi-Dargham & Horga, 2016). For example, given that symptoms of the same disorder are associated with different genetic markers (Thorp et al., 2020), studies seeking to identify the genetic correlates of a disorder will, at best, identify relatively nonspecific genetic markers that are either present with relatively high frequency among the few common symptom phenotypes or distributed evenly across many different phenotypes of the same disorder. On the other hand, even relatively large effects will be missed if they result from a process that is only relevant in a few uncommon phenotypes. Although the use of larger samples will allow correlates with increasingly smaller effects to be detected, the cumulative sizes of these effects are still likely to be limited by the need to represent shared effects across heterogeneous phenotypes. This again, jeopardizes the clinical usefulness of polygenic risk scores.

Second, although random sampling from the population will result in a sample group with a representative share of individuals with rare phenotypes, the selected rare phenotypes themselves are going to vary substantially among different random samples. Because each common phenotype has a high probability of being included in a random sample and each rare phenotype has a very low chance, the most common phenotypes can be found in virtually all random samples of a given population and each rare phenotype is likely to be found only in few samples. The generalizability of the results of such studies across multiple samples is therefore fundamentally limited. This poses a significant challenge to the application of the principles of evidence-based medicine in psychiatry as long as disorders are defined by sets of symptoms.

Third, patients with rare phenotypes are less likely to benefit from standardized treatments, such as manualized psychotherapies. Such treatments typically address problems considered common among individuals with the same diagnosis, and they are rarely offered on an individualized basis in clinical trials investigating their efficacy. These treatments are thus less likely to focus on problems and symptom relevant to individuals with uncommon phenotypes and might therefore be less effective in individuals with rare phenotypes. This might help explain the fact that up to fifty percent of subjects in clinical trials of psychotherapy interventions do not respond to treatment (Hofmann et al., 2012) and why the effect of psychopharmacological treatment is also limited in many disorders.

Our analyses are limited by the fact that they cannot uncover the mechanisms leading to our results. We speculate, that the distribution might be a consequence of the covariance of symptoms. This is, however, out of the scope of this study. An additional limitation is that symptom phenotypes cannot fully represent the complex contextual meaning of the symptoms expressed by each individual: rather, they serve as a simplified model (Nunes et al., 2020). In addition, our results reflect only the five investigated disorders which represent an important but limited selection of mental disorders. Nevertheless, our results highlight that the observed heavy tailed distribution of symptom phenotypes describe an underrepresentation of individuals with rare symptom phenotypes in mental health research, bringing into question the validity of basic and clinical research for such individuals.

