What Predicts Stroop Performance? A Conditional Random Forest Approach

Hans IJzerman
University of Virginia

Thomas Pollet
Leiden University

Charles R. Ebersole
University of Virginia

David Kun
Functional Finances

This manuscript has not yet been published. Cite at own risk.

Abstract

An experimental science relies on solid and replicable results. The last few years have seen a rich discussion on the reliability and validity of psychological science and whether our experimental findings can falsify our existing theoretical models. Yet, concerns have also arisen that this movement may impede new theoretical developments. In this article, we re-analyze the data from a crowdsourced replication project that concluded that lab site did not matter as predictor for Stroop performance, and, therefore, that there were no "hidden moderators" (i.e., context was likely to matter little in predicting the outcome of the Stroop task). The authors challenge this conclusion via a new analytical method– supervised machine learning - that "allows the data to speak." The authors apply this approach to the results from a Stroop task to illustrate the utility of machine learning and to propose moderators for future (confirmatory) testing. The authors discuss differences with some conclusions of the original article, which variables need to be controlled for in future inhibitory control tasks, and why psychologists can use machine learning to find surprising, yet solid, results in their own data.

What Predicts Stroop Performance? A Conditional Random Forest Approach

An experimental science relies on solid and replicable results. The last few years have seen a rich discussion on the reliability and validity of psychological science and whether our experimental findings can falsify our existing theoretical models (Coyne, 2016). Arguably, one of the most important events calling into question the solidity of psychological science was the publication on precognition by Daryl Bem (2011) followed by a paper on data contingent analyses by Simmons, Nelson, and Simonsohn (2011). Following these papers, the times, they have been a-changing: Extensive discussions have been held on power, replicability, pre-registration, and context sensitivity (Bakker, Hartgertink, Wicherts, & Van der Maas, in press; Brandt, IJzerman et al., 2014, Van Bavel, Mende-Siedlecki, Brady, & Reinero, 2016; IJzerman et al., 2016). Though Bem's (2011) paper was perhaps the most prominent, problems seem to be widespread, and these problems are often blamed on the fact that novelty has been championed over solid and replicable findings (Coyne, 2016; Nosek, Spies, & Motyl, 2012). Yet, this begs the question: How can we retain creative and innovative exploration in service of finding out "truths" about human functioning?

Currently, the most popular way to do so is to conduct an experiment, explore the data, write a paper, replicate this study, and, if it fails, update one's theoretical assumptions (and perhaps argue on social media for a while). This approach is incredibly sensitive to post hoc interpretation and is inefficient. Indeed, there are near infinite reasons *why* effects replicate or not and failures to replicate often do little to inform theoretical predictions (but, see Brandt, IJzerman, et al., 2014). A number of recent articles were published with replications, many of which failed to replicate. One example of this is known as "ManyLabs3" (Ebersole et al., 2016). In many of these reports, authors conclude that context matters little, which in ManyLabs3 was concluded due to little heterogeneity across labs (but see IJzerman et al., 2016).

But is it indeed the case that there are no hidden moderators? We challenge this assumption at least for one report (ManyLabs3) showing that experimental context did in fact matter through a novel analytical method that "allows the data to speak" (Jones & Linder, 2015; Yarkoni & Westfall, in press). We prepared this article to show *how* researchers can uncover so-called hidden moderators, but also to demonstrate the utility of machine learning. We thus show how machine learning can be added as a powerful exploratory method to the empirical cycle so as to make the empirical cycle more efficient.

Overfitting and Poor Predictions

What has been the problem in existing empirical approaches in psychological science? Simmons and colleagues (2011) pointed out that flexibility in data analyses often demonstrate effects where none should be detected. This is particularly true if researchers seek to force a model on data that is overly complex and this is one way to create a problem called *overfitting*. Overfitting happens when data are modeled to a specific dataset in an overly complex way of which the hypothesis is then unlikely to be replicated in another dataset (see also Gelman & Loken, 2013). This can cause implausible theoretical models to outperform unbiased models, and such problems are especially noticeable when sample sizes are modest and/or in cases where researchers rely on so-called *researcher degrees of freedom* to understand their data (John, Loewenstein, & Prelec, 2012; but see Fiedler & Schwarz, 2016). These problems are exacerbated when researchers themselves have biased theoretical models (Fiedler, 2000), when researchers are biased in selecting their variables (Fiedler, 2000; Westfall, Judd, & Kenny, 2015), and/or when they apply the wrong model to the data (e.g., a linear model where a nonlinear pattern is more appropriate; for a rich discussion on the topic in ecology, see Ginzburg & Jensen, 2004). In short, when a researcher uses a biased model and *p*-hacks to significant results, he/she will be unlikely to replicate that result in a new dataset because of – at least in part - overfitting.

The problem of overfitting in psychological science is severe and is likely one important reason for relatively low replication rates (for a more technical discussion on overfitting and for the severity of the problem, see Westfall & Yarkoni, in press). One potent solution to reduce the severity of this problem is by adding potent exploratory data analyses, like supervised machine learning, to the empirical cycle. Supervised machine learning is an approach that generates solutions from data, and *explicitly relies on flexibility* in data analysis to specify predictors. It is a very powerful exploratory approach that helps formulate *predictions* for out-of-sample testing. Such predictions can then be used to test predictions in confirmatory studies that follow up on these exploratory data analyses.

A very popular and widely tested approach to supervised machine learning is a method called "random forest" (Breiman, 2001). It is a so-called "ensemble method" which relies on "decision trees" as base learners. This method allows for measuring the *relative* importance of one specific variable, considering for all the variables used in a dataset to predict the outcome, wielding a highly predictive accuracy, as it repeatedly samples the variables used to predict the outcome. The short explanation of this process is that a series of splitting points are created for variables (e.g., the variable is larger or smaller than 1, the values of categorical variables, and so forth; Strobl, Malley, & Tutz, 2009). The order in which these decisions are taken is represented by the "levels" of the tree, captured in a "permutation importance list" to summarize the relative importance of each individual predictor to the next (for more technical discussions see Breiman, 2001; Jones & Lindner, 2015; Yarkoni & Westfall, in press).

As the name implies, the "forest" consists of many such "trees". The method relies on "out of bag estimates" (bagging), which involves repeated sampling to form training datasets from an original dataset (Breiman, 1996; Bylander & Hanzlik, 1999). The rest of the datasets in each case (the test datasets) are used to evaluate the prediction power of the variable importance and trees trained on the training dataset. The forest is then aggregated with each

tree getting a "vote", which constitutes a weight in the ensembled model that summarizes all information from the trees. The type of machine learning we used (*conditional* random forest) improves its predictive power throughout each iteration of the analyses.

Supervised machine learning differs from unsupervised machine learning in that the data patterns are derived by a "supervisory signal" (an outcome variable). In unsupervised machine learning, the algorithm infers a hidden function or pattern from the data, without regard to such a "signal" (which we typically refer to as dependent variable). The type of supervised machine learning we use relies on a regression function, very similar to the correlational analyses psychologists are accustomed to. However, the type of machine learning we used here 1) allows for non-linearity (which standard regressions cannot do without a priori specification), 2) does not presume direction (a positive or negative relationship), 3) has much less problems with collinearity, and 4) is agnostic which type of variable predicts the outcome, thus allowing the researcher to *classify before regressing* onto the signal (i.e., what we typically refer to as dependent variable).

