THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Why do traits come together? The underlying trait and network approaches

Citation for published version:

Mottus, R & Allerhand, M 2018, Why do traits come together? The underlying trait and network approaches. in V Zeigler-Hill & TK Shackelford (eds), *SAGE Handbook of Personality and Individual Differences : The Science of Personality and Individual Differences.*, 6, SAGE, London.
<https://doi.org/20.500.11820/385b5bbd-bf64-4726-9cda-5685e961b032>

Digital Object Identifier (DOI):

[20.500.11820/385b5bbd-bf64-4726-9cda-5685e961b032](https://doi.org/20.500.11820/385b5bbd-bf64-4726-9cda-5685e961b032)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

SAGE Handbook of Personality and Individual Differences

Publisher Rights Statement:

This is the accepted version of the following chapter : "Why do traits come together? The underlying trait and network approaches"
Mottus, R. & Allerhand, M. 26 May 2018 which has been published in final form in V. Zeigler-Hill, & T. K. Shackelford (Eds.), *SAGE Handbook of Personality and Individual Differences : The science of personality and individual differences* (Vol. 1). [6] London: SAGE, <https://uk.sagepub.com/en-gb/eur/the-sage-handbook-of-personality-and-individual-differences/book249310>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Accepted for publication.

To appear as: Möttus, R., & Allerhand, M. (in press). Why do traits come together? The underlying trait and network approaches. In V. Zeigler-Hill & T. K. Shackelford (Eds), *SAGE handbook of personality and individual differences: Volume 1. The science of personality and individual differences* (pp. xx - xx). London: SAGE.

Why do traits come together? The underlying trait and network approaches

René Möttus and Mike Allerhand

Department of Psychology and Centre for Cognitive Ageing and Cognitive Epidemiology,
University of Edinburgh, UK

Abstract

This chapter deals with one of the most pervasive personality-related phenomenon: the coalescence of tendencies for specific thoughts, feelings and behaviors (characteristics) into broader patterns—traits. Two possible explanations are discussed. The more established explanation is that certain characteristics tend to co-exist because they reflect a common underlying cause. A more recent explanation is that they may also hang together because of having direct causal links between them—some characteristics can contribute to, or inhibit, others. However, the chapter offers a more general, mathematically formalized framework, which can, in fact, merge the two explanations. Furthermore, this framework can be used to represent both processes within individuals and individual differences, with the latter emerging from the former. This means potential for a formal bridge between two branches of personality psychology—the social cognitive and trait approaches. Some empirical findings will be reviewed that are consistent with the proposed framework.

Personality characteristics are more or less stable ways of feeling, thinking, and behaving. One of the most persistent observations is that they come in patterns: some characteristic levels are more likely to co-exist than others. We start the chapter by reviewing the long-dominant explanation for this phenomenon: the existence of relatively few and distinct underlying traits, each of which causes multiple personality characteristics. We will go on by arguing that empirical evidence supporting the underlying trait paradigm is currently not unequivocal. Likewise, the explanatory scope of this paradigm may be somewhat limited, as it does not accommodate theories as to how psychological processes happening within individuals intersect with, or lead to, individual differences. These limitations leave room for complementary perspectives. One of them is the network approach, which posits personality as a dynamical system of interconnected characteristics. This approach can account for the clustering of personality characteristics and individual differences by means of within-individual processes. However, this approach is far less parsimonious. Towards the end of the chapter, we will offer a more general, mathematically formalized framework that combines the two approaches.

The broad trait paradigm

For decades, a dominant paradigm of personality research has been that of broad personality traits that encompass individual differences in wide ranges of behaviors, feelings and thoughts (see McCrae, this volume). This is why the term *trait* will be exclusively used to refer to broad personality constructs in this chapter. Journals routinely publish articles on a number of trait-related questions, including their structure (Laverdière, Morin, & St-Hilaire, 2013), variation across demographic groups (Schmitt, Realo, Voracek, & Allik, 2008) or geographical locations (Schmitt, Allik, McCrae, & Benet-Martinez, 2007), correlations with life outcomes (Roberts, Kuncel, Shiner, Caspi, & Goldberg, 2007), or correspondence to neuronal (Bjørnebekk et al., 2013) and genomic variations (de Moor et al., 2012).

Traits are population structures although often interpreted as pertaining to individuals

A major reason for such pervasive focus on traits is that specific behaviors, thoughts, feelings, motives, attitudes, values—the stuff that personality is made of, or becomes manifest through, which we generically refer to as personality *characteristics*—tend to vary across people in certain patterns. For example, a person who is rated as above average in moodiness is also likely to appear commensurately anxious and prone to feeling helpless, whereas a chatty person is also expected to feel positive emotions and be up for adventures. A parsimonious explanation for such clustering is that the co-occurring characteristics reflect some shared underlying etiology. Moody people also being anxious and feeling helpless may be because these specific characteristics reflect a broader underlying propensity (trait) for negative emotions known as Neuroticism. And if so, it makes sense to focus on these underlying traits rather than on the multitude of their manifestations, for both practical (simplicity) and theoretical (parsimony) purposes. If this hypothesis is correct, then personality ought to be a relatively small set of real attributes possessed by individuals that correspond to some stable parameters of their nervous systems and are more or less universal across the human kind—and possibly beyond.

Metaphorically, the underlying traits can be thought of as hidden 'generators' that, perhaps in interaction with situations, produce the observable regularities that appear as traits and are thereby captured using personality trait questionnaires. Individual differences in the trait manifestations largely result from differences in the power of these 'generators.' This metaphor parallels how traits are often modeled using structural equation modeling, where hypothetical latent entities, denoted by circles, cause their indicators, denoted by rectangles (Figure 1). The indicators only co-exist because of sharing the underlying cause: if all individuals were equal on the underlying traits, they would not cluster in any ways.

Figure 1. *The model of a trait as an underlying cause of its associated characteristics (a, b, c, d, and e). This is how personality traits are mostly represented in structural equation models.*

In a vast majority of studies, either some form of factor analysis (FA) or principal component analysis (PCA) is used to study how personality characteristics cluster into what could then be interpreted as reflections of underlying traits. Mostly, differences between individuals measured at one particular point of time are analyzed. Less often, patterns of variation within individuals are investigated (Möttus, Epskamp, & Francis, in press), in which case it is expected that fluctuations in the underlying trait levels over time and situations (or due to experimental manipulations, for that matter) lead to their associated characteristics fluctuating in concert. When individuals are moodier than is typical for them, they should also be commensurately more anxious than usual because moodiness and anxiety are hypothesized to reflect the same underlying trait. Other types of evidence for particular characteristics reflecting underlying common causes such as consistency in their developmental trajectories have been even less frequently used to carve out traits, although developmental consistency also appears to be a necessary property of etiologically unitary traits (Cattell, 1946).

However, as the old adage goes, correlation does imply causation. Patterns of clustering may have different causes. Therefore, to interpret principal components (or factors, for that matter) extracted from correlational data that pertain to individual differences as indicative of underlying attributes *of* these individuals (i.e., causal structures within them that can cause something) is not, in fact, a straightforward inference. One way to see this is to review the mechanics of PCA—and effectively FA. (We will come back to something similar later in the chapter, so there is even more reason to get to grips with PCA.)

We can think of PCA as a procedure that projects individuals into a multidimensional space spanned by k orthogonal dimensions. These dimensions correspond to the k observed personality variables being analyzed. For example, the variables can be items or facets (clusters of similar items) of a personality test. For convenience, we can assume that all variables can be centered at zero and have unit variance. Each individual can be represented as a point in this personality space, with their coordinates being their values on the measured variables. Equivalently, each individual can be thought of as a vector starting from the origin and ending in the location with the said coordinates. Put this way, each individual is characterized by the direction and length of their *person vector*, in relation to those of the other individuals represented in this space. Collectively, the person vectors thus represent the population structure of the personality feature space. PCA attempts to identify orthogonal dimensions in this space along which individuals tend to vary the most. These dimensions are also vectors (eigenvectors), although they are called principal components. The first principal component vector projects into the space in the direction to which the person vectors most

often point and that therefore has the greatest variance in their lengths. Extracting the first component means removing all the variance in person vector lengths in this particular direction—so that we can think of the space being squeezed “flat” in this particular direction. Using the variability in the remaining dimensions, subsequent principal component vectors are extracted in exactly the same way, until all the variance has been removed.

