



OPEN

A prediction-focused approach to personality modeling

Gal Lavi¹, Jonathan Rosenblatt²✉ & Michael Gilead³✉

In the current study, we set out to examine the viability of a novel approach to modeling human personality. Research in psychology suggests that people's personalities can be effectively described using five broad dimensions (the Five-Factor Model; FFM); however, the FFM potentially leaves room for improved predictive accuracy. We propose a novel approach to modeling human personality that is based on the maximization of the model's predictive accuracy. Unlike the FFM, which performs *unsupervised* dimensionality reduction, we utilized a *supervised* machine learning technique for dimensionality reduction of questionnaire data, using numerous psychologically meaningful outcomes as data labels (e.g., intelligence, well-being, sociability). The results showed that our five-dimensional personality summary, which we term the "Predictive Five" (PF), provides predictive performance that is better than the FFM on two independent validation datasets, and on a new set of outcome variables selected by an independent group of psychologists. The approach described herein has the promise of eventually providing an interpretable, low-dimensional personality representation, which is also highly predictive of behavior.

Humans significantly differ from each other. Some people's idea of fun is partying all night long, and others enjoy binging on a TV series while eating snacks; some are extremely intelligent, and others less so; some are hot-headed, and others remain cool, no matter what. Because of this variety, predicting humans' thoughts, feelings, and behaviors is a cumbersome task; nonetheless, we attempt to solve this task on a daily basis. For example, when we decide who to marry, we try to predict whether we can depend on the other person till death do us part; when we choose a career, we must do our best to predict whether we will be successful and fulfilled in a given profession.

In order to predict a person's thoughts, feelings, and behaviors, people often have no other option but to generate something akin to a scientific theory¹—a parsimonious model that attempts to capture the unique characteristics of individuals, and that could be used to predict their behavior in novel circumstances. Indeed, research shows that people employ such theories when predicting their own² and others' behaviors. Unfortunately, theories based strictly on intuition are often highly inaccurate³, even if produced by professional psychological theoreticians⁴. In light of this, ever since the early days of psychology research, scholars have been attempting to devise personality models using the scientific method, giving rise to the longstanding field of personality science.

Personality, when used as a scientific term, refers to the mental features of individuals that characterize them across different situations, and thus can be used to predict their behavior. In the early years of personality research, scientists generated numerous competing theories and measures, but struggled to arrive at a scientific consensus regarding the core structure of human personality. In recent decades, a consensus theory of the core dimensions of human personality has emerged—the Five Factor Model (FFM).

The FFM emerged from the so-called "lexical paradigm", which assumed that if people regularly exhibit a form of behavior that is meaningful to human life, then language will produce a term to describe it⁵. Given this assumption, personality psychologists performed research wherein they asked individuals to rate themselves on lists of common English language trait words (e.g., friendly, upbeat), and then developed and used early dimensionality-reduction methods to find a parsimonious model that can account for much of the variability in each person's trait ratings⁵.

Much research shows that these five factors, often termed the "Big Five" are relatively stable over time and have convergent and discriminant validity across methods and observers⁶. Moreover, research into the FFM has replicated the dimensional structure in different samples, languages, and cultures^{7,8} (but see⁹ for a recent criticism). In light of this, the FFM is taken by some to reflect a comprehensive ontology of the psychological makeup of human beings¹⁰ according to McCrae and Costa¹¹ the five factors are "both necessary and reasonably sufficient for describing at a global level the major features of personality".

¹Department of Industrial Engineering and Management, Ben-Gurion University of the Negev, Beersheba, Israel. ²Pagaya Technologies, Tel Aviv, Israel. ³School of Psychological Sciences and Sagol School of Neuroscience, Tel Aviv University, Tel Aviv, Israel. ✉email: john.ros@gmail.com; michael.gilead@gmail.com

Surely, human beings are complex entities, and their personality is not fully captured by five dimensions; however, the importance of having a parsimonious model of humans' psychological diversity cannot be overstated. As noted by John and Srivasta¹², a parsimonious taxonomy permits researchers to study "specified domains of personality characteristics, rather than examining separately the thousands of particular attributes that make human beings individual and unique." Moreover, as they note, such a taxonomy greatly facilitates "the accumulation and communication of empirical findings by offering a standard vocabulary, or nomenclature".

An additional consequence of having a parsimonious model of the core dimensions of human personality, is that such an abstraction enables the acquisition of novel knowledge via statistical learning (see¹³ for a discussion of the importance of abstract representations in learning); namely, whereas the estimation of covariances between high-dimensional vectors is often highly unreliable (i.e., the so-called "curse of dimensionality"¹⁴), learning the statistical correlates of a low-dimensional structure is a more tractable problem. For example, research has shown that participants' self-reported ratings on the FFM dimensions can be reliably estimated based on their digital footprint¹⁵.

This ability to infer individuals' personality traits using machine learning also raises serious concerns, as it may be used for effective psychological manipulation of the public. In 2013, a private company named *Cambridge Analytica* harvested the data of Facebook users, and used statistical methods to infer the personality characteristics of hundreds of millions of Americans¹⁶. This psychological profile of the American population was supposedly used by the Trump campaign in an attempt to tailor political advertisements based on an individuals' specific personality profile. While the success of these methods remains unclear, given the vast amount of data accumulated by companies such as Alphabet and Meta, the potential dangers of machine-learning based psychological profiling is taken by many to be a serious threat to democracy¹⁷.

Even if dubious entities indeed manage to acquire the Big Five personality profile of entire populations, it is far from obvious that such information could be used to generate actionable predictions. Indeed, the FFM was criticized by some researchers for its somewhat limited contribution to predicting outcomes on meaningful dimensions^{18–20}. In light of such claims, some have argued that the public concern over the Cambridge Analytica scandal was overblown²¹ (but see²² for evidence for potential reasons for concern).

Roberts et al.²³ present counter-argument for critical stances against the predictive accuracy of the FFM and note that: "As research on the relative magnitude of effects has documented, personality psychologists should not apologize for correlations between 0.10 and 0.30, given that the effect sizes found in personality psychology are no different than those found in other fields of inquiry." While this claim is clearly true, there is also no doubt that such correlations (that translate to explained variance in the range of 1%–9%) potentially leave room for improvements in terms of predictive accuracy.

If one's goal is to find a parsimonious representation of personality that has better predictive accuracy than the FFM, it could be instructive to remember that the statistical method by which the FFM was produced—namely, Factor Analysis—is not geared towards prediction. Factor analysis is an unsupervised dimensionality-reduction method (i.e., a method that maps original data to a new lower dimensional space without utilizing information regarding outcomes) aimed at maximizing explanatory coherence and semantic interpretability, rather than maximizing predictive ability. It does so by finding a parsimonious, low-dimension representation (e.g., the five Big Five factors: extraversion, neuroticism and so on) that maximizes the variance explained in the higher-dimension domain (e.g., hundreds of responses to questionnaire items; for example, "I am lazy"; "I enjoy meeting new people"). Advances in statistics and machine learning have opened up new techniques for *supervised* dimensionality-reduction. Namely, methods that reduce the dimensionality of a source domain (i.e., predictor variables, X_1, \dots, X_n ; in the case of personality, hundreds of questionnaire items) by focusing on the objective of maximizing the capacity of the lower-dimensional representation to *predict* outcomes of a target domain (outcome variables, Y_1, \dots, Y_m , for example, depression, risky behavior, workplace performance).

