Bottom Up Construction of a Personality Taxonomy

David M. Condon*, University of Oregon, USA

Dustin Wood, University of Alabama, USA

René Mõttus, University of Edinburgh, UK and University of Tartu, Estonia

Tom Booth, University of Edinburgh, UK

Giulio Costantini, University of Milan-Bicocca, Italy

Samuel Greiff, University of Luxembourg, Luxembourg

Wendy Johnson, University of Edinburgh, UK

Aaron Lukaszewski, California State University, Fullerton, CA, USA

Aja Murray, University of Edinburgh, UK

William Revelle, Northwestern University, USA

Aidan G.C. Wright, University of Pittsburgh, USA

Matthias Ziegler, Humboldt Universität zu Berlin, Germany

Johannes Zimmermann, University of Kassel, Germany

* Correspondence to David Condon, 1227 University St, Eugene, OR, 97403 USA dcondon@uoregon.edu

This manuscript is based on an Expert Meeting jointly supported by European Association of Personality Psychology and European Association of Psychological Assessment and held from 6th to 8th September 2018 in Edinburgh, Scotland (<https://osf.io/fn5pw>). The authors are grateful to Mitja Back, Anna Baumert, Jaime Derringer, Sacha Epskamp, Ryne Sherman, David Stillwell, and Tal Yarkoni for their contributions to the Expert Meeting. Not all authors agree with all arguments put forward in this paper.

Abstract

In pursuit of a more systematic and comprehensive framework for personality assessment, we introduce procedures for assessing personality traits at the lowest level: nuances. We argue that constructing a personality taxonomy from the bottom up addresses some of the limitations of extant top-down assessment frameworks (e.g., the Big Five), including the opportunity to resolve confusion about the breadth and scope of traits at different levels of organization, evaluate unique and reliable trait variance at the item level, and clarify jingle/jangle issues in personality assessment. With a focus on applications in survey methodology and transparent documentation, our procedures contain six steps: (1) identification of a highly inclusive pool of candidate items, (2) programmatic evaluation and documentation of item characteristics, (3) test-retest analyses of items with adequate qualitative and quantitative properties, (4) analysis of cross-ratings from multiple raters for items with adequate retest reliability, (5) aggregation of ratings across diverse samples to evaluate generalizability across populations, (6) evaluations of predictive utility in various contexts. We hope these recommendations are the first step in a collaborative effort to identify a comprehensive pool of personality nuances at the lowest level, enabling subsequent construction of a robust hierarchy -- from the bottom up.

Keywords: personality assessment, taxonomy, nuances, hierarchy, survey methodology

Bottom-Up Construction of a Personality Taxonomy

Introduction

In this article, we discuss procedures for assessing a comprehensive set of specific personality traits that can be interpreted on their own, aggregated into broader constructs, or both. A detailed rationale for such an approach to personality assessment is outlined in a companion piece (Möttus, Wood, Condon et al., in press), published in the joint special issue, *New Approaches for Conceptualizing and Assessing Personality*, of the *European Journal of Personality* and the *European Journal of Psychological Assessment*. In short, a bottom-up approach based on ‘lower-level’ traits would provide a more systematic and comprehensive framework for personality assessment than currently exists while acknowledging the hierarchical organization of traits: personality variance can be studied at multiple levels of abstraction, only one of which is occupied by the commonly measured “Big Few” traits (e.g., the Big Five or HEXACO domains). To achieve this, the approach should focus on the properties of individual items as markers of specific personality traits, rather than on the properties of multi-item scales, as has been customary in personality assessment. Assessing personality as a hierarchical phenomenon facilitates achieving each of the major goals of personality research: description, prediction, and explanation (Möttus, Wood, Condon et al., in press). For example, outcomes can be described by, explained in relation to, and predicted from a few broad traits, whereas for more precise descriptions, predictions, and explanations, numerous narrower traits can be used.

Defining traits and their scope

Defining traits has been an ongoing challenge in personality science for at least several decades (consider, for examples, Allport [1966], Funder [1991], and Tellegen [1991]). Here, we define traits as dimensions of any kind of relatively stable psychological (affective, cognitive, motivational and behavioral) differences among people, independent of their content, breadth, or expected importance (Baumert et al., 2017; Wood, Gardner, & Harms, 2015). This definition allows, in addition to standard approaches, for inclusion of many relatively stable human attributes typically omitted from omnibus personality measures such as individual differences in cognitive abilities, information processing, worldviews, values, goals, strivings, attachment styles, hobbies, habits, behavioral tics, and more. As with existing personality assessment frameworks, it excludes individual differences that are not psychological and/or stable. In practice, of course, researchers may focus their attention on more narrowly defined groups of traits; here, we do not want to impose our criteria for what is and is not in the domain of personality, but simply to emphasize that it can be defined far more broadly than the scope of most current assessment tools (Bouchard, 2016).

Hierarchical framing

Given the number of widely-used hierarchical frameworks for personality assessment, many researchers may take it for granted that traits are hierarchically organized. In practice, the delineation of trait hierarchies has largely been a by-product of the development and refinement of personality assessment

tools that were targeted at different levels of granularity -- that is, the hierarchical representation of traits has mostly been an assessment issue. The procedures proposed herein are fundamentally rooted in the notion that personality traits *themselves* can be organized hierarchically, so that the trait hierarchy has both theoretical and methodological relevance (e.g., DeYoung, 2015; McCrae & Sutin, 2018; Möttus, Wood, Condon et al., in press). Most existing research, however, has only concentrated on the higher levels of the personality hierarchy, most notably the Big Few. In contrast, our bottom-up approach calls for an explicit focus on lower levels as both the building blocks of higher levels and unique sources of descriptive, predictive, and explanatory information (Möttus, Wood, Condon et al., in press).

In taxonomic terminology, two or more objects are said to share the same “level” when they have approximately the same rank or importance (Dawkins, 1976). Generally, personality traits are tacitly ranked based on their breadth, and this has been rarely articulated (for exceptions, see Hogan & Roberts [1996], Ones & Viswesvaran [1996], and Kretzschmar et al. [2018]). This means that traits at the higher levels of the hierarchy are typically broad and subsume multiple lower-level traits. In psychometrics, the breadth of measured trait constructs is often discussed in terms of dimensionality, providing an operational basis for differentiating the hierarchical levels. In theory, the narrow traits at the lowest levels of the hierarchy should be unidimensional and therefore refer to singular characteristics. Traits at the highest levels should *not* be unidimensional (Ziegler & Bäckström, 2016). In practice, the assessment tools designed to capture different levels of the trait hierarchy mix broader and narrower traits within levels rather than clearly separating them based on breadth (Hogan & Roberts, 1996); at the highest levels, unidimensionality is uncommon (Hopwood & Donnellan, 2010).

We refer to the highest personality hierarchy level as the Big Few, which include the Big Five (Goldberg, 1990), the Five-Factor Model (McCrae & John, 1992), and the HEXACO domains (Ashton & Lee, 2020). Higher levels have been proposed (Digman, 1997; DeYoung, 2006; Eysenck, 1994), though these are more often referred to in theoretical discourse than for assessment. As shown in Figure 1, existing hierarchies proceed down from the Domains to the Aspects and then the Facets. Criteria for defining each of these levels have not (yet) been specified, particularly below the level of Domains, and there is little research on the extents to which the traits currently delineated for each of the levels a) exhaust the traits that are potentially identifiable for that level and b) overlap among themselves (Schwaba et al., 2020).