This method thus has the potential to reduce bias in analyses that one may already have and is particularly useful when multiple predictor variables are measured. One potent example comes from political science: Random forests have been applied in political science to identify explanatory variables for respect for human rights. In the "old" days (i.e., before the 2011 Bem paper), one could have gotten away with analyzing the data "simply" using regressions, without pre-registering the hypotheses and relying on flexible data-analyses (i.e., p-hacking), finding that the "right to a fair trial" is an important predictor for respect for human rights (by using a number of tricks to get the results down to significance level). In short, in a biased, small study, researchers could have overfitted their data and have identified the right to a fair trial as an important predictor if that happened to be an intuitive one, or worse, if researchers had an ideological basis to identify this variable as predictor.

In the "new" days (i.e., after the 2011 Bem paper), one could have had a sensible hypothesis, pre-registered, tested the hypothesis, and come up with nothing. With supervised machine learning, one can instead test multiple hypotheses through an iterative process, and not only conclude that variables like having the right to a fair trial and whether a country was a British colony or not do not significantly predict respect for human rights (thereby reducing overfitting), but also that judicial independence or the rule of law in the country *do* (Jones & Lindner, 2015). Notably, supervised machine learning can also help when doing post-hoc analyses in case a pre-registered hypothesis fails (and it is thus wise to be overinclusive in the amount of predictors measured). Random forest in this case do not only reduce overfitting, they also reduced the chance of *underfitting* (i.e., not deducing a model where one *does* exist). This latter part of the empirical process we turn to next.

Random Forests and the Many Labs initiative – The Stroop Effect

How can we apply random forests in psychology? One of the ways to convince fellow researchers of the usefulness of this method is by testing a very popular paradigm. An important benefit is that applying random forests allows the discovery of candidate predictors in 'big data'.[1] Random forests can however already be used on smaller samples (with a greater number of predictors); If samples are larger and more diverse however chances of replicating beyond the given dataset may be larger (indeed, if one tests a population that is sampled from a rather diverse background, the model that comes out is likely more stable across many different populations).

Recently datasets that are considered large in psychological science have been created, mostly as a result of the "replication movement". Most prominently, Klein and colleagues (2014) and Ebersole and colleagues (2016) have adopted a "Many Labs" approach, in which

[1] Big data is not a prerequisite for doing supervised machine learning. Indeed, supervised machine learning is ideal for large p(redictor), small N studies. However, inferences improve (and thus chances for overfitting decrease) once a sample is more diverse. To be clear, and as far as we know, at present, there are no clear sample size recommendations in the literature.

they sought to replicate a variety of psychological studies across labs, with large numbers of participants. In the third iteration of this initiative, Many Labs 3, the well-known Stroop task was included (Stroop, 1935).

During the Stroop task, participants view words written in different colors. Participants are asked to identify the color of the font and this task becomes more difficult when the actual word (e.g., blue) is written in a different color (e.g., a yellow font) than its meaning, as compared to when the word itself is irrelevant for identifying the color (e.g., car) or when the color of the font matches the meaning of the word (e.g., blue printed in blue font). Because psychologists know this task as robust (MacLeod, 1991), but many different predictors have been applied across different studies, it is a great candidate to understand the utility of supervised machine learning in psychological science. Ebersole et al. (2016) found little heterogeneity across labs, and a small effect for the time of the semester (which was the main moderator investigated in the project), which seemed to suggest that context matters little in determining the outcome of Stroop performance.

We suspected that concluding that context did not play a role in Stroop performance was a case of underfitting (i.e., applying too simple of a model to the data). To take one example of underfitting, consider Van Lange, Rinderu, and Bushman (in press), who proposed that people further away from the equator have lower self-control. Using supervised machine learning, IJzerman and colleagues (in press) showed that equator distance was not an important predictor of self-control, concluding thereby that Van Lange et al.'s (in press) conclusion was erroneous. Because their model was too simple for the data (and for real psychological phenomena) their theoretical position was a likely case of *underfitting*.

Past studies *have* identified a number of variables that affect Stroop performance. Several of these were included as individual difference variables in Many Labs 3, such as self-esteem and stress (Rector & Roger, 1997), gender (Bone & Eysenck, 1972; Mekarski,

Cutmore, & Suboski, 1996; but see MacLeod, 1991), age (West, 1999), extraversion (Bone & Eysenck, 1972), agreeableness (Jensen-Campbell et al., 2002), mood (Philips, Bull, Adams, & Fraser, 2002), and temperature (Allan, Gibson, & Green, 1979). Our approach thus does not only permit us to demonstrate the utility of supervised machine learning, it also allows us to weigh important predictors against one another (and thus further understand whether previous empirical papers were overfitted), and identify which of these variables do and do not substantially predict Stroop performance, including whether context mattered for Stroop performance (and thus understand whether Ebersole et al., 2016 underfitted). As far as we know, this is the first study to do so for the Stroop Effect. This thus also allows excluding variables for future research, identifying a first direction for important and more precise follow up predictions, while at the same time identifying important control variables for future work. We will illustrate the latter with a few so-called "partial dependence plots".

Method

Dataset Description

We used the Many Labs 3 dataset to identify predictors of the well-known Stroop Effect (Ebersole et al., 2016). The main focus of Many Labs 3 was to identify whether time of the semester moderated the results of psychology studies. The authors replicated 10 known effects, measured 11 individual difference variables, and measured 3 data quality indicators at 20 different sites ($N = 2,696$) and in one online sample (via mTurk; $N = 737$; Ebersole et al., 2016). The authors took a confirmatory approach, and found seven of the ten effects not to replicate (but see IJzerman et al., 2016). We only address the Stroop Effect (which provided the most solid replication) here, for which the authors only identified a (linear) increase in the Stroop Effect across the semester.

During the version of the Stroop Task in Many Labs 3, participants viewed trials with one of three words (blue, green, or red) printed in blue, green, or red font. After completing

practice trials to orient them to the task, participants completed 63 test trials. Each word was shown 21 times, with 7 trials in each of the three colors. Thus, 2/3 of trials were incongruent with respect to word meaning and font color. The Stroop task was scored using a modified version of the D algorithm for scoring the Implicit Association Task (see Ebersole et al., 2016; Greenwald, Nosek, & Banaji, 2003). A lower score on the D algorithm meant better Stroop performance.

We selected all individual difference variables from the replication effort (global self-esteem, extraversion, openness to experience, conscientiousness, emotional stability/ neuroticism, agreeableness, daily mood, perceived stress, need for cognition, and intrinsic motivation) for inclusion in our analyses.[2] In addition, we took into account participants' gender, year in college, and age, time of semester effect, an attention check, temperature in the lab, the site where participants completed the study (coded as 0 or 1 for each site)[3], and participants' subjective experiences of light, temperature, and space.[4]

Data Handling. Before running the analyses, we inspected the data for anomalies. We removed a value on gender (19), which was not a valid value. In addition, and in accordance with how the authors of the original study requested it, we removed outliers in the subjective temperature estimate (below 50 and above 90 degrees Fahrenheit). We also removed one outlier on the variable age (13), and two improbable values (-2 and 149 years old). We also

---

[2] A full list of variables from Many Labs 3 can be found here: osf.io/rjfkb/. We excluded variables that were asked as dependent variables for other replications (i.e., by design were not relevant for our study), with the exception of subjective experiences of the lab (which were included for the Szymkow, Chandler, IJzerman, Parzuchowski, & Wojciscszke, 2013, replication).