Such techniques where dimensionality-reduction is achieved via maximization of predictive accuracy across a host of target-domain outcomes hold the potential of providing psychologists with parsimonious models of a psychological feature space that serve as relatively "generalizable predictors" of important aspects of human behavior. Moreover, it may demonstrate that privacy leaks, a-lá Cambridge-Analytica, are indeed a serious threat to democracy, despite being dismissed by some as science fiction.

In light of this, we investigated whether a supervised dimensionality-reduction approach that takes into account a host of meaningful can potentially improve the predictive performance of personality models. Such an approach could pave the way to a new family of personality models and could advance the study of personality. Alternatively, it may very well be the case that the FFM indeed "carves nature at its joints" and provides the most accurate ontology of the psychological proclivities of humans. In such a case, the FFM may remain the best predictive model of personality, and our approach will not provide improvements in predictions.

In order to examine this question, we conducted three studies. In Study 1, we built a supervised learning model using big data of personality questionnaire items and diverse, important life outcomes. We reduced the dimensionality of 100 questionnaire items into a set of five dimensions, with the objective of simultaneously minimizing prediction errors across ten meaningful life outcomes. We hypothesized that the resulting five-dimensional representation will outperform the FFM representation—when fitting a new model and attempting to predict the ten important outcomes on a held-out dataset. Next, in Studies 2 and 3, we explored the performance of the resulting model on new outcome variables.

Study 1

Method. *Participants.* The analyses relied on the myPersonality dataset that was collected between 2007 and 2012 via the myPersonality Facebook application. The myPersonality database is no longer shared by its creators for additional use. We received approval to download that data from the administrators of myPersonal-

ity on January 7th, 2018, and downloaded the data shortly thereafter. After the myPersonality database was taken down in 2018, we sent an email to the administrators (on June 8th, 2018), and received confirmation that we can use the data we have already downloaded. The application enabled its users to take various validated psychological and psychometric tests, such as different versions of the International Personality Item Pool (IPIP) questionnaire. Many participants also provided informed consent for researchers to access their Facebook usage details (e.g., liked pages). Participation was voluntary and likely motivated by people's desire for self-knowledge²⁴. The Participants in the myPersonality database are relatively representative of the overall population²⁵. All participants provided informed consent for the data they provided to be used in subsequent psychological studies. We used data from 397,851 participants (210,279 females, 142,497 males, and 44,805 did not identify) who answered all of the questions on the 100-item IPIP representation of Goldberg's²⁶ markers for the FFM which are freely available for all types of use. Participants' mean age was 25.7 years ($SD = 8.84$). The study was approved by the Institutional Review Board of Ben-Gurion University, and was conducted in accordance with relevant guidelines and regulations.

Measures. Dependent variables. We sought to use supervised learning in order to find a low-dimensional representation of personality that can be used to predict psychological consequences across a diverse set of domains. We thus focused on ten meaningful outcome variables that were available in the myPersonality database, that cover many dimensions of human life which psychologists care about:

- (1) Intelligence Quotient (IQ), measured with a brief 20 items version of the Raven's Standard Progressive Matrices test²⁷.
- (2) Well-being, measured with the Satisfaction with Life scale²⁸.
- Personal values, measured using two scores representing the two axes from the Schwartz's Values Survey:
- (3) Self-transcendence vs. Self-enhancement values and
- (4) Openness to Change vs. Conservation values²⁹.
- (5) Empathy, measured with the Empathy Quotient Scale³⁰.
- (6) Depression, measured with The Center for Epidemiologic Study Depression (CES-D) scale³¹.
- (7) Risky behavior, measured with a single-item question concerning illegal drug use.
- (8) Self-reports of legal, yet unhealthy behavior (measured as averaging two single-item questions concerning alcohol consumption and smoking).
- (9) Single item self-report of political ideology.
- (10) The number of friends of participants' had on the social network Facebook.

Independent variables. Our independent variables were the participants' answers to the 100 questions included in the IPIP-100 questionnaire³². In this questionnaire, the participants are asked to rate their agreement with various statements related to different behaviors in their life and their general characteristics and competencies, on a scale from 1 (strongly disagree) to 5 (strongly agree). The original use of this questionnaire is to reliably gauge participants' scores on each of the FFM dimensions. It includes five subscales, each containing 20 items; the factor score for each FFM dimension can be calculated as a simple average of these 20 questions (after reverse coding some items). In the current research we treat each item from this list of 100 questions as a separate independent variable, and seek to reduce the dimensionality of this vector using supervised learning.

Model construction. The problem we set out to solve is to find a good predictive model that is: (a) based on the 100 questions of the existing IPIP-100 questionnaire, and (b) uses five variables only, so we can fairly compare it with the FFM. Reduced Rank Regression (RRR) is a tool that allows just that: it can be used to compress the original 100 IPIP items, to a set of five new variables. These new variables are constructed so that they are good predictors, on average, of a large set of outcomes. Unlike Principal Component Analysis (PCA) or Factor Analysis, RRR reduces data dimensionality by optimizing predictive accuracy.

We randomly divided our data into an independent train and test sets. Each subject in the train and test set had 100 scores of the IPIP questionnaire (X_1, X_2, \dots, X_{100}), as well as their score in each of the ten dependent variables (Y_1, Y_2, \dots, Y_{10}).

X ($n \times 100$) and Y ($n \times 100$) have been centered and scaled. We fitted a linear predictor, with coefficient vector:

$$\hat{Y}_j := \sum_{k=1}^{100} X_k C_{kj} \quad j = 1, 2, \dots, 10 \quad (1)$$

And in matrix notation:

$$\hat{Y} = XC \quad (2)$$

Our linear predictors were fully characterized by the matrix C . We wanted these predictors to satisfy the following criteria: (a) minimize the squared prediction loss (b) consist of 5 predictors, i.e., $rank(C) = r = 5$. Criterion (a) ensures the goodness of fit of the model, and criterion (b) ensures a fair comparison with the FFM. The RRR problem amounts to finding a set of predictors, \hat{C} , so that:

$$\hat{C} := \operatorname{argmin}_C \{ \|Y - XC\|^2, \text{ such that } rank(C) = r \}, \quad (3)$$

where $\| \cdot \|$ denotes the Frobenius matrix norm. The matrix C can be expressed as a product of two rank-constrained matrices:

$$C := BA^T \quad (4)$$

where B is of has p rows and r columns, denoted, $p \times r$, and A is of dimension $q \times r$. The model (2) may thus be rewritten as:

$$\hat{Y} = (X\hat{B})\hat{A}^T \quad (5)$$

The $n \times r$ matrix $X\hat{B}$, which we noted \tilde{X} , may be interpreted as our new low-dimension personality representation. Crucially for our purposes, the same set of r predictors is used for all dependent variables. By choosing dependent variables from different domains, we dare argue that this set of predictors can serve as a set of “generalizable predictors”, which we call henceforth the *Predictive Five* (PF). For the details of the estimation of \hat{B} see the attached code. For a good description of the RRR algorithm see³³.