Figure 1: The combination of content included in four hierarchical personality frameworks. These are the Five-Factor Model (with five domains and 30 facets; Costa & McCrae, 1995), the HEXACO (6 domains and 24 facets; Ashton, Lee, et al., 2004), the Big Five Aspect Scales (10 aspects; DeYoung, Quilty & Peterson, 2007), and the BFI-2 (15 facets; Soto & John, 2017). Duplicated label names among the Facets have been omitted.

Figure 1 illustrates three important points that relate to the case for hierarchy development from the bottom up. First, the specification of content across and within levels is plagued by the jingle/jangle fallacies, with different trait labels often referring to overlapping content and the same labels referring to different content (Block, 1995; Larsen & Bong, 2016). The most obvious jingle examples stem from the use of identical terms across levels: several constructs are both Aspects and Facets (Assertiveness, Order, Volatility, Compassion), and one is both a Domain and an Aspect (Openness). Examples of jangle include various labels used for essentially the same facets (e.g., Liveliness, Energy-Level and Activity Level). Development of a new hierarchy from the bottom up would expose existing cases of jingle and jangle, and provide a way of addressing them: traits with different lower-level constituent parts would be given different labels, whereas if a newly proposed trait has the same constituents as an existing one, it need not be advanced at all. Evaluations of the extent to which two traits have overlapping vs. distinct content could then be empirical rather than based on researchers' intuitions (Schwaba et al., 2020; Siegling, Petrides, & Martskvishvili, 2015).

Second, it is unclear whether the scope of traits at a given level is sufficiently exhaustive of the level (e.g., whether the 30 facets of the NEO Personality Inventory comprehensively sample all relevant facet-level traits). Several researchers have suggested that many traits lie outside the Big Few hierarchies (Paunonen, 2000; DeVries et al., 2009); these claims presumably include content at the Domain level as well as below it. This exposes a problem, however, as most of the proposed hierarchies nest lower-level traits strictly within their Big Few domains by simply decomposing the content at the Domain level (Condon, 2018). Critics have pointed out that this top-down approach precludes the prospect of subsequent expansion of the breadth of content, even at lower levels (Condon, 2018; Markon, 2009). This means that efforts to increase the breadth of personality models necessarily require *de novo* development from the bottom up.

The third and final point is based on the increasing evidence that the items underlying all of the hierarchical levels depicted in Figure 1 contain reliable and valid trait information beyond what they share with other items (Möttus, Wood, Condon et al., in press). That is, many items represent unique personality traits that are narrower than facets, aspects, and domains; these narrow traits have been called *nuances* (McCrae, 2015; McCrae & Möttus, 2019; Möttus, Kandler et al., 2017).

We formally define a *nuance* as the lowest level at which patterns of responses to items continue to have reliable specific variance. In many cases, this will effectively equate items with nuances, as most items have reliable specific variance (Möttus, 2016; Möttus et al., 2017). However, definitionally, the nuance level is broader than the level of items. There will be certain sets of items that contain no reliably distinct information from one another; these can be considered alternative indicators of the same nuance, and are likely to be understood by participants as being synonymous.¹ If items are found to have reliable specific variance from others items in the set, then the item definitionally captures, or is a marker for, a distinct nuance. Nuances have formal trait properties such as stability over extended periods of time and detectability with different assessment methods like self- and informant-reports (Möttus et al., 2014; Möttus, Sinick et al., 2019), besides displaying unique heritability (Möttus, Kandler et al., 2017; Möttus, Sinick et al., 2019), predictive validity (Elleman et al., 2020; Möttus, Bates et al., 2017; Revelle, Dworak, & Condon, 2020; Seeboth & Möttus, 2018; Wessels, Zimmermann, & Leising, 2020) and associations with other demographic variables (Achar-Amankwaa, Oлару, & Schroeders, 2020; Elleman et al., in press; Möttus & Rozgonjuk, 2019). Thus, nuances comprise the lowest identifiable level of the personality trait hierarchy and are the constituents that make up all higher-order traits.

Multiple ways of organizing traits

It may be unreasonable to expect that any single representation could adequately convey the high-dimensional complexity of personality (Yarkoni, 2010). Goldberg has argued "that trait descriptors are not neatly clustered in multivariate space" (p. 27, Goldberg, 1992) but are rather distributed like stars in the night sky, suggesting that the most basic units may not be optimally organized with neat and symmetric hierarchical nesting (Markon, 2009). Viable alternatives include other forms of hierarchy and non-hierarchical structures. Examples of non-nested hierarchies include those by Saucier and Iurino (2019) or Condon (2018), both of which contain independent levels of non-nested factors. Non-hierarchical examples include clusters and lists (Loehlin & Goldberg, 2014), and network structures (Cramer et al., 2012).²

However, nested hierarchies do have a particular strength: the *redundancy reduction* advantage (Dawkins, 1976). Data reduction techniques such as factor, component, and cluster analyses have been embraced for decades because they enable the reduction of vast diversity in psychological individual differences down

¹ As an example, Wood, Lowman, et al. (unpublished) showed how item-level reliability adjustments can be used to indicate that self-ratings of the items *afraid*, *scared*, and *frightened* within the PANAS-X instrument (Watson & Clark, 1994) both had virtually no reliable specific variance and were rated by most participants in a separate rater sample as "having essentially the same meaning" on a semantic similarity scale. Thus, these may not represent meaningfully distinct nuances.

² Networks sometimes contain hierarchical relations in the sense that one or more objects may be subordinate to other objects (Dawkins, 1976), as with directed acyclic graphs (Rohrer, 2018) or explicitly hierarchical networks (Epskamp, Rhemtulla, & Borsboom, 2017).

to (more) manageable numbers. These same methods have been used iteratively to develop omnibus top-down trait hierarchies (DeYoung, Quilty, & Peterson, 2007) as well as for evaluations of specific domains (e.g., Roberts et al., 2005; Olino et al., 2005; Crowe, Lynam, & Miller, 2018). These iterations begin with reduction to a small number of traits (e.g., the Big Few), followed by subsequent factoring of the constituent variables in each (e.g., to identify the aspects). An alternative procedure for atheoretically delineating the appropriate levels is Goldberg's 'bass-ackward' approach (2006), in which a progressive number of non-nested factors are extracted and then evaluated for fit, coherence, and replicability.

Our procedures for developing a bottom-up framework do not assume any specific organizational structure because there are many viable options; for example, a strictly nested hierarchical representation of traits may be implausible because of many-to-many links between traits at different hierarchical levels. Instead, our goal is to set forth a comprehensive set of the fundamental constituent traits that can be organized into various structures -- broader and narrower, nested and non-nested -- just as LEGO™ blocks can be organized into a wide range of structures. But, for this to be possible, we first need a diverse set of strong and functional 'blocks'.

A New Approach

The best path forward, in our view, thus depends upon more detailed specification of the narrow trait space and subsequent construction from the bottom upwards. There is some precedent for this approach, such as the collaborative effort to classify dimensional features of psychopathology ("HiTOP"; Krueger et al., 2018) and efforts supported by the National Institutes of Health to create a bottom-up framework of patient-reported health outcome measures ("PROMIS"; Cella et al., 2019), but the circumstances in personality assessment are somewhat unique.³ The challenges include the need to cover content of a very broad scope while also specifying the lowest level features in unprecedented detail, and to gather multiple large and reasonably representative datasets in order to make defensibly generalizable claims. These challenges can be addressed, though they require some methodological assumptions, as discussed below.