I deliberately used the term vector to denote both people and principal components. This helps to see that the principal component vectors are essentially abstractions of person vectors, describing popular directions among them. They summarize clusters in the population structure and thereby pertain to differences (or similarities, if you will) between people. For example, Figure 2 (left panel) represents a situation where there is no ordered population structure, with no one direction being more popular among the person vectors than any other. In such a case, no principal components would usefully summarize directions of person vectors. In the middle panel of Figure 2, in contrast, person vectors have a popular direction and a principal component can usefully summarize such clustering of person vectors. Naturally, there might be more than one popular direction that can be summarized by corresponding principal components (right panel of Figure 2). This is the most realistic scenario for a personality space.

Figure 2. *Person vectors in three-dimensional feature space.*

The components, however, may say little about the structure of the measured attributes *within* individuals, because the population clustering may result from more than one kind of mechanisms. Personality vectors may point in particular directions because of some universal processes within individuals that cause some of the personality features tend to gravitate toward similar values independently of other characteristics (reflecting their unique underlying common cause), which translates into some configurations of personality features and thereby directions of person vectors being more popular than some others. This would be the underlying trait scenario.

But the clustering may also happen due to other reasons. For example, some configurations of personality characteristics (and thereby directions in personality space) may reflect environmental niches toward which groups of people tend to gravitate. Or, individuals may interact based on their pre-existing personality features and become more similar over time, which may also cause some directions of personality vectors to become more popular than others. If so, the tendencies for particular personality characteristics to co-exist may be rather incidental in the sense that there may be no common cause for specifically these particular characteristics. Finally, the network approach, described in more detail below, posits that personality characteristics might directly influence each other, suggesting that some characteristics may tend to have similar values not because of an underlying common cause, or because of any common cause external to them, but because they “top up” each other.

A recurrent population structure

It is often suggested that population variance in personality characteristics can be efficiently summarized along five directions, known as the Five-Factor Model (FFM) traits of Neuroticism, Extraversion, Openness to Experience, Agreeableness, and Conscientiousness (McCrae & John, 1992). The FFM has been recovered across numerous studies and cultures (Allik, Realo, & McCrae, 2013), at least when the pool of measured trait indicators is pre-specified, as is the case when standardized and carefully adapted personality questionnaires that are pre-designed to measure the particular combinations of personality characteristics are used and data reduction techniques such as PCA are applied (McCrae & Terracciano, 2005).

These findings are not trivial: there is no *a priori* reason why certain personality characteristics should co-exist in a given culture even if researchers have specifically chosen to model the characteristics that tend to co-exist in some other culture. If the popular directions in personality space reflected some sort of environmental niches, these surely could vary across cultural contexts. Therefore, in addition to cross-cultural replicability constituting strong evidence for the usefulness of the FFM traits for *summarizing* individual differences in personality characteristics in a variety of contexts, it may also be evidence for the traits being ontologically “real”. Yet there are some potential concerns.

First, the replication of the FFM appears somewhat poorer when traits are defined *de novo* in each culture rather than being recovered from pre-defined questionnaires (De Raad et al., 2014). This may suggest that some of the apparent cross-cultural replicability may have been designed into the studies (Saucier et al., 2014). The replication may also be poorer when rarely addressed indigenous societies are studied (Gurven, von Rueden, Massenkoff, Kaplan, & Lero Vie, 2013). Some authors have suggested that two (Saucier et al., 2014), three (De Raad et al., 2010, 2014), or six (Ashton & Lee, 2007) components—which sometimes do overlap with the FFM traits—provide the most robust way to describe the population structure of personality characteristics. The degree of replication may also depend on which statistical methods are used. For example, while some studies (e.g., McCrae & Terracciano, 2005) have used PCA followed by targeted rotation (McCrae, Zonderman, Costa, Bond, & Paunonen, 1996), attempts to use arguably more modern statistical methods to test for factorial invariance across cultures point to poorer replication (Thalmayer & Saucier, 2014).

Second, FFM measures that are designed to recover the five FFM traits as clearly as possible demonstrate complex population structures even in their original cultural contexts, with individuals' person vectors systematically deviating from directions that would correspond to the FFM traits. In other words, items or facets correlate with multiple FFM traits at the same time (Hopwood & Donnellan, 2010) and scales lack local independence by correlating even when there appears to be no variance in their ostensibly underlying traits (Gignac, Bates, & Jang, 2007). On one hand, this may seem like some nuisance that can be easily dismissed: measurements could be contaminated by systematic biases, which does not necessarily undermine the trait model that underlies these measurements. After all, no measurement is perfect. On the other hand, the ontological status of the FFM traits hinge on how clearly they can be identified from data because this is how they were postulated in the first place. Therefore, the measurements of FFM also constitute the main means of falsifying the model.

Tackling some of the complexities of trait conceptualization and measurement, it is widely accepted that at different levels of abstraction different personality constructs are appropriate. That is, from a trait perspective personality is thought to have a hierarchical structure (Markon, Krueger, & Watson, 2005), ranging from one or two overarching constructs (DeYoung, 2006; Rushton, Bons, & Hur, 2008) to very specific characteristics reflected in single test items (McCrae, 2015; Möttus, Kandler, Bleidorn, Riemann, & McCrae, in press). This trait representation can account for why most personality characteristics correlate to some extent (van der Linden, te Nijenhuis, & Bakker, 2010), subsets of them correlate more highly, some subsets of these subsets yet more highly, and so on. Arguably, however, when virtually any trait can be split into further parts which then can be split

into yet further parts, and so on, deciding which particular combination of traits to focus on for whatever practical or conceptual purpose will become a very difficult decision. Also, one of the primary strength of the underlying trait paradigm, parsimony, somewhat dissolves in the trait hierarchy, as a series of underlying influences need to be postulated across its levels.

More arguments regarding the reality of traits

Three other types of evidence are often cited for the reality of personality traits: heritability, at least moderate cross-rater agreement, and moderate to high rank-order stability (Funder, 1991; McCrae et al., 2004; McCrae & Costa, 2008a). All of them seem to provide support for the *descriptive* utility of broad traits—that these traits can be useful for describing these properties of human personality. However, one may argue that these types of evidence may not speak for the *reality* of the traits after all.

As for heritability, twin-studies have shown that about 40% of the phenotypic variance in the FFM traits can be attributed to genetic similarity between individuals (Vukasovic & Bratko, 2015). Likewise, the covariation structure of traits has shown to be heritable (McCrae, Jang, Livesley, Riemann, & Angleitner, 2001), such that, for example, cheerfulness of one twin tends to co-vary with talkativeness of another twin (commonly referred as genetic correlation). However, heritability cannot be taken as evidence for one collection of characteristics being more likely to represent a real trait than any other. This is because heritability might be the property of the characteristics that are aggregated into traits rather than pertain to whatever (underlying or not) the aggregate is assumed to represent. Indeed, individual personality test items are heritable, often even when the variance of the FFM traits has been removed from them (Möttus, Kandler, et al., in press). As per the First Law of Behavior Genetics, everything is heritable (Turkheimer, 2000)—as long as behavioral traits are concerned, anyway. For example, even pseudo traits that have been compiled from theoretically unrelated but heritable variables are to a substantial degree heritable (Johnson, Penke, & Spinath, 2011).

Likewise, the presence of a genetic correlation does not necessarily show involvement of overlapping genes in the etiology of two variables as is sometimes mistakenly assumed because the genetic factors underlying one variable may bleed into another via a phenotypic causation (Johnson et al., 2011). For example, one characteristic may mediate the heritable influences of another such as smoking mediates the effect of some genetic variants on lung cancer (that is, what appear as genetic variants for cancer gene are actually genetic variants for smoking that causes cancer). As a result, although the FFM structure can be recovered from genetic correlations, this does not necessarily mean that the genetic effects on measured variables are mediated by latent factors as would be predicted by the underlying trait model (Franic, Borsboom, Dolan, & Boomsma, 2014; but see also Lewis & Bates, 2014). Also, the popular directions among person vectors that the FFM factors summarize might reflect popular genetic configurations of psychological characteristics. That is, it is the *directions* of person vectors that might be heritable rather than something underlying some specific sets of the characteristics that constitute these directions.

As with heritability, cross-rater agreement and rank-order stability may also pertain to the characteristics that are aggregated into traits rather than to the traits themselves. In fact, individual test items are agreed upon by different raters even when trait variance has been removed from them (Möttus, McCrae, Allik, & Realo, 2014). The same argument probably applies for predictive validity of personality traits, which has also been used to “empower” traits (Roberts et al., 2007). For a parallel illustration, although high socioeconomic status (a composite score of, say, educational level, occupational status, income, neighborhood quality) is correlated with just about every good thing in life, it would be presumptuous to think that the indicators of socioeconomic status are only correlated because of an underlying common cause. In fact, it is even possible that pseudo traits consisting of items that measure different (FFM) traits would generally show higher predictive validity than traits that consist of similar items, given that the aggregated items are relevant for the outcome. This is because combinations of different items then encompass more

outcome-relevant content than combinations of similar items.