Model assessment. To assess the predictive performance of the PF, and compare it to the predictive properties of the classical FFM, we used a fourfold cross validation scheme. The validation worked as follows: we learned \hat{B} from a train set (397,851 participants) using RRR; we then divided the independent test set (800 participants) into 4 subsets; we learned \hat{A} from a three-quarters part of the test set (600 participants), and computed the R^2 on the holdout test set (200 participants); we iterated this process over the 4-test subsets. The rationale of this scheme is that: (a) predictive performance is assessed using R^2 on a completely *novel dataset*; (b) when learning the predictive model, we wanted to treat the personality attributes as known. We thus learned \hat{B} and \hat{A} from different sets. The size of the holdout set was selected so that R^2 estimates will have low variance. The details of the process can be found in the accompanying code (<https://github.com/GalBenY/Predictive-Five>).

To examine the performance of the RRR algorithm against another candidate reference model we also performed Principal Component Regression (PCR), where we reduced the IPIP questionnaire to its 5 leading principal components, which were then used to predict the outcome variables. We used the resulting model as a point of comparison in follow-up assessment of predictive accuracy. Like the RRR case, we learned the principal components from the train-set (397,851 participants). Next we divided the independent test set (800 participants) into 4 subsets and used a fourfold cross validation: $\frac{3}{4}$ to learn 5 coefficients, and $\frac{1}{4}$ to compute.

In order to calculate the significance of the difference in the predictive accuracy of the models we took the following approach: predictions are essentially paired, since they originate from the same participant. For each participant, we thus computed the (holdout) difference between the (absolute) error of the PF and FFM models: $|\hat{y}_i^{PF} - |\hat{y}_i^{FFM}|$. Given a sample of such differences, comparing the models collapses to a univariate t-test allowing us to reject the null hypothesis that the mean of the differences is 0.

Results. *PF loadings.* Each of the resulting PF dimensions were a weighted linear combination of IPIP-100 item responses. Despite the fact that the resulting model was based on a questionnaire meant to reliably gauge the FFM, the resulting outcome did *not* fully recapitulate the FFM structure. The detailed loadings for each of the resulting five dimensions appears in the supplementary materials (Fig. 1, Supplementary Materials), can be examined in an online application we have created (<https://predictivefive.shinyapps.io/PredictiveFive>), and can be easily gleaned by examining the correlation of PF scores to the FFM scores (Fig. 2). None of the PF dimensions strongly correlated with demographic variables (Table 1, Supplementary Materials). In Fig. 1, we display the correlations between the ten outcome variables, five principle components of these outcome variables (capturing 86% of the total variance), and the five PF dimensions. For example, it can be observed the PF 3 is inversely related to performance on the intelligence test and to empathy.

Predictive performance. The out-of-sample R^2 of the three models is reported in Table 1. From this figure, we learn that the PF-based regression model is a better predictor of the outcome variables. This holds true on average (over behavioral outcomes), but also for nine of the ten outcomes individually. On 5 of the 10 comparisons, the PF-based model significantly outperformed the FFM, and in a single case the FFM-based model significantly outperformed the PF. The average improvement across all 10 measures was 40.8%.

Reproducibility analysis. If it were the case that our model discovery process produces very different loadings when run on different samples of participants, then the ontological status of the PF representation should be called into question.

In order to assess the reproducibility of the PF we split the training dataset from Study 1 into two datasets: sample A with 198,850 participants and sample B with 198,851 participants. We then learned the rotation matrix, B , on each data part, and applied it. Equipped with two independent copies of the PF, $X_l\hat{B}_l$, $l = \{A, B\}$ replicability is measured by the correlation between data-parts, over participants. Table 2 reports this correlation, averaged over the 5 PFs (column “Correlation between replications”). As can be seen, the correlation between the replications is satisfactory-to-high and ranges from 0.7 to 0.98. This suggests that PF representation replicates well across samples.

Reliability analysis. If the same individuals, tested on different occasions, receive markedly different scores on the PF dimensions, then the ontological status of the PF representation should be called into question. To this end, we exploit the fact that 96,682 users answered the IPIP questionnaire twice. The test–retest correlation between these two answers is reported in Table 2 (column “Test–retest correlation”). It varied from 0.69 for the Dimension 3, to 0.79 for both Dimensions 1 and 5, suggesting that the variance captured by these dimensions is indeed (relatively) stable.




Figure 1. Correlations between the 10 outcome variables, 5 principle components of outcome variables, and the 5 PF dimensions.

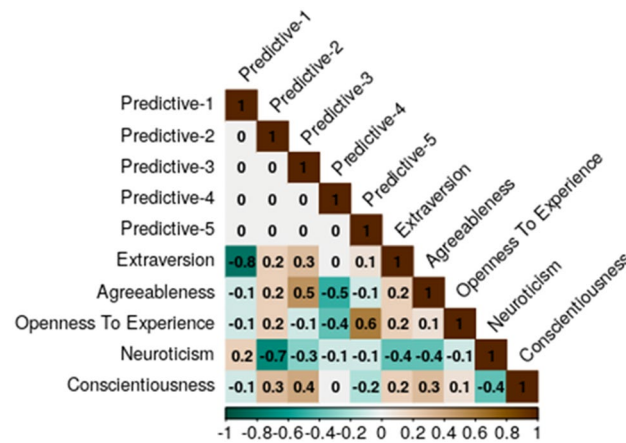


Figure 2. Correlations between the PF and FFM scale scores.

Divergence from the FFM. The superior predictive performance of the PF representation provides evidence that it differs from the FFM. Additionally, as can be gleaned from Fig. 2 (and from the detailed factor loadings’ Supplemental Material), Dimensions 3 and 4 reflect a relatively even combination of several FFM dimensions.

However, these observations do not provide us with an estimate of the degree of agreement between the two *multidimensional* spaces. Prevalent statistical methods of assessment of discriminant validity³⁴ are also not suitable to answer our question regarding the convergence/divergence between the PF and FFM spaces. These various methods only provide researchers with estimates of the agreement between *unidimensional constructs*.