Methodological Considerations

Survey methodology

Of course, personality involves more than an individual's survey response profile (Möttus, Wood, Condon et al., in press) and there are many approaches in personality science that do not make use of surveys at all. We fully agree with calls to study personality and its processes with other methodologies such as direct behavioural observations (e.g., Furr, 2009; Back, 2020; Möttus, Wood, Condon et al., in press; Rauthmann, 2020). That said, a large proportion of research in personality science, along with

³ Many psychologists, including personality scientists working outside the domain of personality structure, are often surprised to discover that the Big Few were identified in the absence of strong consensus regarding their underlying content, or building blocks. Despite the hallowed tales of graduate students pouring over dictionaries to compile lists with person-descriptors, the truth is that the Big Few models -- which varied in the extent to which they were empirically informed -- all were heavily influenced by pragmatic theory-based decisions by a handful of domain experts. Examples include the clustering of traits into paragraph descriptions (Cattell, 1946), and decisions such as whether to include mood-related (e.g., *angry*) or highly evaluative terms (e.g., *impressive*; Goldberg, 1993). Saucier (1997) has shown that such decisions had non-trivial effects on the resulting structures.

many of the other psychology subdisciplines, is conducted through surveys, particularly self-reports but also informant-reports. Arguably, many of the best established personality science findings are based on surveys. Besides, best practices in survey methods are well-established and many in the field have developed expertise in the analysis of survey data. For these reasons, our suggested procedures for developing a bottom-up hierarchy begin with the assumption that data collection will be based on survey methodology.

Trait descriptive adjectives and type nouns vs phrased items

The history of personality trait hierarchies is closely tied to the lexical hypothesis (Goldberg, 1993; John & Srivastava, 1999), as several research programs connected the early psycho-lexical efforts of Allport and Odbert (1936) to the eventual production of Big Few models (Goldberg, 1992; Ashton et al., 2004). These models were identified based on the use of trait-descriptive checklists, where data were collected by instructing respondents to rate themselves and/or close others using terms found in the dictionary. The logic of using single-word trait-descriptors rests upon the essential idea that the finite-but-unstructured universe of these terms reflects the universe of important individual differences (Saucier & Goldberg, 1996; Wood, 2015). In fact, they have been described as “a natural language taxonomy of personality terms” (p. 86, John, Goldberg, Angleitner, 1984), suggesting that the single-word descriptors constitute the lowest-order of the trait hierarchy.

Starting with single-word descriptors is problematic for several reasons. The first of these stems from a lack of consistency with respect to the breadth of these terms, highlighted by single-word descriptors that also serve as labels for broad traits (e.g., agreeable, open). In fact, some of the most frequently used single-word descriptors in everyday language are often rated as being more abstract (Leising et al., 2014).

Further, as Block (1995) has noted, the history of the trait hierarchy was not based, in practice, upon evaluations of the exhaustive sample of the full universe of trait-descriptors. The prospect of collecting data on all possible single-word trait-descriptors was (and continues to be) a formidable challenge.⁴ Block (1995) methodically questioned the “prestructuring” of personality variables at each major milestone in the evolution of the Big Five (Booth & Murray, 2018), and there is evidence that the output of factor analysis can be meaningfully altered by variable selection procedures (for discussion of the specific ways, see Saucier [1997]). The consequences of decisions made to winnow the pool of terms down to more manageable sets remain unclear. Certainly, they created an opportunity for second-guessing, and undermine claims that the trait-descriptor approach is truly comprehensive.

An exclusive focus on single-word trait descriptors may also be problematic in that many single terms are not consistently understood and some traits are not well-captured by single terms. Many trait-descriptors are broadly unfamiliar or may be interpreted in different ways due to having multiple definitions, such that their meaning can only be understood consistently with longer phrases. Consider, for example, the trait-descriptive adjective “arbitrative” (Norman, 1963) vs. “Am often asked by friends and family to help resolve fights” (Goldberg et al., 2006). Similarly, some trait-descriptive terms may be difficult to rate in

⁴ Early work (Allport & Odbert, 1936) suggested the existence of nearly 18,000 terms in English. More recent research suggests that an exhaustive list restricted to widely recognizable terms contains less than 3,000 terms, and perhaps as few as 1,500 (Norman, 1967; Ashton, Lee & Goldberg, 2004).

the absence of situational information. Compare, for example, the trait-descriptive adjective “fearful” against the many items in the International Personality Item Pool using the words “fearful” or “afraid.” The items are more contextualized, as some relate to fears about specific stimuli, the absence of fear (especially when typical), and even the enjoyment of fearful situations (i.e., risk-seeking behaviors). Many individuals are likely to endorse different levels of fear across situations, and the phrased item circumvents this concern. Of course, phrased items are also subject to varying interpretations, and this is a point of concern addressed among the procedures described below.

In any case, more research is needed regarding the comparative advantages of single-word trait descriptors vs. phrased items, regarding their relative breadths, subjective clarity to raters, reliability and incremental predictive value. Meanwhile, restricting taxonomic research to single-word descriptors is almost certainly unwarranted.

Procedures

Our overarching goal for what follows is to detail a set of procedures for developing a more comprehensive bottom-up set of items to ultimately represent the domain of personality at all levels of the trait hierarchy. As we have noted, the creation of any personality ‘structure’ requires theoretical and practical consideration of what content should be included and prioritized. Consequently, we try to make some of our assumptions clear here. Specifically, items should be preferred for inclusion in the set when they have *reference to clear psychological attributes that respondents can understand, high reliability, high agreement between self-reports and reports by knowledgeable informants (unless there are articulated reasons for why different rating sources should not converge for particular items)*, and they *increase the total information within the set* (i.e., are not redundant with other items; have meaningful reliable item-specific variance). That is, the focus is on the properties of items as markers of the lowest measurable level of the personality trait hierarchy, nuances, rather than on the properties of multi-item scales, as has been customary in personality psychology.

The procedures listed below are ordered but we intend them to be carried out iteratively and/or simultaneously as assessment content should be expanded and reduced on an ongoing basis and as more empirical data become available.

Step 1. Identify a comprehensive sample of stimuli

The initial goal of the approach that we propose here is to identify an overly broad pool of items from which to build. Many personality psychologists will consider the International Personality Item Pool (Goldberg, 1999) as an ideal starting point. We agree, though it is also important to set forth a number of qualifications for the sake of subsequent iterations, as is done below with respect to comprehensiveness and domain breadth.

Comprehensive -- not exhaustive

It seems unlikely that there can ever be a finite and ultimate model of personality nuances, as there are (at least) thousands of them. The impossibility of ever fully delineating the domain of nuances becomes even clearer if we allow the list to include aspects such as capacities, preferences, and behavioral tendencies

associated with emerging technologies (e.g., *digital traces*) and “sub-genre” niches of behavior (Fitzpatrick, 2014). But, we take a pragmatic approach.