[3] We coded site as 0 and 1, because the specific R package cannot handle variables with more than 4 levels (while there were 19 sites).

[4] When cleaning the data, choices still remain to be made before one does the analyses (which should be done independently from the analyses), with such logical choices for not including a participant when he/she was drunk, or did not follow instructions by the experimenter. In our case, choices were made by the original paper's authors. For conditional random forests, choices remain to be made. For example, one needs to examine correlations between the different variables before running one's random forest. If correlations are too high, variables may need to be combined or deleted, based on one's theoretical predictions *a priori* (we did not do so here).

checked at what level our predictor variables were correlated. Given that all variables were correlated $r < .48$ (highest one being neuroticism and stress), we decided not to merge any of our predictor variables. Finally, for the comparison of year in college we only maintained the values for freshman through senior (1 through 4), and included them as a categorical variable in our design. All of our data exclusions are described here; we did not exclude any other participants.

Analysis Strategy. For creating the classification trees, we relied on the packages party (Hothorn, Hornik, Strobl, & Zeileis, 2010). We also used the R packages tree (Ripley, 2012), lattice (Sarkar, 2008), plyr (Wickham, 2009), stargazer (Hlavac, 2013), machine learning (Bischl et al., 2016), ParamHelpers (Bischl, Lang, Bossek, Horn, & Schork, 2016), and summarytools (Comtois, 2015). Our analysis scripts are available at the Open Science Framework (https://osf.io/26kr7/).[5]

Explanation of Conditional Random Forest. In the algorithm, the researcher specifies a number of different parameters: "Mtry", "Ntree", and "Seed". Mtry is the parameter that specifies how many predictor variables the algorithm samples at each tree; its value is recommended to be the square root of the total amount of variables (in our case 4; Latinne, Debeir, & Decaestecker, 2001). Ntree is the parameter that specifies how many times the algorithm samples these variables (Ntree is thus the number of trees grown in the forests). The number of trees that are required in part depends on the stability of the model (which can be inferred from the Spearman Rank correlation between the different seeds) as well as computation time. As a general rule of thumb, it is recommended to take Ntree = 1,000, though some authors have shown that Ntree > 200 can already lead to stable results (Latinne et al., 2001). For the demonstration in this paper, one of the versions of our Ntree is set at

_____

[5] When we first completed the random forest, the variable importance was slightly different. Due to a coding error we discovered that our model was not as stable as we would have liked it to be. Because we think this process is useful for others starting to use this method, we have described this process on our OSF page.

10,000[6]; the remainder at 1,000. The seed is a necessary number to be able to reproduce the results. Supervised machine learning randomly samples the data: A "seed" provides a starting point from which it grows the trees in the forests. When using conditional random forests, it is wise to sample a number of seeds (at least 2, but preferably 4) to understand how well the model replicates and to calculate the Spearman Rank Rho of the different rank orders of the predictors to examine stability of the model (we used seeds 1, 2, 3, and 4). At present, there are no guidelines specifying which Spearman rank Rho is sufficient; as a (rather liberal) rule of thumb we recommend a Spearman rank Rho that is at minimum .7.

Variable Selection

Researchers should first seek to produce a stable model following the steps we have outlined above. Once a stable model has been obtained, variables can be selected from a so-called "variable importance list", which is composed of a conditional importance compared to the other variables in the model. This variable importance list ranks the most important predictors of the model from high to low. The number outcome below the graph should never be taken understood as absolute, neither should it be compared as an effect size to other data.

But which variables matter in predicting one's effect? This has proven challenging in running random forests. In random forests, researchers try to eliminate noise from variables that could be important for the signal. Whereas simple significance tests were proposed initially to achieve this, a "backward elimination strategy" was proposed more recently (for a summary, see Strobl, Boulesteix, Kneib, Augustin, & Zeileis, 2008). In a backward elimination strategy, researchers a priori select a fixed number of variables (e.g., 20%) that are at the end of the list of the "variable importance list" and are dropped for a subsequent run of the model from which the relevant predictor variables are selected (i.e., ones that are not

---

[6] The main reason for not growing the trees to too large of a number in the forest is simply computing power. In our case, we had the availability over cloud computing and a far larger number of trees was possible. When testing the algorithm, we recommend initially setting the number of trees at 1; when actually testing the model, we recommend setting the number of trees at 1,000.

noise; see e.g., Genuer, Poggi, & Tuleau-Malot, 2010, for an example). Strobl et al. (2008) instead proposed a more consistent strategy through a conditional importance scheme based on the variables themselves fitted by the random forest model, which produces a threshold within the model itself that has been shown to be more consistent in variable selection to separate noise from potential predictor variables.

Results: Identifying Stroop Predictors

Below we provide the classification trees for the Stroop task. We ran ten different versions of our model in total, to ascertain its stability: Five versions without the mTurk data (all with 1k trees and seeds 1, 2, 3, and 4, and mtry = 4; once with 10k trees and seed 1), and five versions with the mTurk data included (all with 1k trees, seeds 1, 2, 3, and 4, mtry = 4; once with 10k trees and seed 1). Both the models with and without mTurk were stable (Spearman rank Rho with mTurk = .75; lowest Spearman rank Rho without mTurk = .74 meaning that the rank order of the variables, and thus which variables came out as important, was comparable across analyses). All runs and details of the runs can be found on our OSF page (osf.io/q9d56/).

Here we report (permutation) variable importance over the seed that we ran with the most trees with mTurk included (i.e., full dataset; mtry=4, trees=10k, seed=1) and mTurk excluded (mtry=4, trees=10k, seed=1). The outcomes of the other seeds and 1k forests are reported on our project page on the OSF. The variable importance simply indicates to what extent a variable "matters" more than noise and it provides a relative ranking between noise and potential predictor variables. The red line in the graph then indicates to what extend the predictors exceed noise. We report both the variable importance for the full dataset and the dataset excluding the mTurk data.

Full Dataset (mTurk included). For the full dataset, in the original dataset, the variables that

exceeded the Strobl et al.'s (2008) noise threshold (the red line) are ranked as follows (see

Figure 1):

Age > Time of Semester > Temperature in Lab > Openness to Experience > Need for

Cognition > Subjective Temperature > Intrinsic Motivation > Gender > Subjective Space

mTurk Excluded.

For the dataset without mTurkers, the variable importance of the variables that exceeded

Strobl et al.'s (2008) red line is as follows (see also Figure 2):

Openness to Experience > Temperature in Lab

From two analyses with two different versions of the dataset, stable predictors emerge.

Openness to Experience and Lab Temperature were both important predictors for both

datasets. Age, Time of Semester, Need for Cognition, Subjective Temperature, Intrinsic

Motivation, Gender, and Subjective Space were also significant predictors for the entire

dataset (i.e., with mTurk included). Other known predictors (such as self-esteem, stress,

extraversion, agreeableness, or mood) did not play a role in predicting Stroop performance in

the Many Labs 3 data.