Study	Variable	% R ² Change	Sig.	PF		FFM		PCR	
				MAE	R ²	MAE	R ²	MAE	R ²
1	Depression	1.285		8.563	0.3345 [0.275, 0.394]	8.79	0.304 [0.219, 0.389]	8.814	0.3015 [0.208, 0.395]
1	Empathy	39.098	*	5.562	0.318 [0.266, 0.37]	5.896	0.239 [0.212, 0.266]	5.853	0.247 [0.224, 0.27]
1	Intelligence	138.158	*	10.836	0.1365 [0.092, 0.181]	11.547	0.044 [0.012, 0.076]	11.497	0.05 [0.013, 0.087]
1	Nonhealthy Behaviors	12.941		0.489	0.0585 [0.021, 0.096]	0.493	0.0555 [0.026, 0.085]	0.49	0.066 [0.044, 0.088]
1	Number of Facebook Friendences	48.148	*	140.663	0.0185 [-0.003, 0.04]	141.776	0.01 [-0.007, 0.027]	141.645	0.0105 [-0.008, 0.029]
1	Openness Values	5.079		1.074	0.278 [0.225, 0.331]	1.1	0.2485 [0.182, 0.315]	1.102	0.2385 [0.172, 0.305]
1	Political Orientation	145.498	*	0.834	0.4955 [0.473, 0.518]	1.104	0.1665 [0.122, 0.211]	1.125	0.138 [0.094, 0.182]
1	Risky Behavior	44.444		0.348	0.0395 [-0.025, 0.104]	0.351	0.018 [-0.036, 0.072]	0.349	0.024 [-0.04, 0.088]
1	Transcendence Values	-33.634	*	1.144	0.1685 [0.116, 0.221]	1.077	0.2625 [0.192, 0.333]	1.078	0.2575 [0.194, 0.321]
1	Well-being	7.843	*	0.895	0.3375 [0.29, 0.385]	0.931	0.2865 [0.216, 0.357]	0.933	0.284 [0.216, 0.352]
2	Maternal Overprotection	15.385		4.928	0.011 [0.007, 0.015]	4.936	0.0075 [0.002, 0.013]	4.935	0.008 [0.001, 0.015]
2	Maternal Rejection	92.308	*	4.058	0.0635 [0.052, 0.075]	4.126	0.0345 [0.03, 0.039]	4.112	0.04 [0.034, 0.046]
2	Maternal Warmth	72.222		4.95	0.006 [0.002, 0.01]	4.966	0.003 [-0.004, 0.01]	4.967	0.0025 [-0.004, 0.009]
2	Paternal Overprotection	0.000	*	4.061	0.054 [0.045, 0.063]	4.123	0.031 [0.024, 0.038]	4.12	0.0325 [0.026, 0.039]
2	Paternal Rejection	65.789	*	4.088	0.0505 [0.029, 0.072]	4.132	0.0345 [0.021, 0.048]	4.133	0.035 [0.023, 0.047]
3	Paternal Warmth	50.000	*	3.874	0.0465 [0.031, 0.062]	3.922	0.0265 [0.017, 0.036]	3.925	0.026 [0.014, 0.038]
3	Question 1	15.556		0.72	0.374 [0.28, 0.468]	0.727	0.3225 [0.24, 0.405]	0.727	0.331 [0.255, 0.407]
3	Question 2	61.607	*	1.357	0.137 [0.093, 0.181]	1.434	0.063 [0.014, 0.112]	1.417	0.0805 [0.025, 0.136]
3	Question 3	24.590		0.869	0.3855 [0.315, 0.456]	0.923	0.301 [0.236, 0.366]	0.928	0.306 [0.237, 0.375]
3	Question 4	31.065		1.025	0.3425 [0.242, 0.443]	1.056	0.2785 [0.219, 0.338]	0.999	0.3445 [0.284, 0.405]
3	Question 5	-18.723	*	1.081	0.1225 [0.054, 0.191]	1.033	0.179 [0.123, 0.235]	1.048	0.167 [0.113, 0.221]
3	Question 6	34.828	*	1.018	0.346 [0.301, 0.391]	1.072	0.2595 [0.229, 0.29]	1.053	0.2925 [0.257, 0.328]
3	Question 7	72.131	*	1.521	0.066 [0.027, 0.105]	1.581	0.0245 [-0.012, 0.061]	1.565	0.0385 [0, 0.077]
3	Question 8	-25.161		1.671	0.074 [0.032, 0.116]	1.659	0.0765 [-0.002, 0.155]	1.634	0.096 [0.017, 0.175]
3	Question 9	-23.693	*	1.239	0.3975 [0.357, 0.438]	1.04	0.549 [0.524, 0.574]	1.016	0.5705 [0.551, 0.59]
3	Question 10	-21.757		1.375	0.1525 [0.118, 0.187]	1.339	0.182 [0.125, 0.239]	1.307	0.2245 [0.172, 0.277]
3	Question 11	42.657	*	1.513	0.1665 [0.129, 0.204]	1.591	0.115 [0.087, 0.143]	1.562	0.1365 [0.112, 0.161]
3	Question 12	159.567	*	0.922	0.664 [0.609, 0.719]	1.502	0.232 [0.187, 0.277]	1.529	0.214 [0.16, 0.268]

Table 1. Comparison of the predictive performance of the different models. PF—Performance of the model based on the Predictive Five representation; FFM—performance of model based on the Five Factor Model representation; PCR—performance of the Principal Component Regression model; MAE—Mean Absolute Error. For each outcome variable, the most successful model (in terms of predictive accuracy) is colored in green, the second-best model is colored in orange, and the least successful model is colored in pink. Percent R² change and significance pertains to the focal comparison (PF-based MAE vs. FFM-based MAE). Significant focal comparisons are denoted by an asterisk, and survive a threshold of $p < .05$, FDR corrected. In 14 cases, the PF-based model significantly outperformed the FFM-based model, in 3 cases, the FFM-based model was significantly better than the PF-based model.

	Test–retest correlation	Correlation between replications
Dimension 1	0.79	0.98
Dimension 2	0.75	0.77
Dimension 3	0.69	0.70
Dimension 4	0.71	0.79
Dimension 5	0.79	0.86

Table 2. Reproducibility and reliability analysis. Test–retest correlation represents the test–retest reliability within participants. The correlation between replications captures the reproducibility of our dimensional structure across participants.

Nonetheless, the underlying logic behind these methods (i.e., a formalization of a multitrait-multimethod matrix³⁵) is still applicable to our case. We calculated an estimate of agreement between the FFM and the PF spaces using *cosine similarity*, which gauges the angle between two points in a multidimensional space (the smaller the angle, the closer are the points). Our rationale is that if the FFM scores differ from the PF, they should span different spaces. The cosine similarity *within* measures (in our case, first and second measurements, denoted T1 and T2) should thus be larger than the similarity *between* measures (FFM to PF).

We used the data from the 96,682 participants for which we had test–retest data. Instead of computing standard test–retest correlations, we calculated a multidimensional test–rest score as the cosine similarity of participants’ scores on the first and second measurement, for both the FFM and PF. These estimates are expected to be highly similar and provide an upper bound on the similarity measure, partially analogous to the diameter

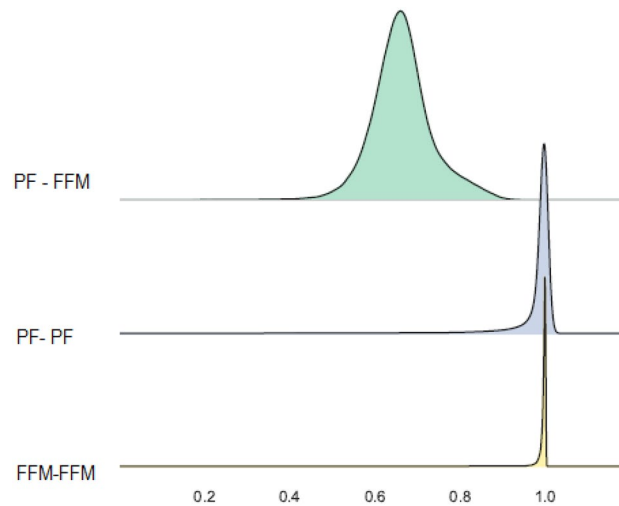


Figure 3. Distribution, over participants, of the multidimensional similarity between the FFM and PF representations.

of the multitrait-multimethod matrix. In a second stage, for each T1 and T2 vector, we measured the extent to which participants' FFM scores are similar to their PF score, thereby calculating a magnitude that is analogous to measures of *divergent validity*. Because cosine similarity is sensitive to the sign and order of dimensions, we extracted the maximal possible similarity between the two spaces, providing the most conservative estimate of divergent validity.