The goal is not to identify a fully exhaustive set of items, but rather to gather a pool of them that can be defended as providing a *sufficiently comprehensive sample of the item universe*. This is feasible as long as researchers recognize that the pool will evolve over time, with new content being added as gaps are identified. The challenge lies in identifying these gaps, for it requires the addition of content not previously considered with a top-down focus that starts with the Big Few. One method, addressed in subsequent steps, calls for the addition of non-redundant items on the basis of empirical evidence, even though this is resource-intensive and time consuming.

A useful alternative to the empirical approach involves the solicitation of input from subject matter experts using procedures similar to those used previously for projects in personality assessment (Block, 1961) and adjacent fields (Krueger et al., 2018; Cella et al., 2019). Block’s (1961) method, for example, was an iterative process involving discussions with a research team to propose content that was insufficiently reflected within each iteration of the instrument. This process can introduce bias to the content -- for instance, some have noted that the Block’s (1961) California Adult Q-Set measure may overemphasize psychoanalytic content given his particular interests (Saucier, 2020) -- but the harm of idiosyncratic biases can be limited if expert input is used to expand rather than constrain the item pool. The involvement of experts from disparate fields may also prove beneficial, including those from fields outside personality with the potential to offer unique perspectives (e.g., social workers, religious leaders, management experts).

We recommend that both expert input and empirical evaluations are used for bottom-up development. Beginning with the existing IPIP, for example, expert input should be solicited in a coordinated way to consider whether the most prominent topics in psychological individual differences research are covered. Where warranted, new content should be added, either from existing scales and item pools or through the creation of new content. Over time, the benefits of additional expert input and empirical evaluation will decrease as the evidence for comprehensiveness increases, but it is important to keep the item pool open to revision. All too frequently, researchers regard assessment content as static or immutable, even when they have reason to suspect that traits which probably relate to the phenomena of interest are not indicated by items within the pool. We suggest that in such cases, there is little reason not to expand the item pool with items that indicate these traits as directly as possible, at least when practically feasible (it may not be in applied assessment contexts). More generally, a benefit to the bottom-up approach is the opportunity for ongoing revision and content expansion.

Inclusion of multiple types of personality content

The dimensions assessed with Big Few measures are historically rooted in psycho-lexical research. However, the extent to which *variance* in psychological and behavioral tendencies is captured by this content remains unclear, and several other lines of research offer promising opportunities for potential expansion. A broad survey of these lines of research would include interests, abilities, motives and goals, strengths, values, strivings, preferences, attitudes, attachment styles, and possibly even more idiosyncratic content such as aspects of identity (Bouchard, 2016).

Most of this is already well-captured by existing survey content, and much of this has already been aggregated -- alongside the IPIP and its supplements (Goldberg & Saucier, 2016) -- from disparate sources into a pool of more than 15,000 items in the Database of Individual Differences Survey Tools (Condon, 2019). Aggregation is only the first step, however, for the disparate nature of the content types complicates the matter of its organization. Cognitive abilities, for example, are typically assessed with performance items rather than endorsement-type survey items. Interests items, by contrast, are often framed in terms of behavioral frequency. Given these distinctions, a key step will be addressing the method effects introduced by assessing different types of content using different response scales.

Step 2. Programmatic evaluation and documentation of item characteristics

Personality test constructors have not historically pursued thorough qualitative analyses of items, but we argue that this should play an important role along with quantitative evaluations of item properties. At the most basic level, this work invokes data collection strategies aimed at evaluating the face validity of individual items; these strategies include cognitive interviewing, focus groups, and/or the use of surveys with open-ended response formats to capture respondents' perceptions about the items (Ziegler et al., 2015). These procedures may also identify items that seem virtually redundant in meaning. The researcher's initial impressions can be made more systematic through collecting small numbers of semantic similarity ratings on the items to see if raters are able to perceive some meaningful distinctions between the most redundant-looking items - essentially, judging whether some participants could meaningfully endorse one but not the other or vice-versa (e.g., Tracy & Robins, 2007; Block, Weiss, & Thorne, 1979). Items that are judged to be highly redundant to others in the set should be revised or the least straightforward one dropped. See MacCann et al. (2009) for examples of such qualitative evaluations.

Quantitative procedures can also be used to evaluate item qualities. For example, surveys can be used to rate items along psychometrically relevant dimensions, such as social desirability and familiarity of the phenomenon being assessed, translation readiness (i.e., extent to which the content is free of idioms, or retains its meaning when translated and back-translated; e.g., Wood et al., 2018; Greiff & Iliescu, 2017), required literacy/reading levels, clarity, brevity, and readiness for usage in both self- and informant-report formats (Funder & Dobroth, 1987; Hampson et al., 1987; Leising et al., 2014; Chandler, 2018; Dumas et al., 2002). Generally, properties such as these only need to be rated by a fairly small number of raters to obtain sufficiently reliable estimates. For instance, Wood (2015) reported that the average inter-rater correlation for trait properties, such as *breadth* or *observability* was about .20, meaning that only about 16 raters should be necessary to achieve a sufficiently high average-rater reliability (e.g., intra-class correlation of about .80).⁵ Following these procedures, items deemed inadequate should be set aside, and a public-record of the data collected about these deprecated items should be retained for future reference in case someone later wants to reinstate it. In addition to use for refining the item pool, these types of ratings can address other substantive research questions (Funder & Dobroth, 1987; Block, 1961) such as cultural differences in what people tend to notice about others, and individuals' levels of and interest in

⁵ Estimated by treating raters as items i in the Spearman-Brown formula:

$$(N_i \times r_{XiXi'}) / (1 + (N_i - 1) \times r_{XiXi'}) = r_{XX'}$$

where $r_{XX'}$ indicates the expected alternative form correlation of the mean rating X with the mean X' formed from a new group of raters of equal size sampled in the same manner.

developing understanding of how one tends to be seen by others (Henry & Möttus, 2020; Wood & Wortman, 2012).

It is also important to consider meta-attributes of the assessment content. A crucial concern is public-domain status. Though many items, including those in the IPIP, are clearly available in the public domain, many existing items have unclear licensing status. For items that are part of proprietary measures with well-documented validity, requests for releases of licensing can be pursued or new/revised items can be written, though the latter option introduces the need for additional validation. While the development of new items is often complicated by the need to respect intellectual property, it can be a beneficial opportunity to improve qualitative item characteristics and/or harmonize the framing of item content. In some cases, this could mean redefining existing constructs completely. There is considerable precedence for such efforts with the International Personality Item Pool (Goldberg, 1999; Goldberg et al., 2006).

Step 3. Conduct retest analyses of items with adequate qualitative and quantitative properties

Test-retest correlations over short spans are particularly good indicators of item quality: for an item to provide reliable and useful information, raters first have to answer it consistently in the short-run -- that is, they have to be able to agree with themselves on the content of the item. The retest interval can be a couple of months (Watson, 2004), a couple of weeks (Möttus, Sinick et al., 2019; Soto & John, 2017), a couple of days (Wood et al., 2010), or even a couple of minutes (Lowman et al., 2018; Wood et al., 2018). What makes these estimates so valuable is that they are particularly good predictors of many standard indicators of item validity simultaneously, such as self-other agreement correlations and stability correlations over longer time periods (McCrae, Kurtz, Yamagata, Terracciano, 2011; Henry & Möttus, 2020), while also being estimable for single items. Item-level retest correlations are also higher than many personality researchers expect. For instance, Wood, Harms, and colleagues (unpublished) found the average BFI-2 item had retest correlations of .58 over a two-month interval, and higher values over shorter intervals.