Plotting the Partial Dependencies

Now that we know that a number of variables are stable predictors of Stroop

performance in the Many Labs 3 dataset, we can start identifying their unique relationships

with the outcome variable. Because we interpret the data including mTurk to be somewhat

more heterogeneous, we plot the partial dependencies for the data excluding mTurk and only

interpret these data (we included a variable that did not exceed the threshold, as it was

interpreted by Ebersole et al., 2016, namely time of semester). By relying on partial

dependency plots, variables are plotted by keeping all other variables at constants and thus

providing a great amount of control over variance through the random forest (see Figure 3).

We again graph the data also for including mTurk (see Figure 4). It is clear that participants above 60 performed significantly poorer, which is why age became a more important predictor in the dataset when including mTurkers.

A number of inferences can be drawn from these plots, and we simply describe the relationships here. First, all relationships appear to be nonlinear. Second, people higher in Openness to Experience perform better on the Stroop task. Furthermore, under lower lab temperatures (67.5 degrees Fahrenheit and below) participants perform better. This performance deteriorates up until 75 degrees Fahrenheit, after which performance remains comparable.[7]

To what degree do random forests exceed random noise?

In the graphs that preceded we included a threshold (the red line) to determine which variables matter in predicting the signal (Stroop performance). This threshold is a moderate threshold suggested by Strobl et al. (2008) as an "embedded" variable selection strategy within the model based on mathematical proof. The value on which this is based is a relative value deduced from the model itself.

Strobl et al.'s (2008) threshold is a moderate threshold based on mathematical proof; we wanted to explore how this threshold compares against random noise, which is important for psychologists working with participant data. We thus decided to replace our dependent variable with noise and compare that with our real data. The distribution of the dependent variable was established as Gaussian; we drew a random sample from a Gaussian distribution with the same mean and standard deviation as the observed Stroop effect and kept the predictors unmodified, used the same seeds (1, 2, 3, 4), the same Ntrees (1k), and the same Mtry (4). Modeling such random noise showed that - for our "mTurk random noise" - in all of

_____

[7] Time of semester became only an important predictor with all participants included, but in the dataset without mTurkers, performance quickly deteriorated. Given that it only became a relevant predictor when mTurkers were included, this variable should be interpreted with some care, and it is likely that Ebersole et al. (2016) were not correct in concluding that this variable mattered for a student sample.

our 4 random forests with 1k trees, a model emerged (with 7, 5, 1, 2 variables exceeding the Strobl et al.'s, 2009, threshold, respectively; without mTurk: 1, 1, 0, 0). However – and importantly – the maximum Spearman Rank Rho's correlation between 4 seeds with 1k trees ("including mTurk") was .21 (excluding mTurk the maximum was .02), thus producing a very instable model. We again provide one example forest of random noise in Figure 5 (the remainder – also for "excluding mTurk" are provided on our project page).

The information from our random noise models thus provides valuable information, in part because some of the variables did exceed Strobl et al.'s noise threshold. We added a number of extra thresholds beyond Strobl et al.'s (2009) guidelines based on real random noise. Specifically, we included the mean of each graph's mean variable importance (in blue). We also included a line with the average of the four highest values of variable importance of the random data (yellow line), which represents a somewhat more stringent threshold. Finally, we included an extreme threshold, with the highest value of the variable importance in all of the seeds from the random data (green line, see Figures 6 and 7 for our revised graphs Figure 1 and 2). Note that on the basis of our blue line, the relevant predictor variables with mTurkers included remain the same; without mTurkers, relevant predictor variables (and its order) is now the following:

Openness to Experience > Temperature in Lab > Time of Semester > Age > Subjective Temperature > Year in College

Just on the basis of visual inspection of our 1k forests, our blue line seems to more consistently incorporate the same variables than the original red line, as the same variables across seeds more consistently exceed the blue (versus the red) threshold (we plotted all the mTurk graphs with 1k trees in Figure 8 for a side-by-side visual inspection). Whether this consistency is true of conditional random forests in general requires more testing with other datasets. Nevertheless, we suspect that our blue line is a more appropriate threshold to be used

to select variables for further confirmatory testing. Furthermore, our random noise thresholds is not yet based on the same proof that Strobl et al. (2008, 2009) presented. We suspect that the yellow and green lines can be used in case when researchers need to make very precise recommendations (as in the case for example in applied research when making policy recommendations), but they still require further examination. How these guidelines will be applied exactly will in large part depend on future research that more accurately discovers the connection between such exploratory and follow up confirmatory work.

Discussion

Psychological science finds itself in a challenging, but interesting period. In the last few years, researchers have recognized the limitation of experimental approaches and the post hoc interpretations that typically went along with them (Brandt, IJzerman et al., 2014; Coyne, 2016; Simmons et al., 2011). At the same time, psychologists are left to wonder how innovation and exploration can be done in a reliable way. Previous scholars have suggested favoring prediction over explanation (Yarkoni & Westfall, in press), and we followed suit and applied a conditional random forest approach to one of the most replicable and well-known paradigms in psychological science: The Stroop Effect. Despite our solid approach, we found effects that may be counterintuitive and that ran against some of the conclusions of Many Labs 3. A key conclusion is that hidden moderators matter more than previously assumed. This did not only come forward in our random forest approach but also when we compared analyses of datasets with and without mTurk participants.

However, these analyses also provided evidence against previous studies that have claimed main effects. For example, some studies have found self-esteem and stress (Rector & Roger, 1997), age (West, 1999), extraversion (Bone & Eysenck, 1972), agreeableness (Jensen-Campbell et al., 2002), mood (Philips, Bull, Adams, & Fraser, 2002), and temperature (Allan, Gibson, & Green, 1979) to matter in Stroop performance. Out of these variables, only

temperature emerged as a relevant and stable predictor in our models. Furthermore, the Many Labs 3 data showed that openness to experience was a reliable predictor of Stroop performance. This is not to say that the other variables do not predict Stroop performance at all. Indeed, it could well be that a variable like individual differences in rumination could be a hidden moderator, interacting with stress to produce Stroop performance. Our approach thus allows for a more directed approach to uncover such hidden moderators. It is clear however that a main effect of stress does *not* play a role in our data.

What was further clear is that including the mTurk sample changed the predictors in the model and this is likely due to the sample becoming more heterogeneous (when comparing the two different datasets, the lowest rank order correlation dropped to Spearman rho =.62). Although Ebersole et al. (2016) inferred that site differences were *not* a relevant predictor, our analyses suggests otherwise: Including a broader age sample with mTurk participants included places "age" at the top of the list for the entire dataset as a predictor of Stroop performance and includes many more predictors exceeding a random noise threshold (but see our added analyses with our own threshold).

Lab temperature and openness to experience were both clear predictors of Stroop performance. Because some labs were significantly warmer or colder than other labs (ranging in $M_{temperature} = 69.06$ for the coldest lab to $M_{temperature} = 77.89$ for the warmest lab in the sample), and because some labs varied in the degree to which their participants scored higher or lower on openness to experience (ranging in $M_{openness} = 4.94$ to $M_{openess} = 5.53$), Stroop performance should vary across these labs. Instead of simply comparing across sites (which also did not turn up as predictor in our model), identifying theoretically relevant variables through supervised machine learning thus helps understanding *what type of context* matters.