Traits are snapshots

A potential limitation of the underlying trait perspective is that its explanatory scope is limited to the description of static individual differences. Hardly anyone would doubt that humans are complex and dynamic systems embedded in environments that they themselves choose, modify, and react to. If so, any comprehensive account of human personality has to account for this complexity and dynamism. Although trait models can provide efficient descriptive snapshots of individual differences, they have little to say on how individuals operate from moment to moment in interaction with each other and their surroundings and how, at any given moment, these processes underlie and spring from their relatively static trait-standings. According to McCrae and Costa (2008a), "... if one wishes to understand the processes that lead to the flow of behavior and experience in individual people, trait psychology is a limited guide" (p. 288). Of course, trait theories can and have been expanded to account for phenomena other than the relatively static traits—for example, the characteristic adaptations and other phenomena encompassed by the Five-Factor Theory (McCrae & Costa, 2008b; McCrae, this volume)—but there are currently few theories that explain how individual differences in personality characteristics and co-variation patterns in these can emerge from the processes happening with individuals.

Of course, the trait approach with its predominant focus on static individual differences is not the only approach to personality, although it, arguably, has dominated the research field for decades. Some other approaches, grounded in social psychology (Cervone, 2004; Mischel & Shoda, 1995), have for the most part been exactly the opposites of the trait approach, focusing on the processes within individuals with relatively little systematic attention to how these specific processes can give rise to the broad phenomena of individual differences that interest trait researchers. The field of human personality research has therefore been characterized by a chasm that has not exactly helped with the emergence of a comprehensive view of what personality is and how it works (Fleeson, 2012; Mischel & Shoda, 1998; McCrae, 2009). One reason for this chasm may be that there has been little in the way of a formal framework for connecting the two approaches. This may be changing, however, as the network approach attempts precisely that—to provide one framework for combining the within- and between-individual sides of personality psychology.

The network approach

According to the network perspective on personality (Cramer et al., 2012), personality characteristics can have "causal, homeostatic or logical" (p. 415) associations between them (Schmittmann et al., 2013). In the network language, the elements of the network are referred to as *nodes* (here, personality characteristics) and positive or negative *edges* between them (here, associations between personality characteristics). The edges can be directed depicting how causality is expected to flow, or undirected, simply denoting associations without specifying the flow of causation. For visual ease, positive edges can be shown in green and negative edges in red. A hypothetical network of the associations between five nodes is depicted in Figure 3.

Figure 3. A hypothetical network of five personality characteristics (*a*, *b*, *c*, *d*, and *e*) and their relationships. The characteristics are referred to as nodes and their relationships as edges. Positive edges are often shown in green and negative edges in red.

For a concrete example, sadness, anger, and impulsive behavior may co-vary because of reciprocal associations between them. Doing things that will be regretted afterwards may contribute to low mood and anger, and lashing out angrily may lead to feeling sadness afterwards. To the extent that all of these combinations of causal associations, or at least some of them, are consistent across individuals, this may contribute to the co-variation pattern that we interpret as Neuroticism. Note that as in Figure 3, not all nodes of this hypothetical network have direct edges between them. As I will describe below, the purpose of the network analysis is delineating the *unique* associations among variables rather than general correlation patterns.

Personality characteristics may co-vary even if they do not have direct links between them but because they share the sets of other elements to which they are connected. First, they may have indirect causal associations: being self-disciplined may be indirectly associated with being gregarious because it contributes to appearing reliable—something that helps with making friends and thereby having people around. Second, there may be spurious correlations due to shared co-variables: being self-disciplined may also correlate with gregariousness because of shared negative contributions from sadness and anxiety. Another example of spurious correlations comes from the functionalist explanation for the coalescence of traits (Wood, Gardner, & Harms, 2015). According to this perspective, what are typically considered personality characteristics co-vary because they are linked with overlapping functionality indicators (self-perceived abilities, expectancies, and valuation): the same goals can be achieved in multiple ways and the same behaviors can serve multiple goals. When we consider the functionality indicators as nodes of the psychological network that appears as personality, they will become the causes of spurious correlations among personality characteristics.

Moreover, characteristics that are traditionally associated with different purported (underlying) traits can have direct or indirect associations between them. For example, keeping promises (a component of Conscientiousness measures) may tend to have a causal association with being seen as co-operative (a component of Agreeableness measures), or consuming culture (tapped by Openness measures) may help to make friends (asked about in Extraversion measures). This may contribute to the commonly observed correlations between the measures of these traits; for example, the correlations between Conscientiousness and Agreeableness measures, or Openness and Extraversion measures, typically exceed .30 (van der Linden et al., 2010; for comparison,

correlations between measures of the same traits are typically around .60, Pace & Brannick, 2010).

Thus, the crucial difference between the underlying trait and network explanations for the coalescence of traits is this. According to the former, the correlation between, say, sadness and anger is *entirely* due to a common cause, the underlying propensity that we may want to call Neuroticism. If the correlation does not vanish after controlling for the underlying trait level (i.e., there is lack of local independence), a new lower-level underlying trait needs to be postulated that accounts for this residual correlation. According to the network perspective, in contrast, the correlation between sadness and anger may result from a direct causal association among them as well as from shared influences from other components of the network. If so, there is no need to worry about local independence or different levels of causal influences (i.e., the trait hierarchy).

Although the network approach can, in principle, do away with the underlying traits altogether, *a priori* precluding any unmodeled causal factors (i.e., underlying traits) seems an unnecessarily restrictive stance and is not inevitably required. Essentially, the shared contributions from other variables—whether explicitly included among the nodes or not—serve the function of common causes for two or more nodes. If the shared contributions do not happen to be explicitly modeled, they become “underlying” by definition. Put differently, should the underlying variables be fleshed out in sufficient detail at some point (e.g., we will know all bits that make up Neuroticism), their elements can become modeled as parts of the personality network. As a result, it appears that the network operationalization of personality is a more general one that can, in principle, accommodate the underlying trait approach.

The network approach combines individual differences and within-individual processes

According to the underlying trait perspective, personality is represented as a set of scores of latent variables and their correlations with the observed personality characteristics (as well as intercepts and [residual] variances, but for simplicity these can be ignored for now). The correlations (factor loadings) are effectively estimates of how reliably the characteristics reflect the latent traits. Because the latent scores are unknown, they are approximated by the sum-scores of appropriate observed variables, possibly weighed by their hypothetical correlations with the latent variables. Essentially, thus, the personality representation amounts to a set of underlying trait scores. Because within-individual processes are not an inherent part of such trait models, the scores are assumed to be relatively stable over moderately long periods of time, even though external (to traits) factors that may cause gradual changes in the scores, via some unknown mechanisms, are often sought after (Specht et al., 2014).

According to the network perspective, personality is a more dynamic phenomena: the nodes, or personality characteristics, are, at least in theory, constantly updated by positive or negative influences from other nodes, be these explicitly modeled or not, or from the respective nodes themselves at previous time-point. As a result, if we choose to model personality processes as discrete steps (if only for convenience), at any given time personality can be represented as its state at the previous time-point that has been updated by the connections among the nodes (including contributions from the nodes themselves). Therefore, this personality representation has two components: node scores and their connections. The former is simply a vector, say, $\mathbf{y}$ of length k for each individual p ($\mathbf{y}_p$). This part of personality representation corresponds to the observed variables in the underlying trait representations before these have been aggregated into trait scores, although there is one important difference: node scores can change from moment-to-moment in the network conceptualization but the observed variables are expected to be stable in the underlying trait paradigm, as otherwise the reliability of measurement would be compromised. The connections between the nodes can be encoded in a k by k weight matrix $\mathbf{W}$. The rule that connects each individual's (p) vector of scores and the weight matrix is $\mathbf{y}_t = \mathbf{y}_{t-1}\mathbf{W}$, where t stands for time. This is how the network representation of personality can formally connect individual differences, encoded in the vector of node scores, and within-individual processes, encoded in the weight matrix. Individual differences are therefore the result of within-individual processes.

This personality representation is consistent with how some personality trait researchers conceive of traits. For example, the Whole Trait Theory (Fleeson & Jayawickreme, 2010) posits that traits have two parts. First, the descriptive part constitutes of the distributions of personality states (personality characteristics, in our terminology). Secondly, the explanatory parts consists of all sorts of underlying mechanisms that contribute to the variance and co-variance of the states over time and across situations. It is just that the network approach formalizes the latter part in more concrete terms.