References

- Abi-Dargham, A., & Horga, G. (2016). The search for imaging biomarkers in psychiatric disorders. *Nature Medicine*, 22(11), 1248–1255. <https://doi.org/10/f9c2pm>
- American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorder* (5th ed.).
- Asmundson, G. J., Frombach, I., McQuaid, J., Pedrelli, P., Lenox, R., & Stein, M. B. (2000). Dimensionality of posttraumatic stress symptoms: A confirmatory factor analysis of DSM-IV symptom clusters and other symptom models. *Behaviour Research and Therapy*, 2(38), 203–214. [https://doi.org/10.1016/s0005-7967\(99\)00061-3](https://doi.org/10.1016/s0005-7967(99)00061-3).
- Barabási, A.-L., & Albert, R. (1999). Emergence of Scaling in Random Networks. *Science*, 286(5439), 509–512. <https://doi.org/10.1126/science.286.5439.509>
- Blevins, C. A., Weathers, F. W., Davis, M. T., Witte, T. K., & Domino, J. L. (2015). The Posttraumatic Stress Disorder Checklist for *DSM-5* (PCL-5): Development and Initial Psychometric Evaluation. *Journal of Traumatic Stress*, 28(6), 489–498. <https://doi.org/10.1002/jts.22059>
- Brewin, C. R., Cloitre, M., Hyland, P., Shevlin, M., Maercker, A., Bryant, R. A., Humayun, A., Jones, L. M., Kagee, A., Rousseau, C., Somasundaram, D., Suzuki, Y., Wessely, S., van Ommeren, M., & Reed, G. M. (2017). A review of current evidence regarding the ICD-11 proposals for diagnosing PTSD and complex PTSD. *Clinical Psychology Review*, 58, 1–15. <https://doi.org/10/gcnjjr>
- Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-Law Distributions in Empirical Data. *SIAM Review*, 51(4), 661–703. <https://doi.org/10.1137/070710111>
- Cuthbert, B. N., & Insel, T. R. (2013). Toward the future of psychiatric diagnosis: The seven pillars of RDoC. *BMC Medicine*, 11(1), 126. <https://doi.org/10.1186/1741-7015-11-126>
- Eeden, W. A., Hemert, A. M., Carlier, I. V. E., Penninx, B. W., & Giltay, E. J. (2019). Severity, course trajectory, and within-person variability of individual symptoms in patients with major depressive disorder. *Acta Psychiatrica Scandinavica*, 139(2), 194–205. <https://doi.org/10/gj6rvb>
- Fletcher, S., & Islam, M. Z. (2018). Comparing sets of patterns with the Jaccard index. *Australasian Journal of Information Systems*, 22. <https://doi.org/10.3127/ajis.v22i0.1538>
- Foss, S., Korshunov, D., & Zachary, S. (2013). Heavy-Tailed and Long-Tailed Distributions. In S. Foss, D. Korshunov, & S. Zachary (Eds.), *An Introduction to Heavy-Tailed and Subexponential Distributions* (pp. 7–42). Springer. https://doi.org/10.1007/978-1-4614-7101-1_2
- Fried, E. I., & Nesse, R. M. (2015). Depression is not a consistent syndrome: An investigation of unique symptom patterns in the STAR*D study. *Journal of Affective Disorders*, 172, 96–102. <https://doi.org/10.1016/j.jad.2014.10.010>
- Galatzer-Levy, I. R., & Bryant, R. A. (2013). 636,120 Ways to Have Posttraumatic Stress Disorder. *Perspectives on Psychological Science*, 8(6), 651–662. <https://doi.org/10.1177/1745691613504115>
- Gillespie, C. S. (2015). Fitting Heavy Tailed Distributions: The **poweRlaw** Package. *Journal of Statistical Software*, 64(2). <https://doi.org/10.18637/jss.v064.i02>
- Hickling, E. J., Barnett, S. D., & Gibbons, S. (2013). The Many Presentations of Posttraumatic Stress Disorder: An Empirical Examination of Theoretical Possibilities. *SAGE Open*, 3(1), 215824401348015. <https://doi.org/10/gghkk8>
- Hofmann, S. G., Asnaani, A., Vonk, I. J. J., Sawyer, A. T., & Fang, A. (2012). The Efficacy of Cognitive Behavioral Therapy: A Review of Meta-analyses. *Cognitive Therapy and Research*, 36(5), 427–440. <https://doi.org/10/f4cm73>