We need to be clear. Our supervised machine learning does not *replace* the existing empirical cycle. Instead, it helps understand 1) which variables differ from random noise in a

dataset, thereby 2) tightening the logic to set up more solid experimental work, 3) widen our horizon in terms of which variables we should take into account, and 4) exclude variables that are *not* theoretically relevant and are simply noise in a dataset. Random forests thus provide one step in reducing overfitting, but hypothesis-driven research remains necessary to separate real predictors from noise. In recent work, Munafò et al. (2017) discuss threats to the empirical cycle (see also Figure 4), which are caused by p-hacking, failure to control for bias, poor study control, and low statistical power (which ultimately led to an overfitted literature). By using supervised machine learning, one has the potential to *reduce* bias in the literature by reducing overfitting, and improve predictive power from one study to the next by understanding which context variables matter (and thus reducing underfitting). Still, analyses, data, and analytical choices should be open so that theoretical progress is facilitated.

In our example, the logic in the empirical cycle is tightened because we were able to exclude some variables, and point to the importance of others. On the basis of our partial dependency plots, future experimental studies can be designed in much more accurate ways than has been done before. Think of lab temperature. Although previous research has already indicated temperature to be of crucial importance to basic cognitive processes (e.g., Halali, Meiran, & Shalev, 2016; IJzerman & Semin, 2009; IJzerman & Hogerzeil, 2017), very few Stroop tasks have controlled for lab temperature in their design. With our conditional random forest approach, it becomes clearer and clearer how important temperature is for basic cognitive functioning (cf. Hancock, 1980; IJzerman et al., 2015). Specifically, future studies could be designed comparing environments varying from 68 to about 74 degrees Fahrenheit. Whether temperature is specific for Americans' (and not Europeans' or Asians') Stroop performance is unlikely, but remains to be seen. Our analyses provide a helpful starting point.

Our findings are not just relevant to design studies with the most central predictors in mind. It is also advisable for future studies that use the Stroop task to control for the

predictors that emerged from our model. Although openness to experience and lab temperature emerged as stable (and significant) predictors from the model, we advise researchers to be somewhat more conservative and aim to control for the top 6 predictors of our model (including Time of Semester, Age, Subjective Temperature, and Year in College). Indeed, like NHST, our cutoff line is an arbitrary choice (although our random noise approach *seem* to have made our conclusions slightly more comparable across models).

Our analyses suggest nonlinear relationships and these nonlinear relationships should be used as a starting point to control for in future studies (e.g., by taking a very specific cutoff point for openness to experience or lab temperature). Further, while our findings might just be limited to research in the United States, at least it helps to be a useful departure point for future studies elsewhere. We thus advise labs outside of the United States to start implementing the controls we specified. In a healthy empirical cycle, labs outside of the United States can then always falsify our proposed generalizability beyond the US (for famous failures, see Henrich, Heine, & Norenzayan, 2010). Given that the logic of the Stroop task is often applied to other interference tasks (like the N-back task or Go/No Go task) we suspect it would be advisable to start reporting the same controls for those tasks.

One of the surprising conclusions for some may be that variables like mood, stress, or self-esteem did not reliably predict Stroop performance. It is possible that when such relationships have been detected in the past, the relationship between these variables and Stroop performance was overfitted (or, as others have called it, p-hacked; Simmons et al., 2011). It is very well possible that such variables interact with other variables (e.g., tendency to ruminate) to predict Stroop performance, but main effects are very unlikely. It is also not at all unlikely that other variables, like openness to experience or lab temperature, were the actual drivers of Stroop performance, rather than mood or stress. What is less obvious however from a random forest approach is what the theoretical mechanism is, or why a

predictor failed. For example, it is unclear why exactly people who are more open to new experiences perform better at the Stroop task. It may be that they get more absorbed in the task; such a relationship can be explored in future studies. Not only do we not have enough data for this, supervised machine learning by no means replaces existing approaches, such as confirmatory studies or 'causal' approaches (such as SEM).

Furthermore, they by no means replace solid theoretical reasoning. Our analyses simply provide a starting point to know with which variables to proceed (for one example of an approach starting with machine learning and proceeding into a mediation model, see IJzerman et al., 2017). What random forests *do* show, however, is whether a variable predicts the outcome measure – Stroop performance in this case – or not in a given dataset. Indeed, although we tentatively conclude that many individual difference variables do not predict Stroop performance, problems with measurement could still play a significant role. Though perhaps less likely for mood, self-esteem or stress, variables like extraversion and/or agreeableness might not have emerged as predictors because they were measured with the 10-item Big Five measurement (Gosling, Rentfrow, & Swann, 2003), rather than with a more complete measure. Other Big Five measures could still reliably predict Stroop performance differently and this can be subjected to future random forest approaches. Whether the latter is a measurement problem or not, it does not matter so much for the utility of random forests. Indeed, at least now it is very clear that neither agreeableness nor extraversion from the Gosling et al. (2003) measure predict Stroop performance in our student population.

Although our analyses provide just a starting point and much remains to be learned, it is clear that with a relatively simple additional analysis we have gained a considerable amount of knowledge about Stroop performance. Our random forests have provided a much more powerful estimation of how Stroop performance is predicted compared to classical analysis approaches, such as the correlational approach used in the original Many Labs 3 manuscript,

and it provided us with some counterintuitive answers. An exploratory approach like random forest is extremely powerful, it is reliable, and adds a considerable amount of information to the existing empirical cycle. And perhaps even more importantly, with our script and explanation from our project page at the Open Science Framework (https://osf.io/x53pu/wiki/home/) supervised machine learning comes within reach of every social psychologist.

References

Allan, J. R., Gibson, T. M., & Green, R. G. (1979). Effect of induced cyclic changes of deep
body temperature on task performances. *Aviation, Space, and Environmental
Medicine*, *50*, 585–589.

Bakker, M., Hartgerink, C. H., Wicherts, J. M., & van der Maas, H. L. (in press).
Researchers' intuitions about power in psychological research. *Psychological Science,
27*, 1069-1077.

Bem, D. J. (2011). Feeling the future: experimental evidence for anomalous retroactive
influences on cognition and affect. *Journal of Personality and Social Psychology*, *100*,
407–425.

Bischl, B., Lang, M., Bossek, J., Horn, D., & Schork, K. (2016). ParamHelpers: Helpers for
parameters in black-box optimization, tuning, and machine learning. Retrieved from
https://cran.r-project.org/web/packages/ParamHelpers/index.html.

Bischl, B., Lang, M., Richter, J., Bossek, J., Judt, L., Kuehn, T., …, & Schiffner, J. (2016).
mlr: Machine learning in R. Retrieved from
https://cran.r-   project.org/web/packages/mlr/.

Brandt, M. J., IJzerman, H., Dijksterhuis, A., Farach, F. J., Geller, J., Giner-Sorolla, R., ... &
Van't Veer, A. (2014). The Replication Recipe: What makes for a convincing
replication? *Journal of Experimental Social Psychology*, *50*, 217-224.

Bone, R. N., & Eysenck, H. J. (1972). Extraversion, field-dependence, and the Stroop test.
Perceptual and Motor Skills, *34*, 873–874.

Breiman, L. (2001). Random forests. *Machine Learning, 45,* 5-32.

Bylander, T., & Hanzlik, D. (1999). Estimating generalization error using out-of-bag
estimates. In *AAAI/IAAI* (pp. 321-327).