As can be seen in Fig. 3, the T1-T2 similarity of the FFM is nearly maximal ($M=0.994$, $SD=0.011$); the T1-T2 similarity of PF is also very high ($M=0.969$, $SD=0.100$). The similarity between the FFM and the PF on both T1 and T2 is much lower ($M=0.730$, $SD=0.111$). The minimal difference between the convergence measures and divergence measures is on the magnitude of Hedge's g of 2.217, clearly representing a substantial divergence between the FFM and PF spaces. In other words, while the PF representation bears some resemblance, it is clearly a different representation.

Discussion. The results of Study 1 provide evidence that a supervised dimensionality reduction method can yield a low-dimensional representation that is simultaneously predictive of a set of psychological outcome variables. We demonstrate that by using a standard personality questionnaire and supervised learning methods, it is possible to improve the overall prediction of a set of 10 important psychological outcomes, even when restricting ourselves to 5 dimensions of personality. RRR allowed us to compress the 100 questions of the personality questionnaire to a new quintet of attributes that optimize prediction across a large set of psychological outcomes. The resulting set of five dimensions differs from the FFM, and has better predictive power on the held-out sample than the classical FFM and an additional comparison benchmark of five dimensions generated using Principal Component Analysis.

A theory of personality should strive to predict humans' thoughts, feelings, and behaviors across different life contexts. Indeed, the representation we discovered in Study 1 was superior to the FFM in terms of its ability to predict a diverse set of psychological outcomes on a set of novel observations. The fact that the same low-dimensional representation was applicable across a set of important outcomes of human psychology suggests that it is a relatively generalizable model, in the sense that it simultaneously applies to several important domains. However, despite the diversity of the outcome measures examined in Study 1, it remains possible that the PF representation is only effective for the prediction of the set of outcome measures on which it was trained. Such a finding would not negate the usefulness of this model, given the wide variety of outcomes captured by the PF. However, it is interesting to see whether the resulting representation can improve prediction on additional sets of outcomes. In light of this, in Study 2 we sought to examine the performance of the PF on a set of novel outcome measures that were present in the myPersonality database, but that were held-out from the model generation process. Specifically, in this study we sought to see whether the PF representation outperforms the FFM in its ability to predict participants' experiences during their *childhood*.

Unlike the outcome measures used in Study 1, this dependent variable does not pertain to participants' lives in the present, rather, it is a measure of their past experiences. As such, "retrodiction" of remote history may be especially challenging. Nonetheless, it is widely held that individuals' psychological properties are shaped, at least to some extent, by the degree to which they were raised in a loving household^{36,37}. Indeed, there is evidence to the fact that many specific psychological attributes are shaped by experiences with primary caregivers (e.g., shared environmental effects on topics such as food preference³⁸, substance abuse³⁹, and aggression⁴⁰). In light of this, we reasoned that it is reasonable to expect that one's personality profile should contain information that is predictive of individuals' retrospective reports of their upbringing.

Study 2

Method. *Participants.* We used data from 3869 participants who answered all of the questions on the 100-item IPIP representation of²⁶ markers for the Big Five factor structure, and answered the short form “My Memories of Upbringing” (EMBU) questionnaire⁴¹.

Measures. The short form of the EMBU includes a total of six subscales: three subscales that contain questions to measure the extent to which the participants’ father was a *warm, rejecting, and overprotecting* parent, and three subscales that measure the extent to which the participants’ mother was *warm, rejecting, and overprotecting*.

Results. As can be seen in Table 1, for all six variables, prediction accuracy was relatively low; however, importantly, in all six cases the PF-based model outperformed the FFM-based model, and was significantly better for four out of the six outcome variables. The average improvement across the six outcome measures was 49.2%.

Discussion. The results of Study 2 further support the idea that the PF representation that was built using the 10 meaningful outcome measures present in the myPersonality database is at least somewhat generalizable. However, Study 2 again relied on myPersonality participants, upon which the PF was built. In light of this, in Study 3 we sought to further test the generality of the PF by examining whether it outperforms the FFM-based model on a set of new participants. Furthermore, we wanted to see whether our model can outperform the FFM-based model on a set of new outcome measures selected by an independent group of professional psychologists, blind to our model-generation procedure.

Study 3

Method. *Participants.* We collected new data using Amazon’s Mechanical Turk (www.MTurk.com). M-Turk is an online marketplace that enables data collection from a diverse workforce who are paid upon successful completion of each task. Our target sample size was 500 participants, which is double the size of what is considered a standard, adequate sample size in individual differences research⁴². In practice, 582 participants participated in the study, 35 of them were omitted for failing attention checks, leaving 547 participants in the final dataset (243 females and 304 men). This number exceed a sample size of 470 participants which provides 95% confidence that a small effect ($\rho = 0.1$) will be estimated with narrow ($w = 0.1$) Corridor of Stability⁴².

Measures. *Dependent variables.* In order to make sure that the PF generalize across different domains of psychological interest, it was important to generate the list of outcome variables in a way that is not biased by our knowledge of the original ten outcome variables on which the PF was designed (i.e., intelligence, well-being, and so on). Therefore, on January 3rd, 2019, we gathered a list of 12 new outcome measures by posting a call on the Facebook group PsychMAP (<https://www.facebook.com/groups/psychmap>) asking researchers: “to name psychological outcome measures that you find interesting, important, and that can be measured on M-Turk using a single questionnaire item on a Likert scale.” Once we arrived at the target number of questions we closed the discussion and stopped collecting additional variables. The 12 items were suggested by eight different psychologists, six of which had a PhD in psychology and five were principal investigators. By using this variable elicitation method, we had no control over the outcome measures, and could be certain that we have gathered a randomly-chosen sample of outcomes that are of interest to psychologists.

This arbitrariness of the outcome generation process (selecting the first 12 outcomes nominated by psychologists, without any consideration of consensus views regarding variable importance)—and the likely low psychometric reliability of single-item measures—can be seen as a limitation of this study. However, our reasoning was that such a situation best approximates the “messiness” of the unexpected, noisy, real-world scenarios wherein prediction may be of interest—and as such, provides a good test of predictive performance of the FFM and PF.