- Insel, T., Cuthbert, B., Garvey, M., Heinssen, R., Pine, D. S., Quinn, K., Sanislow, C., & Wang, P. (2010). Research Domain Criteria (RDoC): Toward a New Classification Framework for Research on Mental Disorders. *American Journal of Psychiatry*, 167(7), 748–751. <https://doi.org/10.1176/appi.ajp.2010.09091379>
- Jaccard. (1901). Distribution de la flore alpine dans le bassin des Dranses et dans quelques régions voisines. *Bull Soc Vaudoise Sci Nat*, 37, 241–272.
- Kapur, S., Phillips, A. G., & Insel, T. R. (2012). Why has it taken so long for biological psychiatry to develop clinical tests and what to do about it? *Molecular Psychiatry*, 17(12), 1174–1179. <https://doi.org/10/f4fbx5>
- Kay, S. R., Fiszbein, A., & Opler, L. A. (1987). The Positive and Negative Syndrome Scale (PANSS) for Schizophrenia. *Schizophrenia Bulletin*, 13(2), 261–276. <https://doi.org/10/gdz73b>
- Klonsky, E. D., & Olino, T. M. (2008). Identifying clinically distinct subgroups of self-injurers among young adults: A latent class analysis. *Journal of Consulting and Clinical Psychology*, 76(1), 22–27. <https://doi.org/10/cpnq29>
- Kroenke, K., Spitzer, R. L., & Williams, J. B. W. (2001). The PHQ-9: Validity of a brief depression severity measure. *Journal of General Internal Medicine*, 16(9), 606–613. <https://doi.org/10.1046/j.1525-1497.2001.016009606.x>
- Lovibond, P. F., & Lovibond, S. H. (1995). The structure of negative emotional states: Comparison of the Depression Anxiety Stress Scales (DASS) with the Beck Depression and Anxiety Inventories. *Behaviour Research and Therapy*, 33(3), 335–343. <https://doi.org/10/fwhsp7>
- Lux, V., & Kendler, K. S. (2010). Deconstructing major depression: A validation study of the DSM-IV symptomatic criteria. *Psychological Medicine*, 40(10), 1679–1690. <https://doi.org/10.1017/S0033291709992157>
- Maslach, C., Jackson, S., & Leiter, M. P. (1996). *The Maslach Burnout Inventory* (3rd ed.). Consulting Psychologists Press.
- Mitzenmacher, M. (2004). A Brief History of Generative Models for Power Law and Lognormal Distributions. *Internet Mathematics*, 1(2), 226–251. <https://doi.org/10.1080/15427951.2004.10129088>
- Nandi, A., Beard, J. R., & Galea, S. (2009). Epidemiologic heterogeneity of common mood and anxiety disorders over the lifecourse in the general population: A systematic review. *BMC Psychiatry*, 9(1), 31. <https://doi.org/10/c76wxw>
- Newman, M. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, 46(5), 323–351. <https://doi.org/10.1080/00107510500052444>
- Nunes, A., Trappenberg, T., & Alda, M. (2020). The definition and measurement of heterogeneity. *Translational Psychiatry*, 10(1), 299. <https://doi.org/10/ghm2pb>
- Olbert, C. M., Gala, G. J., & Tupler, L. A. (2014). Quantifying heterogeneity attributable to polythetic diagnostic criteria: Theoretical framework and empirical application. *Journal of Abnormal Psychology*, 123(2), 452–462. <https://doi.org/10/f6mvmb>
- Rasmussen, A., Verkuilen, J., Jayawickreme, N., Wu, Z., & McCluskey, S. T. (2019). When Did Posttraumatic Stress Disorder Get So Many Factors? Confirmatory Factor Models Since DSM–5. *Clinical Psychological Science*, 7(2), 234–248. <https://doi.org/10.1177/2167702618809370>
- Thorp, J. G., Marees, A. T., Ong, J.-S., An, J., MacGregor, S., & Derks, E. M. (2020). Genetic heterogeneity in self-reported depressive symptoms identified through genetic analyses of the PHQ-9. *Psychological Medicine*, 50(14), 2385–2396. <https://doi.org/10/gj6r2w>
- Ulrich, W., Ollik, M., & Ugland, K. I. (2010). A meta-analysis of species–abundance distributions. *Oikos*, 119(7), 1149–1155. <https://doi.org/10/fpqpvk>

- van Loo, H. M., de Jonge, P., Romeijn, J.-W., Kessler, R. C., & Schoevers, R. A. (2012). Data-driven subtypes of major depressive disorder: A systematic review. *BMC Medicine*, 10(1), 156. <https://doi.org/10.1186/1741-7015-10-156>
- von Känel, R., van Nuffel, M., & Fuchs, W. J. (2016). Risk assessment for job burnout with a mobile health web application using questionnaire data: A proof of concept study. *Biopsychosocial Medicine*, 10. <https://doi.org/10.1186/s13030-016-0082-4>
- Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., ... Yutani, H. (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, 4(43), 1686. <https://doi.org/10.21105/joss.01686>
- World Health Organization. (1992). *The ICD-10 classification of mental and behavioural disorders: Clinical descriptions and diagnostic guidelines*. World Health Organization.
- Young, G., Lareau, C., & Pierre, B. (2014). One Quintillion Ways to Have PTSD Comorbidity: Recommendations for the Disordered DSM-5. *Psychological Injury and Law*, 7(1), 61–74. <https://doi.org/10/gfddmn>

Acknowledgments: Data and/or research tools used in the preparation of this manuscript were obtained from the National Institute of Mental Health (NIMH) Data Archive (NDA). NDA is a collaborative informatics system created by the National Institutes of Health to provide a national resource to support and accelerate research in mental health. Dataset identifier(s): DOI: 10.15154/1522637. This manuscript reflects the views of the authors and may not reflect the opinions or views of the NIH or of the Submitters submitting original data to NDA. This material is the result of work supported with resources and the use of facilities at the West Haven VA Medical Center.

Funding:

Swiss National Science Foundation, Early.Postdoc Mobility Fellowship, grant (P2ZHP3_195191) (TRS)

National Institutes of Health (NIH), grant R01MH112746 (JM)

Author contributions:

Conceptualization: TRS, OD, IHR

Methodology: TRS, OD, MH, JM

Investigation: TRS, OD, MH, JM, RvK, IHR

Visualization: TRS, OD

Funding acquisition: TRS, RvK, IHR

Project administration: TRS, OD, IHR

Supervision: JM, RvK, IHR

Writing – original draft: TRS, OD

Writing – review & editing: TRS, OD, MH, JM, RvK, IHR

Competing interests: Authors declare that they have no competing interests.