We suggest that items with retest correlations below $r = .50$ over a two-week interval may be considered for immediate exclusion, while those with retest correlation below $r = .60$ should only be kept cautiously. Low levels of short-term retest correlations will often indicate items that are hard for participants to interpret or relate to; for instance the BFI-2 item “feels little sympathy for others” showed retest correlations below .40 across two months and even 15-minute retest intervals (Wood, Harms, et al., unpublished), suggesting that participants find this item hard to interpret, and/or judge its applicability to them. In a carefully selected item pool, average two-week retest reliability could realistically be approximately $r = .70$ (Möttus, Wood, Condon et al., 2020).

Step 4. Collect cross-ratings of items with adequate retest reliability from multiple raters

For a characteristic to capture a valid and useful piece of information, it is also highly desirable that multiple ways of measuring it yield convergent information. As a result, we argue that cross-rater agreement is one of the most straightforward indicators of the degrees to which items represent valid and method-independent information as opposed to reflecting unique method-specific influences (McCrae & Möttus, 2019). Of course, information only known privately to the person providing self-reports can be uniquely valuable for some subsets of traits, as can be information only externally visible to informants (Vazire, 2010). But, high agreement across these sources is suggestive of particularly valid and therefore

useful information (Möttus, Wood, Condon et al., in press; Kendrick & Funder, 1988). It is unlikely that any other validity criteria is as universally applicable across items. This hypothesis is consistent with the finding that items with the highest retest reliability also tend to have the highest cross-rater agreement, on average (e.g., McCrae et al., 2011). As a result, items with relatively high correlations across multiple ratings should generally be preferred to those with lower cross-rater correlations, at least for some purposes.

Step 5. Aggregate ratings of retained items across diverse participant samples to evaluate their utility across populations

The collection of self-report data is a routine part of most personality research protocols, particularly with cross-sectional data. This step is focused on the aggregation of data sets with at least overlapping item pools (it is not required that all items overlap; Möttus, Wood, Condon et al., 2020). While the inclusion of many different data sets will increase the diversity of participant samples, the inclusion of samples with very large item pools is useful for building up the number of pairwise comparisons that can be made across the comprehensive set of items used. The Eugene-Springfield Community Sample, for example, provides a strong basis for evaluating the associations among many items, though this can be improved with the addition of more diverse samples and an even broader range of items. If a portion of item content overlaps across samples, it is even possible to estimate associations for items that were not administered to both samples using, for example, factor extension (Dwyer, 1937; Horn, 1973) and multiple imputation procedures (Azur et al., 2011).

With a sufficiently large number of data sets, it will also become possible to consider the effect of situational or grouping variables on the associations between items, as first suggested with Brunswik's (1955) representative designs. For example, it would be possible to evaluate the extent to which item associations are robust to differences among cohorts, life stages, cultures, or geographic regions (Elleman et al., in press; Möttus, Wood, Condon et al., 2020). Eventually, it may also become feasible to consider effects like those described in the "occasions" dimension of Cattell's data box (1966). These analyses would extend beyond the retest effects discussed in Step 3 to include evaluations of item stability over time (Revelle & Condon, 2019), as may be affected by factors such as aging, interventions, or disease courses.

Step 6. Evaluations of predictive utility

Following or during the aggregation of informant- and self-report data, it will also be important to consider the extents to which items differentially predict outcomes of interest; it is particularly important that the accuracies of these predictions be estimated in (sub)samples other than those in which the models were initially created (Möttus, Wood, Condon et al., in press; Yarkoni & Westfall, 2017). The outcomes to be predicted may include cross-sectional associations and/or predictions over time. Important outcomes could include a wide range of criteria such as how much people are liked by others, their socioeconomic success, their health, how their romantic relationships are going (or if they even have them), their overall life satisfaction, how well they are doing at their jobs, and so on (Saucier, in press; Seebboth & Möttus, 2018; Wessels et al., 2020). Items' associations with demographic variables such as age (Möttus & Rozgonjuk, 2019), gender (Möttus, Sinick et al., 2019) or geographic residency (Achaa-Amankwaa, Oлару, & Schroeders, 2020) are also useful for establishing the degrees to which they capture unique

information with potential descriptive, predictive, or explanatory value. The identification of overlap between predicted outcomes and item content will be needed to guard against circular causal reasoning (Vainik et al., 2015). Similarly, we may determine that certain items have reliable unique variance but virtually never correlate with important outcomes. When assessment time is limited, such items may be removed from the item pool to be rated by participants.

Conclusion

In this work, we have introduced a rationale and procedures for developing building blocks for a personality taxonomy that is unlike the many top-down frameworks currently used in personality research, but instead proposes a bottom-up approach, with the aim of more comprehensively canvassing the domain of personality traits. These recommendations are only the beginning of a long development process, for this taxonomy will likely require collaboration across many research groups over several years. Along the way, the procedures proposed here will benefit from the input of other research teams and they will inevitably evolve as they are carried out. We believe and have argued that the benefits of this endeavor, and the specific procedures we have suggested for carrying it out, are likely to be substantial for the descriptive, predictive, and explanatory approaches to personality research outlined by Möttus, Wood, Condon et al. (in press).

A key limitation of the proposed procedures is the focus on survey methods, to the exclusion of other assessment methods in personality research. Several researchers have raised concerns that personality psychology (and psychology more broadly) has become overly reliant on the use of surveys and have begun calling for greater use of observational techniques (Baumeister, Vohs, & Funder, 2007; Back, 2020) and other methods. We acknowledge that there are good reasons to prospectively move away from the over-use of surveys in personality science, and suggest that new technologies may serve to accelerate such efforts. For example, rapidly-evolving advances in natural language processing are enabling the extraction of clues about personality from virtually all types of text (Kjell, 2019), and access to digital footprint data makes it increasingly possible to observe behavioral expression virtually (Rauthmann, 2020).

The use of such research methodologies poses exciting opportunities to extend our knowledge of personality in the near future. If, in time, the proportion of data collected by personality psychologists swings dramatically towards the use of some other method -- say, natural language processing -- the procedures proposed here can be adapted for dealing with information signals of other varieties. This may prove important given that much of the objective behavioral data (e.g., digital traces) does not have inherent psychological meaning by itself; the data typically acquire meaning through validation against survey-based self- and informant-reports (Möttus, Wood, Condon et al., in press). We leave the treatment of other signal types for future research given that the popularity of surveys is currently well-entrenched.

In closing, it is worth acknowledging that the pursuit of procedures like those proposed here poses an unprecedented opportunity for collective action within the field. Though wide collaborative efforts often require considerable organizational effort and expense, they also can provide unique benefits in hastening the pace of research and enhancing our understanding.