The off-diagonal elements of the weight matrix $\mathbf{W}$ can be either positive or negative, in which case the corresponding characteristics top up or take off from each other, respectively, or zero, in which case there is no link between the corresponding characteristics. The diagonal of the weight matrix represents “self-loops”, which represent the (autoregressive) stability of the corresponding characteristics—or, put alternatively, their resistance to the influences from the other characteristics. If personality is completely stable, then all off-diagonal elements of $\mathbf{W}$ are zero, while all diagonal elements are 1.

If the matrix denotes causal processes, it is non-symmetric, with rows encoding connections that the corresponding characteristics send out to other characteristics and columns encoding incoming connections from them. A symmetric matrix encodes non-directional connections, which may be a useful way of representing associations when we cannot, or prefer not to, specify the direction of causal flow. For example, when the data being analyzed is cross-sectional, directional connections cannot be estimated because causal processes are likely to take time. In this case, $\mathbf{W}$ represents “average” connection strengths over time, in whatever direction the causality flows. Alternatively, in longitudinal data the testing intervals may be too long to meaningfully delineate causal processes (e.g., doing something silly may not be particularly likely to contribute to sadness 3 years later unless the silly deed was something particularly dramatic).

$\mathbf{W}$ can be specified as invariant across individuals, but it can also be allowed to vary across them ($\mathbf{W}_p$). In the latter case, we end up with having two kinds of individual differences: scores of personality characteristic at any point of time *and* processes that connect and thereby update these characteristics. In other words, individuals may not only differ in the values of their personality characteristic at any given time but also in the processes that give rise to these values. However, because the scores depend on the connections, the latter are more important for describing and understanding personality. For individual differences in $\mathbf{W}$ to become estimable, time-series data can be used, where individuals are measured repeatedly over time-intervals that are short enough for meaningfully revealing such causal processes. For example, increasingly popular experience sampling studies, wherein individuals are asked to report on their personality characteristics several times a day, are well suited for this purpose.

Recently, Epskamp and colleagues (2016) proposed a multi-level framework that allows for simultaneously representing, and empirically estimating, three different kinds of relationships between psychological characteristics. First, the framework consists of a directed network encoding autoregressive and cross-lagged associations among the nodes; this network is suggestive of the stability and change-causing causal processes happening within individuals. As is typical in multi-level models, this network is estimated as a set of fixed effects (representing average associations across individuals) and (individual-specific) deviations from these. Second, the framework represents within-individual contemporaneous associations among the variables, conditional on the autoregressive and cross-lagged links; again, fixed effects and their variability can be estimated. This (undirected) network of associations is suggestive of whatever processes that unfold quickly and therefore cannot be captured by the lag intervals. Third, the framework allows estimating associations that pertain to individual differences—how average levels of characteristics tend to co-vary across individuals. These estimates correspond to the factor-analytic work done within the underlying trait approach. Because these processes are not necessarily indicative of causality—*average* sadness may not be causally related to *average* anxiety, over and above the time-lagged associations already modeled at within-individual variance level—these associations are

represented as an undirected network.

Epskamp and colleagues (2016) exemplified their approach by analyzing the associations among 17 personality characteristics pertaining to the Neuroticism, Extraversion, and Conscientiousness domains and physical exercising. Each of these variables was measured several times each day. Among the notable findings, the personality characteristics clustered according to their FFM traits in all three networks. This is consistent with the possibilities that the coalescence of the measured characteristics into FFM-type traits may partially result from direct causal associations among (consistently with the basic idea of the network approach) them and partly reflect unmodeled influences shared among them (consistently with the underlying trait idea), and that between-individual associations may reflect the results of these processes. Further, physical activity was associated with several characteristics, but the associations varied according to the network type. For example, individuals who on average reported more energy exercised more often and (around the time) when individuals were exercising they also felt more energetic than usually. However, temporal associations revealed a more intricate pattern: while exercising tended to be preceded by feeling more energetic than usual it was followed by feeling less energetic. Indeed, exercising may be exhausting.

Therefore, compared to the underlying trait approach, the network approach allows for a richer way of representing associations between personality characteristics and their associations with variables outside the personality domain (e.g., exercising or whatever other situational factors). This representation combines between-individual differences with within-individual processes that unfold over time or happen (almost) contemporaneously. This way of approaching personality can yield insights that would be hard to gain by relying only on an underlying trait-only approach. When appropriate data is available, these representations can also be empirically operationalized and tested.

The network approach can yield new insights from existing data

The network approach is flexible in that it can also be used to re-model the kinds of data that dominates the underlying trait paradigm: cross-sectional responses to questionnaire items (Costantini et al., 2015). Specifically, the correlation matrix of the personality characteristic can be represented as a network. Although these analyses offer less added value over traditional techniques compared with analyses based on data with repeated measurements, they offer some possible benefits such as outlining the unique associations among variables (by striving for sparse networks) or identifying the most “central” characteristics.

Sparsity

As said above, personality characteristics can be correlated for various reasons according to the network approach: directly, indirectly, or spuriously due to overlapping influences. Many, if not most, of the correlations are therefore uninformative when the purpose is to disentangle plausible causal associations. As a result, network-approach based representations of personality seek to get rid of extensive correlation patterns. The desired property of networks to contain only relatively few edges is called sparsity. In a sparse network, the edges that arise from confounded or indirect causal relationships are removed. Note that the strive for sparsity contrasts with the underlying trait approach, which *assumes* that the manifestations of the same underlying traits are all correlated—this phenomena is the very evidence for the trait.

One way to make a network sparse is to transform the correlation matrix into a partial correlation matrix. Partial correlations represent residual associations when *all* other correlations in the network have been taken into account. For example, this can be done by using the Gaussian Graphical Model (Lauritzen, 1996). In order to even more effectively shrink weak edges that are likely to reflect nuisance co-variance or otherwise uninteresting information to zero, LASSO-family procedures such as the graphical LASSO (Friedman, Hastie, & Tibshirani, 2008) or adaptive LASSO (Zou, 2006) can be used. For networks that contain temporal information—autoregressive

and cross-lagged associations—multi-level vector autoregressive models can be used (Epskamp et al., 2016).

By focusing on sparse networks and therefore the unique associations among personality characteristics, the network approach can also add value to the analyses of cross-sectional data, over and above methods stemming from the underlying trait approach such as PCA or FA. For example, pockets of characteristics that are associated over and above the variance they share with any other characteristics can be identified and subjected to further conceptual or empirical analysis. Such pockets can suggest facets or nuances for traits.

Centrality and other network properties

Networks can be characterized by various local (pertaining to some areas of the network) or global (pertaining to the whole network) properties, which can also add value to representing and understanding personality processes. For example, nodes can be characterized by their centrality, which quantifies their relative importance: some nodes emit and/or transmit more information than others and removing (or changing) them will likely have consequences that spread throughout the network. For instance, intervention attempts may want to focus on characteristics that are particularly central in the personality systems of individuals because changing these characteristics is especially likely to cause widespread and possibly even lasting change. At the same time, it is exactly the central characteristics that may be particularly hard to change as they are connected to so many other characteristics and are thereby constantly pulled back towards their typical values based on the state of system as a whole.

Note that there is some correspondence between the centrality of nodes in the network approach and component/factor loadings in PCA/FA because characteristics that are more strongly linked with others are also likely to have higher correlations with other characteristics, thereby entailing higher loadings. However, centrality and loadings are conceptually very different properties, because the latter are indicators of how reliably the characteristic *reflects* an underlying trait rather than indicators of direct relationships among them. In the underlying trait approach, changing a characteristic does not change the trait that causes it and therefore changing characteristics with high loadings is expected to be of no more consequence than changing characteristics with low factor loadings. For a parallel situation, imagine that Wellington boots and, in particular, umbrellas are results, and thereby indicators, of rain: alas, putting on the boots or even carrying an umbrella are unlikely to have any influence on whether it rains or not. For a different example, imagine that an underlying trait corresponds to something like human immunodeficiency virus (HIV): treating its symptoms will not remove the virus.

Among different types of centrality indices, degree centrality quantifies the strength of the connections incident to a node (absolute connection strength across all nodes). Nodes with relatively higher degree centralities may have larger direct influences on other nodes or receive more of such influences. On the other hand, closeness centrality is the inverse of the distances of the focal node to all other nodes in the network: if something happens in the network and the influence starts spreading, nodes with high closeness centrality will be affected more quickly than those with low closeness centrality. Another popular index of centrality is the betweenness centrality, which quantifies the extent to which a node “sits” on the shortest connections between any pair of nodes in the network. Removal of a node high in betweenness centrality may be particularly consequential because the shortest paths between many pairs of nodes will increase and therefore the propagation of influences throughout the network slows. For more detailed information on centrality, see Costantini and colleagues (2014).