In the M-turk study, participants rated their agreement with 12 statements (1- Strongly Disagree to 7- Strongly Agree). The elicited items were:

- (1) “I care deeply about being a good person at heart”.
- (2) “I value following my heart/intuition over carefully reasoning about problems in my life”.
- (3) “Other people’s pain is very real to me”.
- (4) “It is important to me to have power over other people”.
- (5) “I have always been an honest person”.
- (6) “When someone reveals that s/he is lonely I want to keep my distance from him/her”.
- (7) “Before an important decision, I ask myself what my parents would think”.
- (8) “I have math anxiety”.
- (9) “I am typically very anxious”.
- (10) “I enjoy playing with fire”.
- (11) “I am a hardcore sports fan”.
- (12) “Politically speaking, I consider myself to be very conservative”.

Independent variables. The independent variables were participants’ answers to the 100 questions of the IPIP questionnaire.

Model assessment. Similarly to Study 1, we use a fourfold cross validation scheme in order to assess the predictive performance of the PF on new data set and outcome variables. Next, we compared it to the predictive

performance of the FFM. The validation worked as follows: we had \hat{B} from Study 1, we learned \hat{A} from a part of the new sample (400 ~ participants) and computed the R^2 on the holdout test set (130 ~ participants). In the spirit of the fourfold cross-validation, we iterated this process over the 4-test sets and calculated the average test R^2 for each model.

Results. Similarly to Studies 1–2, the results showed that the predictive performance of the PF was again better than that of the Big Five, although the improvements were more modest (average 30% improvement across the 12 measures). In 5 out of 12 cases, the PF-based model was significantly better than the FFM-based model, and the opposite was true in 2 cases.

Discussion. The out-of-sample R^2 of the two models (PF\Big Five) in Study 3 show a consistent trend with the results presented earlier in Study 1 and Study 2, that is, a somewhat higher percentage of explained variance in the models with the PF as predictors. This improvement observed in Study 3 was more modest than that observed earlier, but is nonetheless non-trivial—given that the set of outcome variables was different from the one the PF representation was trained on, and given that the PF representation was trained on items from questionnaires designed to measure the FFM. As such, the results of Studies 1–3 clearly demonstrate the generalizability of the PF.

A potential criticism of these findings is that the success of the PF model was more prominent on variables that were more similar to the 10 dependent measures upon which the PF was trained. However, it is important to keep in mind that the 12 outcome measures in this study were selected at random by an external group of psychologists. As such, this primarily means that the 10 psychological outcomes used to train the PF indeed provide good coverage of psychological processes that are of interest to psychologists, and thereby, overall, generalize well to novel prediction challenges.

General discussion

In this contribution, we set out to examine the viability of a novel approach to modeling human personality. Unlike the prevailing Five-Factor Model (FFM) of personality, which was developed by relying on *unsupervised* dimensionality reduction techniques (i.e., Factor Analysis), we utilized *supervised* machine learning techniques for dimensionality reduction, using numerous psychologically meaningful outcomes as data labels (e.g., intelligence, well-being, sociability). Whereas the FFM is optimized towards discovering an ontology that explains most of the variance on self-report measures of psychological traits, our new approach devised a low-dimensional representation of human trait statements that is optimized towards *prediction* of life outcomes. Indeed, the results showed that our model, which we term the Predictive Five (PF), provides predictive performance that is better than the one achieved by the FFM in independent validation datasets (Study 1–2), and on a new set of outcome variables, selected independently of the first study (Study 3). The main contribution of the current work is explicating and demonstrating a methodological approach of generating a personality representation. However, the results of this work is a specific representation that is of interest and of potential use in and of itself. We now turn to discuss both our general approach and the resulting representation.

Interpreting the PF. The dimensional structure that emerged when using our supervised-dimensionality reduction approach differed from the FFM. Two dimensions (Dimension 1 and 2) largely reproduced the original FFM factors of Extraversion and Neuroticism. Interestingly, these two dimensions are the ones that were highlighted in early psychological research as the “Big Two” factors of personality (Wiggins, 1966). Dimension 5 was also highly related to an existing FFM dimension, namely, *Openness to Experience*.

The third and fourth dimensions in the model did not correspond to a single FFM trait, but were composed of a mixture of various items. An inspection of the loadings suggests that Dimension 4 is related to some sort of a combative attitude, perhaps captured best by the construct of Dominance^{44–45}. The items that loaded highly on this dimension related to hostility (“Do not sympathize with others”; “Insult people”), a right-wing political orientation (“Do not vote for liberal political candidates”), and an approach-oriented⁴⁶ stance (“Get chores done right away”; “Find it easy to get down to work”).

Like PF Dimension 4, Dimension 3 also seemed to capture approach-oriented characteristics (with high loadings for the items “Get chores done right away” and “Find it easy to get down to work”), however, this dimension differed from Dimension 4 in that it represented a harmony-seeking phenotype⁴⁷. The items highly loaded on this dimension were those associated with low levels of narcissism (“keep in the background”; “do not believe I am better than others”) but with a stable self-worth (“am pleased with myself”). Additional items that were highly loaded on this dimension were those that reflect cooperativity (“concerned with others” and “sympathize with others”).

These two dimensions may seem like dialectical opposites. Indeed, the item “sympathize with others” strongly loaded on both factors, but with a different sign. However, the additional items that strongly loaded on these two dimensions appear to have provided a context that altered the meaning of this item. This is evident in the fact that Dimensions 3 and 4 are not correlated with each other. A possible speculative interpretation is that the two phenotypes captured by Dimensions 3 and 4 can be thought of as two strategies that may have been adaptive throughout human evolution. The first, captured by Dimension 4 seems to represent aggressive traits that may have been especially useful in the context of *inter*-group competition and conflict; the second, captured by Dimension 3, seems to represent traits that may be associated with *intra*-group cooperation and peace.

In general, the interpretability of the PF representation is lower than that of the FFM, with some surprising items loaded together on the same dimension. For example, the two agreeableness items that “do not believe I am better than others” and “respect others” that are strongly correlated with each other were highly loaded onto

Dimension 1 (that is related to introversion), but with opposite signs. To a certain extent, this is a limitation of the predictive approach in psychology. However, such confusing associations may lead us towards identifying novel insights. For example, it is possible that some individuals adopt an irreverent stance towards both self and others, and such a stance could be predictive of various psychological outcomes, and correlated with introversion.

Towards a more predictive science of personality. As noted, the reasons that people seek models of personality are twofold: first, we want models that allow us to understand, discuss and study the differences between people; second, we need these models in order to be able to predict and affect people's choices, feelings and behaviors⁴⁸. Current approaches to personality modeling succeeded on the former, providing highly comprehensible dimensions of individual differences (e.g., we can easily understand and communicate the contents of the dimension of “Neuroticism” by using this sparse semantic label). However, the ability of the FFM to accurately predict outcomes in people's lives is at least somewhat limited^{19–20,49}.

The significance of the current work is that it describes a new approach to modeling human personality, that makes the prediction of behavior an explicit and fundamental goal. Our research shows that supervised dimensionality reduction methods can generate relatively generalizable, low-dimensional models of personality with somewhat improved predictive accuracy. Such an approach could complement the unsupervised dimensionality reduction models that have prevailed for decades in personality research. Moreover, this research can complement attempts to improve the predictive validity of psychology by using non-parsimonious (i.e., facets and item-level) questionnaire-based predictive models⁵⁰.