References

- Achaa-Amankwaa, P., Oлару, G., & Schroeders, U. (2020, April 16). Coffee or Tea? Examining Cross-Cultural Differences in Personality Nuances Across Former Colonies of the British Empire. <https://doi.org/10.31234/osf.io/dpqrx>
- Allport, G. W. (1966). Traits revisited. *American Psychologist*, 21(1):1.
- Allport, G. W., & Odbert, H. S. (1936). Trait-names: A psycho-lexical study. *Psychological Monographs: General and Applied*, 47(1):1-170.
- Ashton, M. C., & Lee, K. (2020). Objections to the hexaco model of personality structure - and why those objections fail. *European Journal of Personality*.
- Ashton, M. C., Lee, K., Perugini, M., Szarota, P., de Vries, R. E., Di Blas, L., Boies, K., & De Raad, B. (2004). A six-factor structure of personality-descriptive adjectives: Solutions from psycholexical studies in seven languages. *Journal of Personality and Social Psychology*, 86(2):356-366.
- Ashton, M. C., Lee, K., & Goldberg, L. R. (2004). A hierarchical analysis of 1,710 english personality-descriptive adjectives. *Journal of Personality and Social Psychology*, 87(5):707.
- Azur, M. J., Stuart, E. A., Frangakis, C., & Leaf, P. J. (2011). Multiple imputation by chained equations: What is it and how does it work?: Multiple imputation by chained equations. *International Journal of Methods in Psychiatric Research*, 20(1), 40–49.
- Back, M. D. (2020). A Brief Wish List for Personality Research. *European Journal of Personality*, 34(1), 3-7.
- Baumeister, R. F., Vohs, K. D., & Funder, D. C. (2007). Psychology as the Science of Self-Reports and Finger Movements: Whatever Happened to Actual Behavior? *Perspectives on Psychological Science*, 2, 396–403.
- Baumert, A., Schmitt, M., Perugini, M., Johnson, W., Blum, G., Borkenau, P., Costantini, G., Denissen, J.J., Fleeson, W., Grafton, B., & Jayawickreme, E. (2017). Integrating personality structure, personality process, and personality development. *European Journal of Personality*, 31(5), pp. 503-528.
- Block, J. (1961). *The Q-sort method in personality assessment and psychiatric research*. Charles C Thomas Publisher.
- Block, J., Weiss, D. S., & Thorne, A. (1979). How relevant is a semantic similarity interpretation of personality ratings? *Journal of Personality and Social Psychology*, 37, 1055–1074.

- Block, J. (1995). A contrarian view of the five-factor approach to personality description. *Psychological Bulletin*, 117(2):187-215.
- Booth, T. and Murray, A. L. (2018). How Factor Analysis Has Shaped Personality Trait Psychology. In Irwing, P., Booth, T., and Hughes, D. J., editors, *The Wiley Handbook of Psychometric Testing: A Multidisciplinary Reference on Survey, Scale and Test Development*, chapter 29, pages 933-951. John Wiley & Sons, Ltd.
- Bouchard, T. J. (2016). Experience producing drive theory: Personality “writ large”. *Personality and Individual Differences*, 90:302-314.
- Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, 62(3), 193.
- Cattell, R. B. (1946). *Description and measurement of personality*. World Book Company, Oxford, England.
- Cattell, R. B. (1966). The data box: its ordering of total resources in terms of possible relational systems. In: Cattell R. B. (Ed.), *Handbook of multivariate experimental psychology*. pp. 67–128. Chicago, IL: Rand McNally.
- Cella, D., Choi, S. W., Condon, D. M., Schalet, B., Hays, R. D., Rothrock, N. E., ... & DeWalt, D. A. (2019). PROMIS® adult health profiles: efficient short-form measures of seven health domains. *Value in health*, 22(5), 537-544.
- Chandler, J. (2018). Likeableness and meaningfulness ratings of 555 (+487) person-descriptive words. *Journal of Research in Personality*, 72, 50-57.
- Condon, D. M. (2018). The SAPA Personality Inventory: An empirically-derived, hierarchically-organized self-report personality assessment model. *PsyArXiv*, 1-444.
- Condon, D. (2019). Database of Individual Differences Survey Tools, <https://doi.org/10.7910/DVN/T1NQ4V>, Harvard Dataverse.
- Costa, P. T. , & McCrae, R. R. (1995). Domains and facets: Hierarchical personality assessment using the Revised NEO Personality Inventory. *Journal of Personality Assessment*, 64(1):21-50.
- Cramer, A. O., 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(4), 414-431.
- Crowe, M. L., Lynam, D. R., & Miller, J. D. (2018). Uncovering the structure of agreeableness from self-report measures. *Journal of Personality*, 86(5), 771-787.

- Dawkins, R. (1976). *Hierarchical organisation: A candidate principle for ethology*. In P. P. G. Bateson & R. A. Hinde (Eds.), *Growing points in ethology*. Cambridge University Press.
- De Vries, R. E., De Vries, A., De Hoogh, A., and Feij, J. (2009). More than the Big Five: Egoism and the HEXACO model of personality. *European Journal of Personality*, 23(8), 635-654.
- DeYoung, C. G. (2006). Higher-order factors of the Big Five in a multi-informant sample. *Journal of Personality and Social Psychology*, 91(6):1138-1151.
- DeYoung, C. G. (2015). Cybernetic big five theory. *Journal of research in personality*, 56, 33-58.
- DeYoung, C. G., Quilty, L. C., and Peterson, J. B. (2007). Between facets and domains: 10 aspects of the Big Five. *Journal of Personality and Social Psychology*, 93(5):880-896.
- Digman, J. M. (1997). Higher-order factors of the Big Five. *Journal of personality and social psychology*, 73(6), 1246.
- Dumas, J. E., Johnson, M., & Lynch, A. M. (2002). Likableness, familiarity, and frequency of 844 person-descriptive words. *Personality and individual differences*, 32(3), 523-531.
- Dwyer, P. S. (1937). The determination of the factor loadings of a given test from the known factor loadings of other tests. *Psychometrika*, 2(3):173-178.
- Elleman, L., Condon, D. M., Holtzman, N. S., Allen, V. R., & Revelle, W. (in press). Smaller is better: Associations between personality and demographics are improved by examining narrower traits and regions. *Collabra*.
- Elleman, L. G., McDougald, S. K., Condon, D. M., & Revelle, W. (2020). That takes the BISCUIT: A comparative study of predictive accuracy and parsimony of four statistical learning techniques in personality data, with data missingness conditions. *European Journal of Psychological Assessment*.
- Epskamp, S., Rhemtulla, M., & Borsboom, D. (2017). Generalized network psychometrics: Combining network and latent variable models. *Psychometrika*, 82(4), 904-927.
- Eysenck, H. J. (1994). The big five or the giant three: Criteria for a paradigm. In Halverson, C. F., Kohnstamm, G. A., and Martin, R. P., editors, *The developing structure of temperament and personality from infancy to adulthood*, pages 37-51. Lawrence Erlbaum Associates, Hillsdale, N.J.
- Fitzpatrick, R. (2014). From charred death to deep filthstep: The 1,264 genres that make modern music. *The Guardian*, 4.