Comtois, D. (2015). Dataframe summaries, frequency tables and numerical summaries with

  customizable output. Retrieved from

  https://cran.r-project.org/web/packages/summarytools/summarytools.pdf.

Coyne, J. C. (2016). Replication initiatives will not salvage the trustworthiness of

psychology. *BMC Psychology*, *4*, 1-11.

Ebersole, C. R., Atherton, O. E., Belanger, A. L., Skulborstad, H. M., Adams, R. B., Allen, J.,

  … & Nosek, B. A. (2016). Many Labs 3: Evaluating participant pool quality across the

  academic semester via replication. *Journal of Experimental Social Psychology, 67*,

  68-82.

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to

  judgment biases. *Psychological Review*, *107*, 659-676.

Fiedler, K., & Schwarz, N. (2016). Questionable Research Practices revisited. *Social

  Psychological and Personality Science*, *7*, 45-52.

Gelman, A., Loken, E., 2013. The garden of forking paths: Why multiple comparisons can be

  a problem, even when there is no "fishing expedition"or "p-hacking"and the research

  hypothesis was posited ahead of time [WWW Document]. URL http://

  www.stat.columbia.edu/~gelman/research/unpublished/p_hacking.pdf.

Genauer, R., Poggi, J.-M., Tauleau-Malot, C. (2010). Variable selection using Random

  Forests. *Pattern Recognition Letters*, *31*, 2225-2236.

Ginzburg, L. R., & Jensen, C. X. (2004). Rules of thumb for judging ecological

  theories. *Trends in Ecology & Evolution*, *19*, 121-126.

Gosling, S. D., Rentfrow, P. J., & Swann Jr, W. B. (2003). A very brief measure of the Big-

  Five personality domains. *Journal of Research in Personality*, *37*, 504-528.

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, *85*, 197-216.

Halali, E., Meiran, N., & Shalev, I. (2016). Keep it cool: temperature priming effect on cognitive control. *Psychological Research*, *81,* 1-12.

Hancock, P. A. (1986). Sustained attention under thermal stress. *Psychological Bulletin*, *99*, 263 – 281.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences, 33*, 61-135.

Hlavac, M. (2013). Stargazer: LaTeX code and ASCII text for well-formatted regression and summary statistics tables. Retrieved from *http://CRAN. R-project. org/package= stargazer*.

Hothorn, T., Hornik, K., Strobl, C., & Zeileis, A. (2010). Party: A laboratory for recursive partitioning. Retrieved from https://cran.r-project.org/web/packages/party/party.pdf.

IJzerman, H., & Hogerzeil, L. (forthcoming). People as penguins: Thermoregulation as part of the human essence. In M. van Zomeren & J. Dovidio (Eds.). *The Handbook of Human Essence*. Oxford University Press. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2844963.

IJzerman, H., & Semin, G. R. (2009). The thermometer of social relations: Mapping social proximity on temperature. *Psychological Science, 10,* 1214–1220.

IJzerman, H., Szymkow, A., & Parzuchowski, M. (2016). Warmer hearts, and warmer, but noisier rooms: Communality does elicit warmth, but only for those in colder ambient temperatures–Commentary on Ebersole et al. (2016). *Journal of Experimental Social Psychology.*

IJzerman, H., Coan, J. A., Wagemans, F. M., Missler, M. A., van Beest, I., Lindenberg, S., & Tops, M. (2015). A theory of social thermoregulation in human primates. *Frontiers in Psychology*, *6*.

IJzerman, H., Čolić, M., Hennecke, M., Hong, Y., Hu, C. P., Joy-Gaba, J., ... & Schubert, T. W. (in press). Does distance from the equator predict self-control? Lessons from the Human Penguin Project. *Behavioral and Brain Sciences*. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2821804.

IJzerman, H., Lindenberg, S. M., Dalgar, I., Weissgerber, S. C., Vergara, R., Cairo, A., … & Zickfeld, J. (2017). The Human Penguin Project: Social integration protects against cold climates. Available at PsyArxiv: https://osf.io/preprints/psyarxiv/6b7ne.

Jensen-Campbell, L. A., Rosselli, M., Workman, K. A., Santisi, M., Rios, J. D., & Bojan, D. (2002). Agreeableness, conscientiousness, and effortful control processes. *Journal of Research in Personality*, *36*, 476-489.

John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of Questionable Research Practices with incentives for truth telling. *Psychological Science, 23*, 524 – 532.

Jones, Z., & Linder, F. (2015, April). Exploratory data analysis using random forests. *Prepared for the 73rd annual MPSA conference*.

Klein, R. A., Ratliff, K., Vianello, M., Adams Jr, R. B., Bahnik, S., Bernstein, M. J., ... & Nosek, B. A. (2014). Investigating variation in replicability: A "Many Labs" replication project. *Social Psychology*, *45*, 142-152.

Latinne, P., Debeir, O., & Decaestecker, C. (2001). Limiting the number of trees in random forests: In J. Kittler & F. Roli (Eds.). *Multiple Classifier Systems*. *MCI 2001*. *Lecture Notes in Computer Science, Vol. 2096*. Springer: Berlin, Heidelberg.

MacLeod, C. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, *109*, 163-203.

Mekarski, J. E., Cutmore, T. R. H., & Suboski, W. (1996). Gender differences during processing of the Stroop task. *Perceptual and Motor Skills*, *83*, 563-568.

Munafò, M. R., Nosek, B. A., Bishop, D. M. V., Button, K. S., Chambers, C. D., Percie du Sert, N. …& Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behavior, 1,* 1–9.

Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific utopia II. Restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, *7*, 615-631.

Phillips, L. H., Bull, R., Adams, E., & Fraser, L. (2002). Positive mood and executive function: Evidence from Stroop and fluency tasks. *Emotion*, *2*, 12-22.

Rector, N. A., & Roger, D. (1997). The stress buffering effects of self-esteem. *Personality and Individual Differences*, *23*, 799-808.

Ripley, B. (2012). Classification and regression trees. R package version 1.0–33.

Sarkar, D. (2008). *Lattice: multivariate data visualization with R*. Springer Science & Business Media.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science, 22,* 1359–1366.

Strobl, C., Malley, J., Tutz, G., (2009). An introduction to recursive partitioning: rationale, application, and characteristics of classification and regression trees, bagging, and random forests. *Psychological Methods*, 14, 323–348.

Strobl, C., Boulesteix, A.-L., Kneib, T., Augustin, T., & Zeileis, T. (2008). Conditional variable importance for random forests. *BMC Bioinformatics, 9,* 307.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643–662.

Szymkow, A., Chandler, J., IJzerman, H., Parzuchowski, M., & Wojciszke, B. (2013). Warmer hearts, warmer rooms: How positive communal traits increase estimates of ambient temperature. *Social Psychology*, *44*, 167-176.

Van Bavel, J. J., Mende-Siedlecki, P., Brady, W. J., & Reinero, D. A. (2016). Contextual sensitivity in scientific reproducibility. *Proceedings of the National Academy of Sciences, 113*, 6454-6459.

Van Lange, P. A., Rinderu, M. I., & Bushman, B. J. (in press). Aggression and violence around the world: a model of CLimate, Aggression, and Self-control in Humans (CLASH). *Behavioral and Brain Sciences*, 1-63.

West, R. (1999). Age differences in lapses of intention in the Stroop task. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, *54*, P34-P43.