Data and materials availability: Data for sample 1 (PANSS) is available through the National Institute of Mental Health Data Archive website using the DOI 10.15154/1522637. However, approval of the NIMH is required, which can be requested here: <https://nda.nih.gov/get/access-data.html>. Data for sample 2 and 5 (PCL-5 & PHQ-9) are subject to legal restrictions on confidential data of the United States Department of Veterans Affairs (VA) and cannot be shared by the authors. Free access to the data can be requested from the VA, however, access might be denied. More information can be found online (<https://www.virec.research.va.gov>) or by contacting Dr. Harpaz-Rotem or the VA research office at the West Haven VA Medical Center (<https://www.connecticut.va.gov/services/Research.asp>). Data for sample 3 (DASS) can be obtained online at https://openpsychometrics.org/_rawdata/. Data for sample 4 (MBI) are available upon reasonable request from the corresponding author but might only be used to verify the presented analysis and cannot be made available for commercial or other research purposes.

Disclaimer: The contents do not represent the views of the U.S. Department of Veterans Affairs or the United States Government.

	1. PANSS	2. PCL	3. DASS	4. MBI	5. PHQ
Characteristics					
Sample Size	4,616	52,609	39,700	24,323	165,397
Questionnaire used	PANSS	PCL-5	DASS	MBI	PHQ-9
Number of items	14	20	13	16	9
Phenotype (range of each symptom)	binary	binary	binary	binary	Four levels
Phenotypes					
Max. possible	$2^{14} = 16,384$	$2^{20} = 1,048,576$	$2^{13} = 8,192$	$2^{16} = 65,536$	$4^9 = 262,144$
Number of phenotypes found	1,185	8,174	4,528	5,768	39,341
Frequency of most common phenotype	950	20,575	6,999	1,682	5,858
Frequency of phenotypes, median	16	2453	49	21	12
Diversity Indices					
Jaccard-Index, median [Q1, Q3]	0.798 [0.72, 0.87]	0.92 [0.86, 0.98]	0.74 [0.67, 0.74]	0.61 [0.50, 0.72]	NA
Parameters of Power-Law					
$x_{\min}$	1	8	7	1	16
α	2.07	1.86	2.1	2.05	2.43
p -value	0.222	0.865	0.123	0.455	0.272
Comparison with Lognormal Distribution					
Two-sided, p -value	0.752	0.954	0.579	0.854	0.748
One-sided, p -value	0.376	0.522	0.438	0.573	0.625

Table 1. Characteristics of the included samples, their phenotype frequencies, and the parameters of the power-law models.

	PHQ		
Symptom burden	Low	Med	High
Characteristics			
Sample Size	6646	8283	7645
Phenotypes			
Number of phenotypes found	150	339	167
Frequency of most common phenotype	939	352	1318
Frequency of phenotypes, median	7	7	5
Diversity Indices			
Jaccard-Index, median [Q1, Q3]	0 [0,0.33]	0.42 [0.33,0.6]	0.75 [0.62,0.85]
Parameters of Power-Law			
$x_{\min}$	105	20	14
α	2.62	2.04	1.77
p -value	0.277	0.026	0.242
Comparison with Lognormal Distribution			
Two-sided, p	0.799	0.056	0.252
One-sided, p	0.399	0.971	0.873

Table 2. Characteristics, phenotype frequencies and the parameters of the power-law models of the subsamples investigated in the second complimentary analysis.

Figure 1. Distribution of frequency of symptom phenotypes. (A) The fifty most common phenotypes of each study are represented by bars and are ordered with descending frequency along the x-axis. The y-axis denotes the number of times a given phenotype was reported. Symptom phenotypes in studies 1-4 were based on binarized symptoms. In study 5, the symptom phenotypes were defined on the basis of the original, four-level coding, of the individual symptoms. (B) For each study, the x-axis denotes the number of times a phenotype was reported and the y-axis the complementary cumulative distribution function (CDF), indicating the probability of a given value X (the frequency of a given phenotype) to take a value above x (a frequency higher than the one of a given phenotype). The black circles represent all phenotypes with the same specific frequency. The red (solid) line and the blue (dashed) line represent the maximum likelihood power-law and log-normal distributions respectively.