- Funder, D. C. (1991). Global traits: A neo-Allportian approach to personality. *Psychological Science*, 2(1):31-39.
- Funder, D. C., & Dobroth, K. M. (1987). Differences between traits: Properties associated with interjudge agreement. *Journal of personality and social psychology*, 52(2), 409.
- Furr, R. M. (2009). Personality psychology as a truly behavioural science. *European Journal of Personality*, 23(5), 369-401.
- Goldberg, L. R. (1990). An alternative "description of personality": The Big-Five factor structure. *Journal of Personality and Social Psychology*, 59(6):1216-1229.
- Goldberg, L. R. (1992). The development of markers for the Big-Five factor structure. *Psychological Assessment*, 4(1), 26-42. doi:10.1037/1040-3590.4.1.26
- Goldberg, L. R. (1993). The structure of phenotypic personality traits. *American Psychologist*, 48:26-34.
- Goldberg, L. R. (1999). A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several Five-Factor Models. In Mervielde, I., Deary, I., De Fruyt, F., & Ostendorf, F., editors, *Personality and Individual Differences*, pages 7-28. Tilburg University Press, The Netherlands.
- Goldberg, L. R. (2006). Doing it all bass-ackwards: The development of hierarchical factor structures from the top down. *Journal of Research in Personality*, 40(4):347-358.
- Goldberg, L. R., Johnson, J. A., Eber, H. W., Hogan, R., Ashton, M. C., Cloninger, C. R., & Gough, H. G. (2006). The international personality item pool and the future of public-domain personality measures. *Journal of research in personality*, 40(1):84-96.
- Goldberg, L. R., & Saucier, G. (2016). The Eugene-Springfield Community Sample: Information Available from the Research Participants. Technical Report 1, Oregon Research Institute.
- Greiff, S. & Iliescu, D. (2017). A test is much more than just the test. Some thoughts on adaptations and equivalence. *European Journal of Psychological Assessment*, 33, 145-148.
- Hampson, S. E., Goldberg, L. R., & John, O. P. (1987). Category-breadth and social-desirability values for 573 personality terms. *European Journal of Personality*, 1(4), 241-258.
- Henry, S., & Möttus, R. (2020). Traits and Adaptations: A Theoretical Examination and New Empirical Evidence. *European Journal of Personality*, 34, 265–284.
- Hogan, J., & Roberts, B. W. (1996). Issues and non-issues in the fidelity–bandwidth trade-off. *Journal of organizational behavior*, 17(6), 627-637.

- Hopwood, C. J., & Donnellan, M. B. (2010). How should the internal structure of personality inventories be evaluated?. *Personality and Social Psychology Review*, 14(3), 332-346.
- Horn, J. L. (1973). On extension analysis and its relation to correlations between variables and factor scores. *Multivariate Behavioral Research*, 8(4):477 – 489.
- John, O. P., Goldberg, L. R., & Angleitner, A. (1984). Better than the alphabet: Taxonomies of personality-descriptive terms in English, Dutch, and German. In *Personality psychology in Europe, volume 1*.
- John, O. P., & Srivastava, S. (1999). The Big Five trait taxonomy: History, measurement, and theoretical perspectives. In *The Handbook of Personality: Theory and Research*, pages 102-138. Guilford Press, London, UK.
- Kenrick, D. T., & Funder, D. C. (1988). Profiting from controversy: Lessons from the person-situation debate. *American Psychologist*, 43(1), 23–34.
- Kjell, O. N., Kjell, K., Garcia, D., and Sikström, S. (2019). Semantic measures: Using natural language processing to measure, differentiate, and describe psychological constructs. *Psychological Methods*, 24(1):92
- Kretzschmar, A., Spengler, M., Schubert, A. L., Steinmayr, R., & Ziegler, M. (2018). The Relation of Personality and Intelligence—What Can the Brunswik Symmetry Principle Tell Us? *Journal of Intelligence*, 6, 30.
- Krueger, R. F., Kotov, R., Watson, D., Forbes, M. K., Eaton, N. R., Ruggero, C. J., ... & Zimmermann, J. (2018). Progress in achieving quantitative classification of psychopathology. *World Psychiatry*, 17(3), 282-293.
- Larsen, K. R., & Bong, C. H. (2016). A tool for addressing construct identity in literature reviews and meta-analyses. *MIS Quarterly*, 40(3):529-551.
- Leising, D., Scharloth, J., Lohse, O., & Wood, D. (2014). What types of terms do people use when describing an individual's personality?. *Psychological science*, 25(9), 1787-1794.
- Loehlin, J. C., & Goldberg, L. R. (2014). Do personality traits conform to lists or hierarchies? *Personality and individual differences*, 70, 51-56.
- Lowman, G. H., Wood, D., Armstrong, B. F. I., Harms, P. D., & Watson, D. (2018). Estimating the reliability of emotion measures over very short intervals: The utility of within-session retest correlations. *Emotion*, 18(6):896.
- MacCann, C., Duckworth, A. L., & Roberts, R. D. (2009). Empirical identification of the major facets of conscientiousness. *Learning and Individual Differences*, 19(4), 451-458.

- Markon, K. E. (2009). Hierarchies in the structure of personality traits. *Social and Personality Psychology Compass*, 3(5):812-826.
- McCrae, R. R. (2015). A More Nuanced View of Reliability: Specificity in the Trait Hierarchy. *Personality and Social Psychology Review*, 19, 97–112.
- McCrae, R. R., Kurtz, J. E., Yamagata, S., & Terracciano, A. (2011). Internal consistency, retest reliability, and their implications for personality scale validity. *Personality and social psychology review*, 15(1), 28-50.
- McCrae, R. R., & John, O. P. (1992). An introduction to the five-factor model and its applications. *Journal of Personality*, 60(2):175-215.
- McCrae, R. R., & Möttus, R. (2019). What Personality Scales Measure: A New Psychometrics and Its Implications for Theory and Assessment. *Current Directions in Psychological Science*, 28, 415–420.
- McCrae, R. R., & Sutin, A. R. (2018). A five-factor theory perspective on causal analysis. *European Journal of Personality*, 32(3), 151-166.
- McCrae, R. R., Möttus, R., Hřebíčková, M., Realo, A., & Allik, J. (2019). Source method biases as implicit personality theory at the domain and facet levels. *Journal of Personality*, 87(4), 813–826.
- Möttus, R. (2016). Towards more rigorous personality trait-outcome research. *European Journal of Personality*, 30(4):292-303.
- Möttus, R., Bates, T. C., Condon, D. M., Mroczek, D., & Revelle, W. (2017). Leveraging a more nuanced view of personality: Narrow characteristics predict and explain variance in life outcomes. <https://doi.org/10.31234/osf.io/4q9gv>
- Möttus, R., Kandler, C., Bleidorn, W., Riemann, R., & McCrae, R. R. (2017). Personality traits below facets: The consensual validity, longitudinal stability, heritability, and utility of personality nuances. *Journal of Personality and Social Psychology*, 112(3):474.
- 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.
- Möttus, R., & Rozgonjuk, D. (2019). Development is in the details: Age differences in the Big Five domains, facets and nuances. *Journal of Personality and Social Psychology*.
- Möttus, R., Sinick, J., Terracciano, A., Hřebíčková, M., Kandler, C., Ando, J., Mortensen, E. L., Colodro-Conde, L., & Jang, K. L. (2019). Personality characteristics below facets: A replication and meta-