Aside from providing a general approach for the generation of personality models, the current research also provides a potentially useful instrument for psychologists across different domains of psychological investigation. Our findings suggest that psychologists who are interested in predicting meaningful consequences (e.g., workplace or romantic compatibility) or in optimizing interventions on the basis of individuals' characteristics (e.g., finding out which individuals will best respond to a given therapeutic technique)—may benefit from incorporating the PF dimensions in their predictive models. To facilitate such future research, we provide the R code that calculates the five dimensions based on answers on the freely available IPIP-100 questionnaire (<https://github.com/GalBenY/Predictive-Five>). The use of an existing, open-access, widely-used questionnaire means that researchers can now easily apply the PF coding scheme alongside with the FFM coding scheme to their data, and compare the utility of the two models in their own specific research domains.

One avenue of potential use of the PF representation is in clinical research. The PF showed improved prediction of depression and well-being; moreover, the PF substantially outperformed the FFM in the prediction of two known resilience factors (intelligence and empathy). Specifically, PF Dimension 3 (which, as noted above, seems to represent some harmony-seeking phenotype) significantly contributed to the prediction of all of four outcomes. As such, future work could further investigate the incremental validity of this dimension (and the PF representation more generally) as a global resilience indicator.

Across a set of 28 comparisons, the predictions derived from the PF-based model were significantly better in 15 cases, and significantly worse in 3 cases. The average improvement in R^2 across the 28 outcomes was 37.7%. However, it is important to note that the PF representation described herein is just a first proof of concept of this general approach, and it is likely that future attempts that are untethered to the constraints undertaken in the current study can provide models of greater predictive accuracy. Specifically, in the current research we relied on the IPIP-100, a questionnaire designed by researchers specifically in order to reliability measure the factors of the FFM, and limited ourselves to a five-dimension solution, to allow comparison with the FFMs. The PF representation outperformed the FFM representation despite these constraints. These results provide a very conservative test for the utility of our approach.

Future directions. Future attempts to generate generalizable predictive models will likely produce even stronger predictive performance if they relax the constraint of finding exactly five dimensions and perform dimensionality-reduction based on the raw data used to generate the FFM itself—namely, the long list of trait adjectives that exist in human language, and that were reduced into the five dimensions of the FFM.

For the sake of simplicity comparability to the FFM, the current work employed a linear method for supervised dimensionality reduction. Recent work in machine learning has demonstrated the power of Deep Neural Networks as tools for dimensionality reduction (e.g., language embedding models). In light of this, it is likely that future work that utilizes non-linear methods for supervised dimensionality reduction could generate ever more predictive representations (i.e., “personality embeddings”).

A limitation of the current work is that the PF was trained on a relatively limited set of 10 important life outcomes (e.g., IQ, well-being, etc.). While these outcome measures seem to cover many of the important consequences humans care about (as evident by the predictive performance on Study 3), it is likely that training a PF model on a larger set of outcome variables will improve the coverage and generalizability of future (supervised) personality models. A potential downside of extending the set of outcome measures used for training, is that at some point (e.g., 20, 100 outcomes) it is possible that the “blanket will become too short”: namely, that it will be difficult to find a low-dimensional representation that arrives at satisfactory prediction performance simultaneously across all outcomes. Thus, future research aiming at generating more predictive personality models may need to find a “sweet spot” that allows the model to fit to a sufficiently comprehensive array of target outcomes.

What may be the most important consequence of the current approach is that whereas previous attempts of modeling human personality necessarily limited by their reliance on the subjective products of the human mind (i.e., were predicated on human-made psychological theories, or subjective ratings of trait words), our approach holds the unique potential of generating personality representations that are based on objective inputs.

A final question concerning predictive models of personality is whether we even want to generate such models, given the potential of their misuse. While the current results still show the majority of variance in psychological outcomes remain unexplained—in the era of social networks and commercial genetic testing, the predictive approach to personality modeling could theoretically lead to models that render human behavior highly predictable. Such models give rise to both ethical concerns (e.g., unethical use by governments and private companies, as in the Cambridge-Analytica scandal) and moral qualms (e.g., if behavior becomes highly predictable, what will it mean for notions of free will and personal responsibility?). While these are all valid concerns, we believe that like all other scientific advancements, personality models are tools that can provide a meaningful contribution to human life (e.g., predicting suicide in order to avoid it; predicting which occupation will make a person happiest). As such, the important, inescapable quest towards generating even more effective models that will allow us to predict and intervene in human behavior is only just the beginning.

Data availability

The data for Study 1, 3 and 4 rely on the myPersonality database (www.mypersonality.org) which is an unprecedented big-data repository for psychological research, used in more than a hundred publications. We achieved permission from the owners of the data to use it for the current research—but we do not have their permission to share it for wider use. The data for Study 2 is available upon request. We also share the complete code and the full model with factor loadings (<https://github.com/GalBenY/Predictive-Five>).