Among the global network properties, small-worldness (Watts & Strogatz, 1998) should be mentioned. Small-world networks are characterized by clusters of densely connected nodes and links that connect these otherwise distant clusters. If a network demonstrates small-worldness, this means that the changes generally propagate quickly throughout the network because it takes only a

few intermediate nodes to get from any node to any other node. For example, it could be hypothesized that individuals whose personality networks demonstrate the small-world property are more reactive to changes (such as psychotherapy or stress) because influences targeted at specific characteristics are more likely to propagate throughout their personality systems—perhaps even when the target characteristics are not particularly central. The symptoms of psychiatric disorders demonstrate small-worldness which may explain the high rates of co-morbidity among the disorders (Borsboom, Cramer, Schmittmann, Epskamp, & Waldorp, 2011). In the same way, small-worldness might, in principle, explain the often-substantial inter-correlations of personality traits such as those of the FFM (van der Linden et al., 2010) or the general factor of psychopathology (Caspi et al. 2014).

Equilibrium with environment as a cause of stability

According to the underlying trait approach, personality is represented as stable individual differences and can be modeled through data that has been generated by one or a few measurement occasions for each participant. In such a model, associations between personality characteristics tend to be represented (or approximated) as linear. As said above, the network approach conceives of personality as a dynamical system that evolves over time and across situations. When thinking of dynamic systems represented by potentially very large numbers of interactions, one immediately sees the need to drop stationary linear representations. The elements of a dynamic system simply cannot continuously influence each other in an invariant way. For example, if some personality characteristics contribute to one another monotonically, their scores will grow unbounded. Likewise, if the connections are negative and characteristics inhibit each other, they could quickly almost 'vanish'. These are unlikely scenarios for psychological processes, at least pertaining to normal development. To avoid this, the characteristics being modeled have to have either natural boundaries, negative feedback loops within traits, or both positive and negative connections that balance each other.

Cramer and colleagues (2012) describe personality as a system that strives for equilibria with environment: people actively select environments that match some of their characteristics and receive reinforcing feedback from these environments. However, settling for an equilibrium which requires and sustains some characteristics may imply negative feedback loops for other characteristics because the kinds of behavioral, cognitive, and affective activity that personality characteristics reflect is likely to be a limited resource. There is only so much a person can do at a time. For example, when a person is gregarious and often settles for the situations that activate characteristics that tend to go with social activity such as being friendly, co-operative, and sympathetic, he or she almost inevitably has less time for practicing musical instruments or reading poetry—behaviors reflecting Openness characteristics—or working extra hours to get a fast-track promotion, and the person may have less of a chance for cultivating a melancholic mood. Positive connections between some characteristics in the personality system automatically implying negative connections elsewhere prevents the system from “exploding” or “vanishing.” When individuals settle for an equilibrium with their environment that chronically promotes some characteristics and inhibits others, then this can account for the stability of individual differences, among other reasons such as genetic influences on personality “wiring.” Another illustration of the kinds of negative feedback loops that sustain stability is the above-described finding that feeling energetic precedes (physical) activity, whereas the level of the characteristic decreases afterwards, possibly reflecting exhaustion.

Combining the personality feature space, networks and underlying traits

Here, we describe an idea concerning how one can think of personality as a system of characteristics and forces that interact with them. This is a conceptual representation that is (currently) not empirically estimable, but it shows how the network and latent trait approaches can, in principle, be combined into one framework.

As discussed above, people can be represented as person vectors in a multidimensional personality feature space. If we allow the feature space to evolve over time, it can accommodate personality change (short term fluctuations and general developmental trends alike) in addition to individual differences. This is because person vectors that can change their length and direction. For the change to happen, including for person vectors to get to the positions they would then tend to settle in for much of the time (which causes the stability in personality), there have to be some forces that “pull them.” We argue that (a) these forces can also be thought of as vectors in this feature space and (b) the network connections among personality features that make up the space constitute a “bridge” between the forces and person vectors.

The force vectors may represent general environmental niches that pertain to all personality characteristics (pull them all in a particular direction) or specific influences that only pull some characteristics. The force vectors may represent other people: one person’s vector can be a force vector for others, allowing people and their social environments to transact. For example, if people interact with others, their personality automatically becomes part of their own social environment by influencing it and this environment will then automatically reinforce their pre-existing characteristics. Such inter-person transactions may, among other things, cause the clustering of person vectors that can then be identified as factors or principal components. The force vectors can also represent time-invariant genetic influences: person vectors’ baseline positions in the personality space that they tend to gravitate toward. But the force vectors can also include latent factors that pull only *particular* personality characteristics toward particular values (DeYoung, 2015). It was discussed above that components/factors identified by PCA/FA can be represented as vectors in the multidimensional feature space, so considering them as forces that pull person vectors in particular directions is a straightforward step. Of course, when latent factors become better understood, they can be taken apart into separate force vectors or, when the components of latent factors become observed, they can be turned into additional dimensions of the personality space. In principle, there is no limit to the flexibility of this personality representation.

All these forces can simultaneously act on a person vector at any given time: they pull the person vector toward a “target” vector in personality feature space that represents the combined influence of these forces. This target vector of length k ($\mathbf{v}$) can be operationalized as a resultant of the relevant force vectors, which can be appropriately weighed (e.g., the force vector representing genetic influences can be weighed by a heritability estimate). Since updating person vector $\mathbf{y}$ requires a connection matrix $\mathbf{W}$ (because $\mathbf{y}_t = \mathbf{y}_{t-1}\mathbf{W}$), the target vector $\mathbf{v}$ has been turned into a connection matrix. As a simple solution, this can be done by treating $\mathbf{v}$ as the principal eigenvector of the connection matrix (the eigenvector that has a corresponding eigenvalue of 1 whereas all other possible eigenvectors correspond to eigenvalues of 0). So, if $\mathbf{Q}$ denotes a k by k matrix of eigenvectors such that the first column contains the target vector $\mathbf{v}$ and the other columns contain just random vectors orthogonal to $\mathbf{v}$ (e.g., reflecting some stochastic noise), and $\mathbf{\Lambda}$ denotes a diagonal matrix whose diagonal elements are the eigenvalues of the eigenvectors in $\mathbf{Q}$ such that the first eigenvalue is 1 and all others zero, then $\mathbf{W} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}$ (the random vectors in the second to k^{th} column of $\mathbf{Q}$ are really only required to make $\mathbf{Q}$ invertable). In technical terms, given the updating rule of $\mathbf{y}$, $\mathbf{W}$ projects it orthogonally onto the direction of $\mathbf{v}$ (eigenvalues of 1 and 0 is a defining property of an orthogonal projection matrix). In simple terms, such projection makes the connections among personality characteristics a bridge between a person vector’s current state and any force that act on them at the time, whether these forces reflect underlying propensities, influences from other people, or anything else.

Therefore, what we consider to be underlying traits—some directions in personality feature space that pull person vectors toward them—can in principle be accommodated within the network approach, because they can contribute to the connections among characteristics. Furthermore, when more specific, but still unobserved, components of the underlying traits are delineated, they can become additional force vectors in the feature space and thereby separate contributors to the network connections. Or, when the components of the underlying traits become observed, they can

be explicitly represented as dimensions of the feature space.

The strengths of the network approach and the proposed common framework

Richness and flexibility

In comparison to the underlying trait paradigm, the network approach provides a richer and more flexible representation of personality and so does our proposed framework that can combine the underlying trait and network approaches. It is richer in that it combines individual differences with processes happening within individuals and between individuals and their environments. It is more flexible in that it *does not have to be* richer: it can only be used for re-assembling the kinds of data that the underlying trait paradigm typically generates, sometimes potentially offering additional insights. In fact, the network approach appears more generic than the underlying trait one in that it can accommodate the latter, especially when the proposed framework is considered.

An account of complex personality structure

Not only can the network approach account for the coalescence of personality characteristics, but it can also explain why their structural relations are as “messy” as they appear. As discussed above, factor analysis of personality characteristics wrestles with characteristics being associated with multiple ostensible underlying traits and having residual correlations that require postulating increasingly more specific underlying traits underneath the broader ones (the personality hierarchy). According to the network approach, broader factors represent bigger “chunks” of the interconnected network and it is natural that there are also direct edges between the nodes in these somewhat separate chunks—simply because it is a network. Furthermore, as nodes that are closer to each other are more strongly connected (directly or via other nodes) than more distant nodes, this explains the residual correlations *within* the chunks and can account for what appears as the personality hierarchy. If so, there is no inevitable need for postulating different levels of hierarchy such as, for example, the General Factor of Personality (Rushton et al., 2008), the Stability/Plasticity level (DeYoung, 2006), the FFM level (R. R. McCrae & John, 1992), the aspects level (DeYoung, Quilty, & Peterson, 2007), the facets level (Costa & McCrae, 1992) or the nuances level (Möttus, Kandler, et al., in press).