- analysis of cross-rater agreement, rank-order stability, heritability, and utility of personality nuances. *Journal of Personality and Social Psychology*, 117(4):e35.
- Möttus, R., Wood, D., Condon, D. M., Back, M., Baumert, A., Costantini, G., Epskamp, S., Greiff, S., Johnson, W., Lukaszewski, A., Murray, A., Revelle, W., Wright, A. G. C., Yarkoni, T. Ziegler, M., & Zimmermann, J. (in press). Descriptive, predictive and explanatory personality research: Different goals, different approaches, but a shared need to move beyond the Big Few traits. *European Journal of Personality*.
- Norman, W. T. (1963). Toward an adequate taxonomy of personality attributes: Replicated factor structure in peer nomination personality ratings. *Journal of Abnormal and Social Psychology*, 66(6):574-583.
- Norman, W. T. (1967). *2800 Personality Trait Descriptors: Normative Operating Characteristics for a University Population*. Unpublished manuscript, University of Michigan, Department of Psychology, Ann Arbor.
- Olino, T. M., Klein, D. N., Durbin, C. E., Hayden, E. P., & Buckley, M. E. (2005). The structure of extraversion in preschool aged children. *Personality and Individual Differences*, 39(2), 481-492.
- Ones, D. S., & Viswesvaran, C. (1996). Bandwidth–fidelity dilemma in personality measurement for personnel selection. *Journal of organizational behavior*, 17(6), 609-626.
- Paunonen, S. V., & Jackson, D. N. (2000). What is beyond the big five? Plenty!. *Journal of personality*, 68(5), 821-835.
- Rauthmann, J. (2020). A (More) Behavioral Science of Personality in the Age of Multi-Modal Sensing, Big Data, Machine Learning, and Artificial Intelligence. *European Journal of Personality*.
- Revelle, W., & Condon, D. M. (2019). Reliability from α to ω : A tutorial. *Psychological assessment*, 31(12), 1395.
- Revelle, W., Dworak, E. M., & Condon, D. M. (2020). Exploring the persome: The power of the item in understanding personality structure. *Personality and Individual Differences*.
- Roberts, B. W., Chernyshenko, O. S., Stark, S., & Goldberg, L. R. (2005). The structure of conscientiousness: An empirical investigation based on seven major personality questionnaires. *Personnel Psychology*, 58(1):103-139.
- Rohrer, J. M. (2018). Thinking clearly about correlations and causation: Graphical causal models for observational data. *Advances in Methods and Practices in Psychological Science*, 1(1):27-42.
- Saucier, G. (1997). Effects of variable selection on the factor structure of person descriptors. *Journal of Personality and Social Psychology*, 73(6):1296.

- Saucier, G. (2020). Language, Subjectivity, Culture, Comprehensiveness, and Structure. In Rauthmann, J. F., Sherman, R. A., Funder, D. C., editors, *The Oxford Handbook of Psychological Situations*, pages 375-388. Oxford University Press.
- Saucier, G., & Goldberg, L. R. (1996). The language of personality: Lexical perspectives on the five-factor model. In J. S. Wiggins (Ed.), *The five-factor model of personality: Theoretical perspectives*. (pp. 21–50). New York: Guilford.
- Saucier, G., & Iurino, K. (2019). High-dimensionality personality structure in the natural language: Further analyses of classic sets of english-language trait-adjectives. *Journal of Personality and Social Psychology*.
- Saucier, G., Iurino, K., & Thalmayer, A. G. (in press). Comparing predictive validity in a community sample: High-dimensionality and traditional domain-and-facet structures of personality variation. *European Journal of Personality*.
- Schwaba, T., Rhemtulla, M., Hopwood, C. J., & Bleidorn, W. (2020). A facet atlas: Visualizing networks that describe the blends, cores, and peripheries of personality structure. *PloS one*, 15(7):e0236893.
- Seeboth, A., & Möttus, R. (2018). Successful Explanations Start with Accurate Descriptions: Questionnaire Items as Personality Markers for More Accurate Predictions. *European Journal of Personality*, 32, 186–201.
- Siegling, A. B., Petrides, K. V., & Martskvishvili, K. (2015). An Examination of a New Psychometric Method for Optimizing Multi-Faceted Assessment Instruments in the Context of Trait Emotional Intelligence. *European Journal of Personality*, 29(1), 42-54.
- Soto, C. J., & John, O. P. (2017). The Next Big Five Inventory (BFI-2): Developing and Assessing a Hierarchical Model With 15 Facets to Enhance Bandwidth, Fidelity, and Predictive Power. *Journal of Personality and Social Psychology*, 113(1):117-143.
- Tellegen, A. (1991). Personality traits: Issues of definition, evidence, and assessment. In D. Cicchetti & W. M. Grove (Eds.), *Thinking clearly about psychology: Essays in honor of Paul E. Meehl, Vol. 2: Personality and psychopathology*. (pp. 10–35). Minneapolis, MN: University of Minnesota Press.
- Tracy, J. L., & Robins, R. W. (2007). The psychological structure of pride: A tale of two facets. *Journal of Personality and Social Psychology*, 92, 506–525
- Vainik, U., Möttus, R., Allik, J., Esko, T., & Realo, A. (2015). Are trait–outcome associations caused by scales or particular items? Example analysis of personality facets and BMI. *European Journal of Personality*, 29(6), 622-634.

- Vazire, S. (2010). Who knows what about a person? The self-other knowledge asymmetry (SOKA) model. *Journal of Personality and Social Psychology*, 98, 281–300.
- Watson, D. (2004). Stability versus change, dependability versus error: Issues in the assessment of personality over time. *Journal of Research in Personality*, 38, 319–350.
- Watson, D., & Clark, L. (1994). *The PANAS-X: Manual for the Positive and Negative Affect Schedule - Expanded Form* (p. 24 pages). University of Iowa.
- Wessels, N. M., Zimmermann, J., & Leising, D. (2020). Who Knows Best What the Next Year Will Hold for You? The Validity of Direct and Personality-based Predictions of Future Life Experiences Across Different Perceivers. *European Journal of Personality*.
- Wood, D. (2015). Testing the lexical hypothesis: Are socially important traits more densely reflected in the English lexicon? *Journal of Personality and Social Psychology*, 108(2):317.
- Wood, D., Nye, C. D., & Saucier, G. (2010). Identification and measurement of a more comprehensive set of person-descriptive trait markers from the English lexicon. *Journal of Research in Personality*, 44(2):258-272.
- Wood, D., Gardner, M. H., & Harms, P. D. (2015). How functionalist and process approaches to behavior can explain trait covariation. *Psychological review*, 122(1):84.
- Wood, D., Harms, P. D., Lowman, G., Soto, C. J., Qiu, L., John, O. P., & Jiahui, L. (n.d.). Evaluating the utility of within-session retest correlations as reliability estimates. University of Alabama, Tuscaloosa, AL.
- Wood, D., Lowman, G., Armstrong, B., Harms, P. D., Denissen, J. J. A., & Chung, J. M. (n.d.). Building better similarity detectors through repeated administrations of the same inventory. University of Alabama, Tuscaloosa, AL.
- Wood, D., Qiu, L., Lu, J., Lin, H., & Tov, W. (2018). Adjusting bilingual ratings by retest reliability Improves estimation of translation quality. *Journal of Cross-Cultural Psychology*, 49, 1325–1339.
- Wood, D. & Wortman, J. (2012). Trait means and desirabilities as artifactual and real sources of differential stability of personality traits. *Journal of Personality*, 80(3):665-701.
- Yarkoni, T. (2010). The abbreviation of personality, or how to measure 200 personality scales with 200 items. *Journal of research in personality*, 44(2):180-198.
- Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science*, 12, 1100–1122.

Ziegler, M., & Bäckström, M. (2016). 50 facets of a trait—50 ways to mess up? *European Journal of Psychological Assessment*, 32(2), 105-110.

Ziegler, M., Kemper, C. J., & Lenzner, T. (2015). The issue of fuzzy concepts in test construction and possible remedies. *European Journal of Psychological Assessment*, 31(1), 1-4.