Received: 19 December 2021; Accepted: 5 July 2022

Published online: 25 July 2022

References

- Newcomb, T. & Heider, F. The psychology of interpersonal relations. *Am. Sociol. Rev.* **23**, 742 (1958).
- Dweck, C. S. *Self-Theories: Their Role in Motivation, Personality, and Development* (Psychology press, London, 2013).
- Swann, W. B. Jr. Quest for accuracy in person perception: A matter of pragmatics. *Psychol. Rev.* **91**, 457–477 (1984).
- Ægisdóttir, S. *et al.* The meta-analysis of clinical judgment project: Fifty-six years of accumulated research on clinical versus statistical prediction. *Couns. Psychol.* **34**, 341–382 (2006).
- Allport, G. W. & Odbert, H. S. Trait-names: A psycho-lexical study. *Psychol. Monogr.* **47**, i (1936).
- Costa, P. T. Jr. & McCrae, R. R. Personality: Another ‘hidden factor’ is stress research. *Psychol. Inq.* **1**, 22–24 (1990).
- Allik, I. & Allik, I. U. *The Five-Factor Model of Personality Across Cultures* (Springer, Berlin, 2002).
- Benet-Martínez, V. & John, O. P. Los Cinco Grandes across cultures and ethnic groups: Multitrait-multimethod analyses of the Big Five in Spanish and English. *J. Pers. Soc. Psychol.* **75**, 729–750 (1998).
- Laajaj, R. *et al.* Challenges to capture the big five personality traits in non-WEIRD populations. *Sci. Adv.* **5**, eaaw5226 (2019).
- John, O. P. The ‘Big Five’ factor taxonomy: Dimensions of personality in the natural language and in questionnaires. *Handbook of personality: Theory and research* (1990).
- McCrae, R. R. & Costa, P. T. Clinical assessment can benefit from recent advances in personality psychology. *Am. Psychol.* **41**, 1001–1003 (1986).
- John, O. P. & Srivastava, S. The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research* (1999).
- Gilead, M., Trope, Y. & Liberman, N. Above and beyond the concrete: The diverse representational substrates of the predictive brain. *Behav. Brain Sci.* **43**, e121 (2019).
- Chen, L. Curse of dimensionality. *Encycl. Database Syst.* https://doi.org/10.1007/978-1-4614-8265-9_133 (2018).
- Kosinski, M., Stillwell, D. & Graepel, T. Private traits and attributes are predictable from digital records of human behavior. *Proc. Natl. Acad. Sci. USA* **110**, 5802–5805 (2013).
- Confessore, N. Cambridge Analytica and Facebook: The Scandal and the Fallout So Far. *The New York Times* (2018).
- Harari, Y. N. *Homo Deus: A Brief History of Tomorrow* (Random House, London, 2016).
- Hough, L. M. The ‘Big Five’ personality variables-construct confusion: Description versus prediction. *Hum. Perform.* **5**, 139–155 (1992).
- Sibley, C. G., Osborne, D. & Duckitt, J. Personality and political orientation: Meta-analysis and test of a threat-constraint model. *J. Res. Pers.* **46**, 664–677 (2012).
- Morgeson, F. P. *et al.* Are we getting fooled again? Coming to terms with limitations in the use of personality tests for personnel selection. *Pers. Psychol.* **60**, 1029–1049 (2007).
- Gibney, E. The scant science behind Cambridge Analytica’s controversial marketing techniques. *Nature* <https://doi.org/10.1038/d41586-018-03880-4> (2018).
- Matz, S. C., Kosinski, M., Nave, G. & Stillwell, D. J. Psychological targeting as an effective approach to digital mass persuasion. *Proc. Natl. Acad. Sci. USA* **114**, 12714–12719 (2017).
- Roberts, B. W., Kuncel, N. R., Shiner, R., Caspi, A. & Goldberg, L. R. The power of personality: The comparative validity of personality traits, socioeconomic status, and cognitive ability for predicting important life outcomes. *Perspect. Psychol. Sci.* **2**, 313–345 (2007).
- Kosinski, M., Matz, S. C., Gosling, S. D., Popov, V. & Stillwell, D. Facebook as a research tool for the social sciences: Opportunities, challenges, ethical considerations, and practical guidelines. *Am. Psychol.* **70**, 543–556 (2015).
- Stillwell & Kosinski. myPersonality project: Example of successful utilization of online social networks for large-scale social research. *Am. Psychol.*
- Goldberg, L. R. The development of markers for the Big-Five factor structure. *Psychol. Assess.* **4**, 26–42 (1992).
- Stillwell, D. J. & Kosinski, M. myPersonality Project website. (2015).
- Diener, E. D., Emmons, R. A., Larsen, R. J. & Griffin, S. The satisfaction with life scale. *J. Pers. Assess.* **49**, 71–75 (1985).
- Schwartz, S. H. Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. In *Advances in Experimental Social Psychology* Vol. 25 (ed. Zanna, M. P.) 1–65 (Academic Press, New York, 1992).
- Baron-Cohen, S. & Wheelwright, S. The empathy quotient: An investigation of adults with Asperger syndrome or high functioning autism, and normal sex differences. *J. Autism Dev. Disord.* **34**, 163–175 (2004).
- Radloff, L. S. The CES-D Scale: A self-report depression scale for research in the general population. *Appl. Psychol. Meas.* **1**, 385–401 (1977).
- Cucina, J. M., Goldenberg, R. & Vasilopoulos, N. L. Confirmatory factor analysis of the NEO-PI-R equivalent IPIP inventory. *PsycEXTRA Dataset* <https://doi.org/10.1037/e518612013-349> (2005).

33. Chen, L. Sparse Reduced-Rank Regression for Simultaneous Dimension Reduction and Variable Selection in Multivariate Regression.
34. Henseler, J., Ringle, C. M. & Sarstedt, M. A new criterion for assessing discriminant validity in variance-based structural equation modeling. *J. Acad. Mark. Sci.* **43**, 115–135 (2015).
35. Campbell, D. T. & Fiske, D. W. Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychol. Bull.* **56**, 81–105 (1959).
36. Winnicott, D. W. *The Child, the Family, and the Outside World: By DW Winnicott* (Penguin Books, London, 1969).
37. Bowlby, J. *A Secure Base: Clinical Applications of Attachment Theory* (Taylor & Francis, New York, 2005).
38. Smith, A. D. *et al.* Food fussiness and food neophobia share a common etiology in early childhood. *J. Child Psychol. Psychiatry* **58**, 189–196 (2017).
39. Verhulst, B., Neale, M. C. & Kendler, K. S. The heritability of alcohol use disorders: A meta-analysis of twin and adoption studies. *Psychol. Med.* **45**, 1061–1072 (2015).
40. Chen, J., Yu, J., Zhang, J., Li, X. & McGue, M. Investigating genetic and environmental contributions to adolescent externalizing behavior in a collectivistic culture: A multi-informant twin study. *Psychol. Med.* **45**, 1989–1997 (2015).
41. Arrindell, W. A. *et al.* The development of a short form of the EMBU: Its appraisal with students in Greece, Guatemala, Hungary and Italy. *Pers. Individ. Dif.* **27**, 613–628 (1999).
42. Schönbrodt, F. D. & Perugini, M. At what sample size do correlations stabilize?. *J. Res. Pers.* **47**, 609–612 (2013).
43. Murray, H. A. Explorations in personality: a clinical and experimental study of fifty men of college age. **761**, (1938).
44. Jackson, D. N. *Personality research form* (Research psychologists Press, 1965).
45. Pratto, F., Sidanius, J., Stallworth, L. M. & Malle, B. F. Social dominance orientation: A personality variable predicting social and political attitudes. *J. Pers. Soc. Psychol.* **67**, 741–763 (1994).
46. Higgins, E. T., Kruglanski, A. W. & Pierro, A. Regulatory mode: Locomotion and assessment as distinct orientations. In *Advances in Experimental Social Psychology* Vol. 35 (ed. Zanna, M. P.) 293–344 (New York, Elsevier Academic Press, 2003).
47. Leung, K., Koch, P. T. & Lu, L. A dualistic model of harmony and its implications for conflict management in Asia. *Asia Pac. J. Manag.* **19**, 201–220 (2002).
48. Saucier, G. & Srivastava, S. What makes a good structural model of personality? Evaluating the Big Five and alternatives. *APA Handbook Person. Soc. Psychol.* **4**, 283–305 (2015).
49. Salgado, J. F. The big five personality dimensions and counterproductive behaviors. *Int. J. Select. Assess.* **10**, 117–125 (2002).
50. Stewart, R. D., Möttus, R., Seeboth, A., Soto, C. J. & Johnson, W. The finer details? The predictability of life outcomes from Big Five domains, facets, and nuances. *J. Pers.* **90**, 167–182 (2022).

Acknowledgements

The study was supported by ISF grant 1113/18 to M.G.

Author contributions

G.L., J.D.R., and M.G. conducted this work jointly.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-16108-3>.

Correspondence and requests for materials should be addressed to J.R. or M.G.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022