Consistency with how personality is linked with its possible causes and consequences

Constituents of the same traits often demonstrate different links with factors that possibly contribute to personality variance as well as with phenomena that might receive causal contributions from personality (outcomes; Möttus, 2016). For example, Möttus and colleagues (2015) found that facets of the same FFM traits and items of the same facets tended show different correlations with age, which could not be accounted for by difference in the degrees to which the facets reflected the FFM traits or items reflected the facets. To the extent that such findings hold, the only way to reconcile them with the underlying trait perspective is to conclude that it tends to be the item-specific variance that is related to the etiological factors or outcomes of personality, rather than whatever the items share and that thereby could reflect the underlying traits. In many cases, this may make the latent traits redundant for any practical or conceptual purposes. Such findings are, however, consistent with the network perspective because it conceives of personality characteristics as autonomous components with their own etiology and consequences (i.e., incoming and outgoing connections).

Consistency with what we know about the genetics of personality

The specific genetic variants contributing to population variance in personality have remained elusive, although twin studies have consistently shown at least moderate heritability estimates for personality traits, regardless of their breadth or flavor (Turkheimer, Pettersson, &

Horn, 2014; Vukasovic & Bratko, 2015). Collectively, hundreds of thousands of genetic variants can explain up to about 15% of variance in the traits, sometimes much less (Okbay et al., 2016; van den Berg et al., 2016), but few, if any, specific genetic variants that are robustly linked with the traits have been identified to date.

There are numerous explanations for this. Conceiving of personality as a network of characteristics is one of them. At every stage of development, individuals' node scores are partly influenced by their genetic make-up, but these genetic influences may not be aligned with particular groups of characteristics (i.e., chunks of personality network that we conceive of as traits) but with individual characteristics and their interconnections (Cramer et al., 2012). If so, genetic influences could be more fruitfully identified for either these specific characteristics or for the network connections. This explanation is consistent with the findings that random collections of personality test items show heritability estimates similar to “real” trait scores (Johnson et al., 2011) and that even the residual variance of single test items is often heritable (Mõttus, Kandler, et al., in press).

A reason for why heritability of personality traits appears higher in twin studies than family/adoption studies (Vukasovic & Bratko, 2015) or based on molecular genetic data (Okbay et al., 2016; van den Berg et al., 2016) is that the genetic variance is to some extent non-additive (Keller, Coventry, Heath, & Martin, 2005; Vinkhuyzen, van der Sluis, Maes, & Posthuma, 2012). Non-additive genetic variance depends on interactions within the same genetic locus (dominance) or across genetic loci (epistasis). The network perspective on the coalescence of personality characteristics is consistent with the trait scores appearing to reflect non-additive genetic variance in addition to additive influences. In particular, genetic effects on nodes or direct connections between nodes may be more likely to contribute to additive variance, whereas indirect connections may result in what appears as (epistatic) non-additive effects.

Figure 4. Genetic variance in a trait consisting of characteristics x , z , y , and u may appear as partly non-additive. Characteristic x contributes to characteristic y via gene set A additively along with direct connection from characteristics u and z . However, x also contributes to y indirectly via characteristics z and u and characteristic z contributes indirectly via characteristic x . The realizations of these indirect connections simultaneously depends on the combination of multiple sets of genes.

For example, imagine the network depicted in Figure 4. Node y may be directly connected to nodes x , z , and u and these connections are respectively controlled by gene sets A , C , and E (for

simplicity, let's assume that these gene sets are independent). As a result, x , z , and u can additively either contribute or inhibit y . However, y can also be influenced by x indirectly via z and u , whereas these connections are controlled by gene sets B and C, and D and E, respectively. Whether these indirect contributions from x reach y , may depend on the multiplicative effect of the gene sets involved in the indirect pathways and this may result in some epistatic genetic effects. For example, node x may positively contribute to y via direct connection controlled by gene set A but if either gene set B or D cause node x to inhibit nodes z and y , respectively, then y may indirectly inhibit y , depending on the sign of the connections from z and u to y . Likewise, node z can indirectly contribute to y via x , but this depends on whether both gene sets F and A allow for this.

Lack of robust environmental correlates

Although the far-from-unity heritability of personality traits suggests that there are ubiquitous environmental effects on personality, the specific factors that robustly contribute to the environmental variance in traits have remained elusive. This is consistent with the network-based explanation for trait coalescence. The influences specific to single characteristics can produce only temporary effects, because once the exogenous force is removed from the characteristics, their scores may gradually return to the values that are influenced by their connections with other characteristics (i.e., the “wiring” of personality). Effects on the *connections* among nodes can produce more lasting changes as they could entail changes in the equilibrium state. However, as the effects on connections entail changes in how the personality activity (the limited resource discussed above) is distributed among the characteristics, they likely propagate throughout the whole network. This makes it hard to associate external effects with changes in specific groups of characteristics (traits).

The limitations of the network approach

The underlying trait paradigm has been the workhorse of personality psychology for decades, whereas the network approach is a new kid on the block. Because of its novelty, there is very little research that has even attempted to operationalize the network representation of personality. As a result, there is little empirical support for it and even the methodology for gaining the empirical support is underdeveloped. But as studies of within-individual variability become increasingly popular, aided by mobile technology, and the statistical tools for handling the complex dataset develop (e.g., Epskamp et al., 2016), this situation is poised to change. And it could change rapidly.

The network approach lacks the conceptual simplicity and parsimony of the underlying trait approach. Thinking of personality, measuring it, linking it with variables outside the personality domain and communicating the findings would be easy indeed if everything that is important about personality was a few underlying but indirectly measurable trait structures, possible reducible to some simple psychobiological principles with identifiable genetic and environmental etiology. The network representation of personality is so much more complex as the number of autonomous components it must encompass is potentially very large and the number of potential associations among them increase exponentially as the number of components grows. Furthermore, to entangle these associations, researchers would need to see how they change over time and in interaction with environments. This is a daunting task indeed and one can only hope that the apparent complexity will be reducible to some overarching principles that govern the dynamic processes. For example, the conceptual model representing personality as a system characteristics and forces that have mathematically formalized relationships may go some way in this direction.

One of the most compelling arguments for the underlying trait perspective is the cross-cultural replicability of trait structure (see above). This evidence is parsimoniously explained by there being universal underlying structures that cause variance and co-variance of particular types of characteristics. To the extent that the network approach accommodates the underlying trait approach, it can benefit from this parsimonious explanation, but then the rest of what the network

approach allows may become redundant. However, it is plausible that the direct causal connections between personality characteristics, the crux of the network approach, are also recurrent across specific sociocultural contexts. It is plausible that researchers' tend to overestimate the degree to which different cultures vary from these of their own. Perhaps doing something foolish leads to later guilt and sadness everywhere. Future research may address this question.

Conclusion

The chapter compared two explanations for the coalescence of personality characteristics into broader traits. The paradigm of broad underlying traits is well established, but suffers from some limitations such as being inconclusively supported by empirical data and being unable to account for processes within individuals that give rise to observable individual differences. The network approach has only been around for a very short time and has therefore been able to attract a limited amount of empirical research. In theory, it can account for some of the limitations of the underlying broad trait paradigm. For example, the very foundational idea of the approach—that of direct causal connections between the specific characteristic that constitute personality—can account for the complex structure of personality characteristics without any need to bring in *ad hoc* explanations such as numerous new underlying causal influences stepping in at increasingly more specific levels of personality hierarchy. Likewise, the approach, in principle, explain how individual differences emerge from processes within individuals and individual's transactions with environment. We also reviewed a number of other empirical findings that are consistent with the network approach-based explanation to personality. Moreover, some data analytic tools that originate from the network approach can be used to gain novel insights from the kinds of data that typically pertain to the underlying trait approach.

However, despite the two approaches appearing very different at the outset, we provided a more general, mathematically formalized framework that can, in principle, combine the two. To us, it seems that there is no inevitable need to see the two approaches as opposing each other. Combining what has already been established with what may be achieved with novel approaches rather than pitting the established and new approaches against each other may be the best way forward.

Acknowledgements

Authors are grateful to Jeff McCrae for his helpful comments on a draft of the chapter.

References

- Allik, J., Realo, A., & McCrae, R. R. (2013). Universality of the Five-Factor Model of Personality. In P. T. Costa & T. A. Widiger (Eds.), *Personality disorders and the Five Factor Model of personality* (pp. 61–74). Washington, DC: American Psychological Association.
- Ashton, M. C., & Lee, K. (2007). Empirical, theoretical, and practical advantages of the HEXACO model of personality structure. *Personality and Social Psychology Review*, *11*, 150–166. <https://doi.org/10.1177/1088868306294907>
- Bjørnebekk, A., Fjell, A. M., Walhovd, K. B., Grydeland, H., Torgersen, S., & Westlye, L. T. (2013). Neuronal correlates of the five factor model (FFM) of human personality: Multimodal imaging in a large healthy sample. *NeuroImage*, *65*, 194–208. <https://doi.org/10.1016/j.neuroimage.2012.10.009>
- Borsboom, D., Cramer, A. O. J., Schmittmann, V. D., Epskamp, S., & Waldorp, L. J. (2011). The small world of psychopathology. *PLoS ONE*, *6*, e27407. <https://doi.org/10.1371/journal.pone.0027407>
- Cattell, R. B. (1946). Personality structure and measurement. *British Journal of Psychology. General Section*, *36*, 88–103. <https://doi.org/10.1111/j.2044-8295.1946.tb01110.x>

- Caspi, A., Houts, R. M., Belsky, D. W., Goldman-Mellor, S. J., Harrington, H., Israel, S., ... Moffitt, T. E. (2014). The p Factor: One general psychopathology factor in the structure of psychiatric disorders? *Clinical Psychological Science*, 2, 119–137. <https://doi.org/10.1177/2167702613497473>
- Cervone, D. (2004). The architecture of personality. *Psychological Review*, 111, 183–204. <https://doi.org/10.1037/0033-295X.111.1.183>
- Costa, P. T., & McCrae, R. R. (1992). *Revised NEO Personality Inventory (NEO PI-R) and NEO Five-Factor Inventory (NEO-FFI) professional manual*. Odessa, FL: Psychological Assessment Resources.
- Costantini, G., Epskamp, S., Borsboom, D., Perugini, M., Mõttus, R., Waldorp, L. J., & Cramer, A. O. J. (2015). State of the aRt personality research: A tutorial on network analysis of personality data in R. *Journal of Research in Personality*, 54, 13–29. <https://doi.org/10.1016/j.jrp.2014.07.003>
- Cramer, A. O. J., van der Sluis, S., Noordhof, A., Wichers, M., Geschwind, N., Aggen, S. H., ... Borsboom, D. (2012). Dimensions of normal personality as networks in search of equilibrium: You can't like parties if you don't like people. *European Journal of Personality*, 26, 414–431. <https://doi.org/10.1002/per.1866>
- de Moor, M. H. M., Costa, P. T., Terracciano, A., Krueger, R. F., de Geus, E. J. C., Toshiko, T., ... Boomsma, D. I. (2012). Meta-analysis of genome-wide association studies for personality. *Molecular Psychiatry*, 17, 337–349. <https://doi.org/10.1038/mp.2010.128>
- De Raad, B., Barelds, D. P. H., Levert, E., Ostendorf, F., Mlacić, B., Di Blas, L., ... Katigbak, M. S. (2010). Only three factors of personality description are fully replicable across languages: a comparison of 14 trait taxonomies. *Journal of Personality and Social Psychology*, 98, 160–173. <https://doi.org/10.1037/a0017184>
- De Raad, B., Barelds, D. P. H., Timmerman, M. E., De Roover, K., Mlacić, B., & Church, A. T. (2014). Towards a pan-cultural personality structure: Input from 11 psycholexical studies. *European Journal of Personality*, 28, 497–510. <https://doi.org/10.1002/per.1953>
- DeYoung, C. G. (2006). Higher-order factors of the Big Five in a multi-informant sample. *Journal of Personality and Social Psychology*, 91, 1138–1151. <https://doi.org/10.1037/0022-3514.91.6.1138>
- DeYoung, C. G. (2015). Cybernetic Big Five Theory. *Journal of Research in Personality*, 56, 33–58. <https://doi.org/10.1016/j.jrp.2014.07.004>
- DeYoung, C. G., Quilty, L. C., & Peterson, J. B. (2007). Between facets and domains: 10 aspects of the Big Five. *Journal of Personality and Social Psychology*, 93, 880–896. <https://doi.org/10.1037/0022-3514.93.5.880>
- Epskamp, S., Waldorp, L. J., Mõttus, R., & Borsboom, D. (2016). Discovering psychological dynamics: The Gaussian graphical model in cross-sectional and time-series data. *arXiv:1609.04156 [Stat]*. Retrieved from <http://arxiv.org/abs/1609.04156>
- Fleeson, W. (2012). Perspectives on the person: Rapid growth and opportunities for integration. In K. Deaux & M. Snyder (Eds.), *The Oxford handbook of personality and social psychology*. (pp. 33–63). New York, NY: Oxford University Press.
- Fleeson, W., & Jayawickreme, E. (2015). Whole Trait Theory. *Journal of Research in Personality*, 56, 82–92. <https://doi.org/10.1016/j.jrp.2014.10.009>
- Franic, S., Borsboom, D., Dolan, C. V., & Boomsma, D. I. (2014). The Big Five Personality Traits: Psychological Entities or Statistical Constructs? *Behavior Genetics*, 44, 591–604. <https://doi.org/10.1007/s10519-013-9625-7>
- Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9, 432–441. <https://doi.org/10.1093/biostatistics/kxm045>
- Funder, D. C. (1991). Global traits: A Neo-Allportian approach to personality. *Psychological Science*, 2, 31–39. <https://doi.org/10.1111/j.1467-9280.1991.tb00093.x>
- Gignac, G. E., Bates, T. C., & Jang, K. L. (2007). Implications relevant to CFA model misfit,

- reliability, and the five-factor model as measured by the NEO-FFI. *Personality and Individual Differences*, 43, 1051–1062. <https://doi.org/10.1016/j.paid.2007.02.024>
- Gurven, M., von Rueden, C., Massenkoff, M., Kaplan, H., & Lero Vie, M. (2013). How universal is the Big Five? Testing the five-factor model of personality variation among forager-farmers in the Bolivian Amazon. *Journal of Personality and Social Psychology*, 104, 354–370. <https://doi.org/10.1037/a0030841>
- Hopwood, C. J., & Donnellan, M. B. (2010). How should the internal structure of personality inventories be evaluated? *Personality and Social Psychology Review*, 14, 332–346. <https://doi.org/10.1177/1088868310361240>
- Johnson, W., Penke, L., & Spinath, F. M. (2011). Heritability in the era of molecular genetics: Some thoughts for understanding genetic influences on behavioural traits. *European Journal of Personality*, 25, 254–266. <https://doi.org/10.1002/per.836>
- Keller, M. C., Coventry, W. L., Heath, A. C., & Martin, N. G. (2005). Widespread evidence for non-additive genetic variation in Cloninger's and Eysenck's personality dimensions using a twin plus sibling design. *Behavior Genetics*, 35, 707–721. <https://doi.org/10.1007/s10519-005-6041-7>
- Lauritzen, S. L. (1996). *Graphical models*. Oxford: Claredon Press.
- Laverdière, O., Morin, A. J. S., & St-Hilaire, F. (2013). Factor structure and measurement invariance of a short measure of the Big Five personality traits. *Personality and Individual Differences*, 55, 739–743. <https://doi.org/10.1016/j.paid.2013.06.008>
- Lewis, G. J., & Bates, T. C. (2014). How genes influence personality: Evidence from multi-facet twin analyses of the HEXACO dimensions. *Journal of Research in Personality*, 51, 9–17. <https://doi.org/10.1016/j.jrp.2014.04.004>
- Markon, K. E., Krueger, R. F., & Watson, D. (2005). Delineating the structure of normal and abnormal personality: An integrative hierarchical approach. *Journal of Personality and Social Psychology*, 88, 139–157.
- McCrae, R. R. (2009). The Physics and Chemistry of Personality. *Theory & Psychology*, 19, 670–687.
- McCrae, R. R. (2015). A more nuanced view of reliability: Specificity in the trait hierarchy. *Personality and Social Psychology Review*, 19, 97–112. <https://doi.org/10.1177/1088868314541857>
- McCrae, R. R., & Costa, P. T. (2008a). Empirical and theoretical status of the five-factor model of personality traits. In B. Boyle, G. Matthews, & D. Saklofske (Eds.), *The SAGE handbook of personality theory and assessment: Volume 1 — Personality theories and models* (pp. 273–295). London: SAGE.
- McCrae, R. R., & Costa, P. T. (2008b). The five-factor theory of personality. In O. P. John, R. W. Robins, & L. A. Pervin (Eds.), *Handbook of personality: Theory and research* (3rd ed.; pp. 159–181). New York, NY: Guilford Press.
- McCrae, R. R., Costa, P. T., Martin, T. A., Oryol, V. E., Rukavishnikov, A. A., Senin, I. G., ... Urbanek, T. (2004). Consensual validation of personality traits across cultures. *Journal of Research in Personality*, 38, 179–201. [https://doi.org/10.1016/S0092-6566\(03\)00056-4](https://doi.org/10.1016/S0092-6566(03)00056-4)
- McCrae, R. R., Jang, K. L., Livesley, W. J., Riemann, R., & Angleitner, A. (2001). Sources of structure: Genetic, environmental, and artifactual influences on the covariation of personality traits. *Journal of Personality*, 69, 511–535.
- McCrae, R. R., & John, O. P. (1992). An introduction to the five-factor model and its applications. *Journal of Personality*, 60, 175–215. <https://doi.org/10.1111/j.1467-6494.1992.tb00970.x>
- McCrae, R. R., & Terracciano, A. (2005). Universal features of personality traits from the observer's perspective: Data from 50 cultures. *Journal of Personality and Social Psychology*, 88, 547–561.
- McCrae, R. R., Zonderman, A. B., Costa, P. T., Bond, M. H., & Paunonen, S. V. (1996). Evaluating replicability of factors in the Revised NEO Personality Inventory: Confirmatory factor

- analysis versus Procrustes rotation. *Journal of Personality and Social Psychology*, 70, 552–566. <https://doi.org/10.1037/0022-3514.70.3.552>
- Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological Review*, 102, 246–268.
- Mischel, W., & Shoda, Y. (1998). Reconciling processing dynamics and personality dispositions. *Annual Review of Psychology*, 49, 229–258. <https://doi.org/10.1146/annurev.psych.49.1.229>
- Möttus, R. (2016). Towards more rigorous personality trait–outcome research. *European Journal of Personality*, 30, 292–303. <https://doi.org/10.1002/per.2041>
- Möttus, R., Epskamp, S., & Francis, A. (in press). Within- and between individual variability of personality characteristics and physical exercise. *Journal of Research in Personality*.
- Möttus, R., Kandler, C., Bleidorn, W., Riemann, R., & McCrae, R. R. (in press). Personality traits below facets: The consensual validity, longitudinal stability, heritability, and utility of personality nuances. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspp0000100>
- Möttus, R., McCrae, R. R., Allik, J., & Realo, A. (2014). Cross-rater agreement on common and specific variance of personality scales and items. *Journal of Research in Personality*, 52, 47–54. <https://doi.org/10.1016/j.jrp.2014.07.005>
- Möttus, R., Realo, A., Allik, J., Esko, T., Metspalu, A., & Johnson, W. (2015). Within-trait heterogeneity in age group differences in personality domains and facets: Implications for the development and coherence of personality traits. *PLoS ONE*, 10, e0119667. <https://doi.org/10.1371/journal.pone.0119667>
- Okbay, A., Baselmans, B. M. L., De Neve, J.-E., Turley, P., Nivard, M. G., Fontana, M. A., ... Cesarini, D. (2016). Genetic variants associated with subjective well-being, depressive symptoms, and neuroticism identified through genome-wide analyses. *Nature Genetics*, 48, 624–633. <https://doi.org/10.1038/ng.3552>
- Pace, V. L., & Brannick, M. T. (2010). How similar are personality scales of the “same” construct? A meta-analytic investigation. *Personality and Individual Differences*, 49, 669–676. <https://doi.org/10.1016/j.paid.2010.06.014>
- Roberts, B. W., Kuncel, N. R., Shiner, R., Caspi, A., & Goldberg, L. R. (2007). The power of personality: The comparative validity of personality traits, socioeconomic status, and cognitive ability for predicting important life outcomes. *Perspectives on Psychological Science*, 2, 313–345. <https://doi.org/10.1111/j.1745-6916.2007.00047.x>
- Rushton, J. P., Bons, T. A., & Hur, Y.-M. (2008). The genetics and evolution of the general factor of personality. *Journal of Research in Personality*, 42, 1173–1185. <https://doi.org/10.1016/j.jrp.2009.01.005>
- Saucier, G., Thalmayer, A. G., Payne, D. L., Carlson, R., Sanogo, L., Ole-Kotikash, L., ... Zhou, X. (2014). A basic bivariate structure of personality attributes evident across nine languages. *Journal of Personality*, 82, 1–14. <https://doi.org/10.1111/jopy.12028>
- Schmitt, D. P., Allik, J., McCrae, R. R., & Benet-Martinez, V. (2007). The geographic distribution of Big Five personality traits: Patterns and profiles of human self-description across 56 nations. *Journal of Cross-Cultural Psychology*, 38, 173–212.
- Schmitt, D. P., Realo, A., Voracek, M., & Allik, J. (2008). Why can't a man be more like a woman? Sex differences in big five personality traits across 55 cultures. *Journal of Personality and Social Psychology*, 94, 168–182.
- Schmittmann, V. D., Cramer, A. O. J., Waldorp, L. J., Epskamp, S., Kievit, R. A., & Borsboom, D. (2013). Deconstructing the construct: A network perspective on psychological phenomena. *New Ideas in Psychology*, 31, 43–53. <https://doi.org/10.1016/j.newideapsych.2011.02.007>
- Specht, J., Bleidorn, W., Denissen, J. J. A., Hennecke, M., Hutteman, R., Kandler, C., ... Zimmermann, J. (2014). What drives adult personality development? A comparison of theoretical perspectives and empirical evidence. *European Journal of Personality*, 28, 216–

230. <https://doi.org/10.1002/per.1966>
- Thalmayer, A. G., & Saucier, G. (2014). The Questionnaire Big Six in 26 nations: Developing cross-culturally applicable Big Six, Big Five and Big Two inventories. *European Journal of Personality*, 28, 482–496. <https://doi.org/10.1002/per.1969>
- Turkheimer, E. (2000). Three laws of behavior genetics and what they mean. *Current Directions in Psychological Science*, 9, 160–164. <https://doi.org/10.1111/1467-8721.00084>
- Turkheimer, E., Pettersson, E., & Horn, E. E. (2014). A phenotypic null hypothesis for the genetics of personality. *Annual Review of Psychology*, 65, 515–540. <https://doi.org/10.1146/annurev-psych-113011-143752>
- Valchev, V. H., van de Vijver, F. J. R., Meiring, D., Nel, J. A., Hill, C., Laher, S., & Adams, B. G. (2014). Beyond agreeableness: Social-relational personality concepts from an indigenous and cross-cultural perspective. *Journal of Research in Personality*, 48, 17–32. <https://doi.org/10.1016/j.jrp.2013.10.003>
- van den Berg, S. M., de Moor, M. H. M., Verweij, K. J. H., Krueger, R. F., Luciano, M., Arias Vasquez, A., ... Boomsma, D. I. (2016). Meta-analysis of genome-wide association studies for extraversion: Findings from the Genetics of Personality Consortium. *Behavior Genetics*, 46, 170–182. <https://doi.org/10.1007/s10519-015-9735-5>
- van der Linden, D., te Nijenhuis, J., & Bakker, A. B. (2010). The General Factor of Personality: A meta-analysis of Big Five intercorrelations and a criterion-related validity study. *Journal of Research in Personality*, 44, 315–327. <https://doi.org/10.1016/j.jrp.2010.03.003>
- Vinkhuyzen, A. A. E., van der Sluis, S., Maes, H. H. M., & Posthuma, D. (2012). Reconsidering the heritability of intelligence in adulthood: Taking assortative mating and cultural transmission into account. *Behavior Genetics*, 42, 187–198. <https://doi.org/10.1007/s10519-011-9507-9>
- Vukasovic, T., & Bratko, D. (2015). Heritability of personality: A meta-analysis of behavior genetic studies. *Psychological Bulletin*, 141, 769–785. <https://doi.org/10.1037/bul0000017>
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, 393, 440–442. <https://doi.org/10.1038/30918>
- Wood, D., Gardner, M. H., & Harms, P. D. (2015). How functionalist and process approaches to behavior can explain trait covariation. *Psychological Review*, 122, 84–111.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101, 1418–1429.