Westfall, J., Judd, C. M., & Kenny, D. A. (2015). Replicating studies in which samples of participants respond to samples of stimuli. *Perspectives on Psychological Science*, *10*, 390-399.

Wickham, H. (2009). plyr: Tools for splitting, applying and combining data. *R package version 0.1*, *9*, 651.

Yarkoni, T., & Westfall, J. (in press). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives in Psychological Science*.
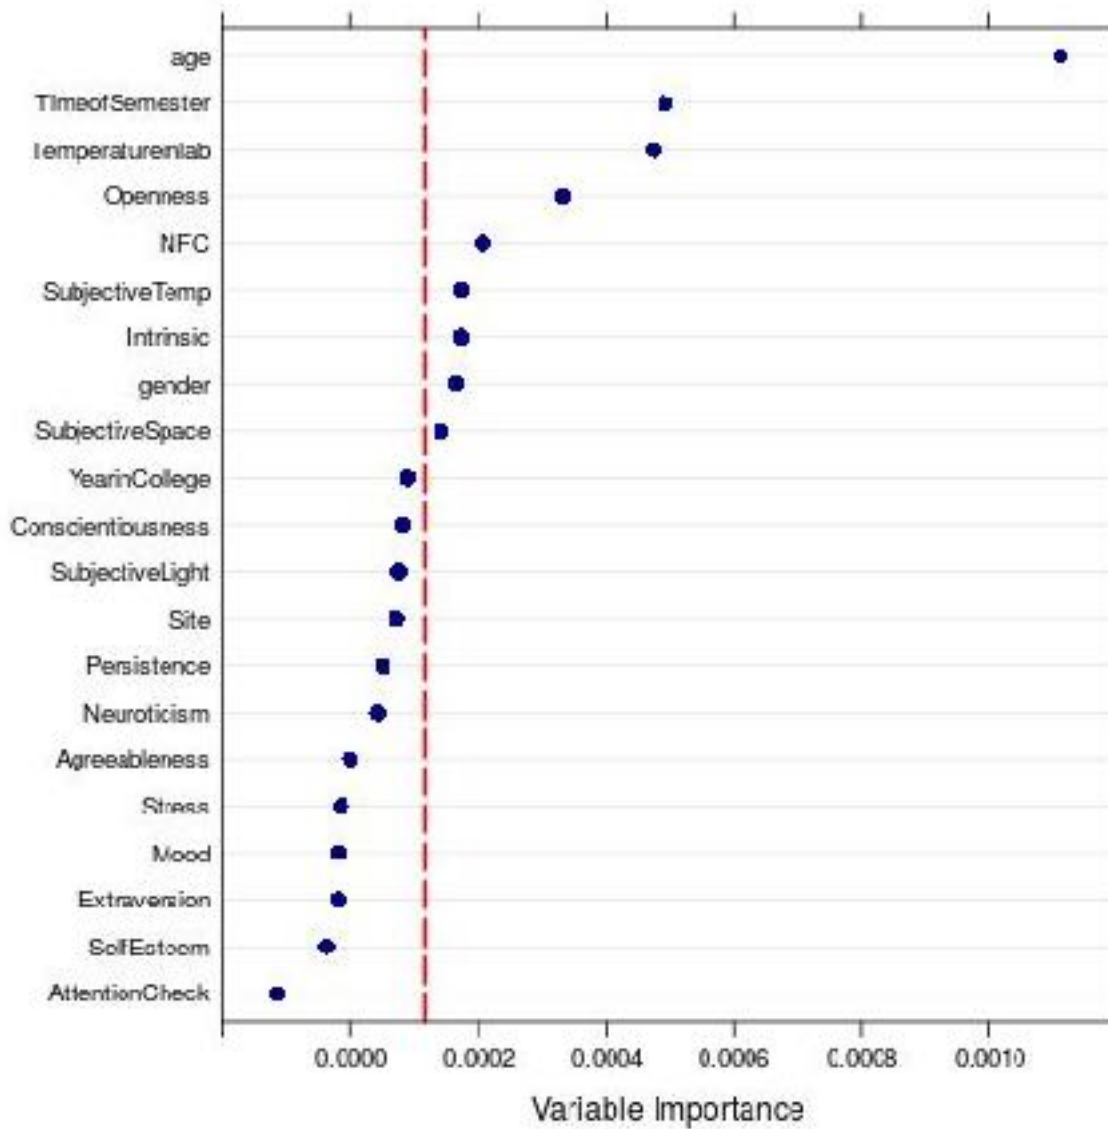
Figure 1. Permutation variable importance obtained from original analyses in the full dataset (including mTurkers; seed=1, mtry=4, trees=10k). Variation in Stroop performance from variables to the left of the red line is due to random variation; variation in Stroop performance to the right of the red line is due to real variation. Figures for other seeds are reported on our project page on the Open Science Framework.
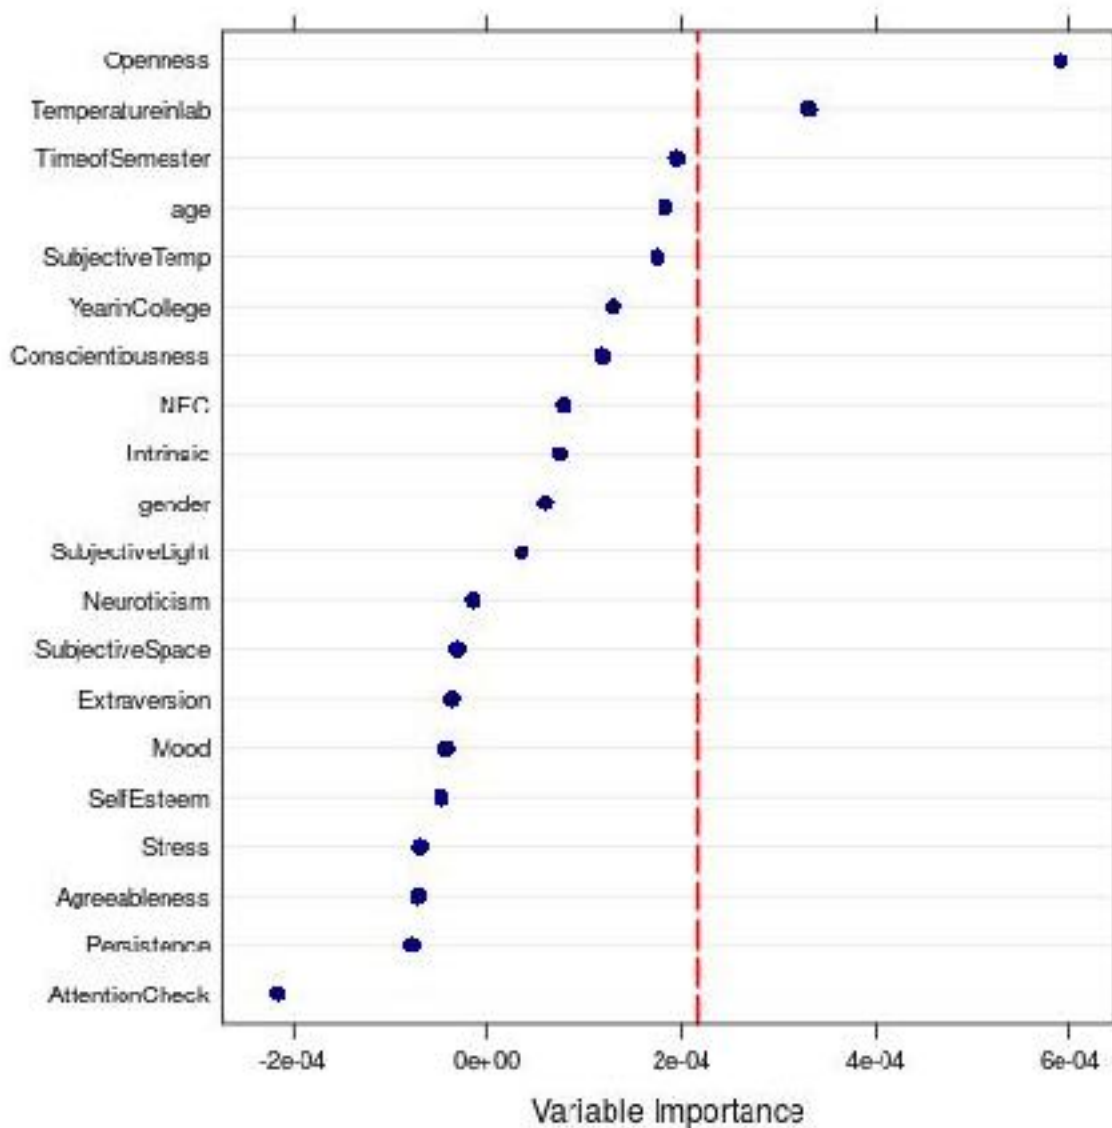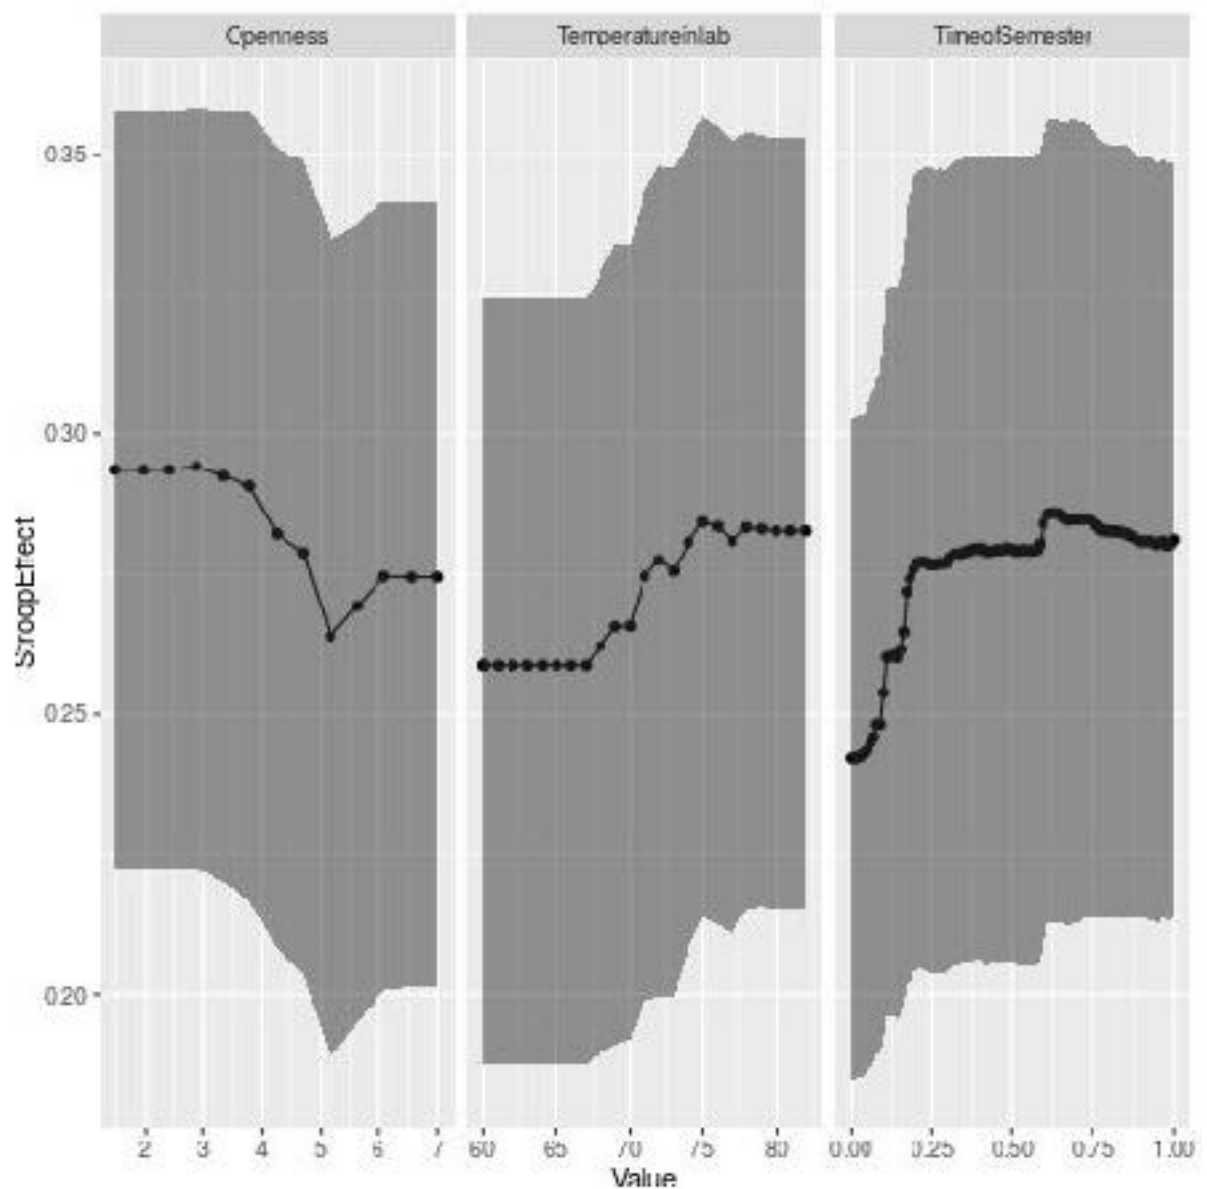
Figure 2. *Permutation variable importance obtained from original analyses in dataset excluding mTurkers (seed=1, mtry=4, trees=10k). Variation in Stroop performance from variables to the left of the red line is due to random variation; variation in Stroop performance to the right of the red line is due to real variation. Figures for other seeds are reported on our project page on the Open Science Framework.*

Figure 3. Partial dependency plots for the dataset excluding mTurkers with Openness to Experience, Lab Temperature, and Time of Semester as predictors, with Stroop performance as outcome variable (we included Time of Semester as variable that did not exceed "the red line" because of its importance in the original paper). Variables are plotted while keeping all other variables in the random forest at a constant.

Figure 4. Partial dependency plots for the dataset including mTurkers with Age, Time of Semester, Gender, Need for Cognition, Intrinsic Motivation, Need for Cognition, Openness to Experience, Lab Temperature, Subjective Temperature, and Subjective Space as predictors, with Stroop performance as outcome variable. Variables are plotted while keeping all other variables in the random forest at a constant.
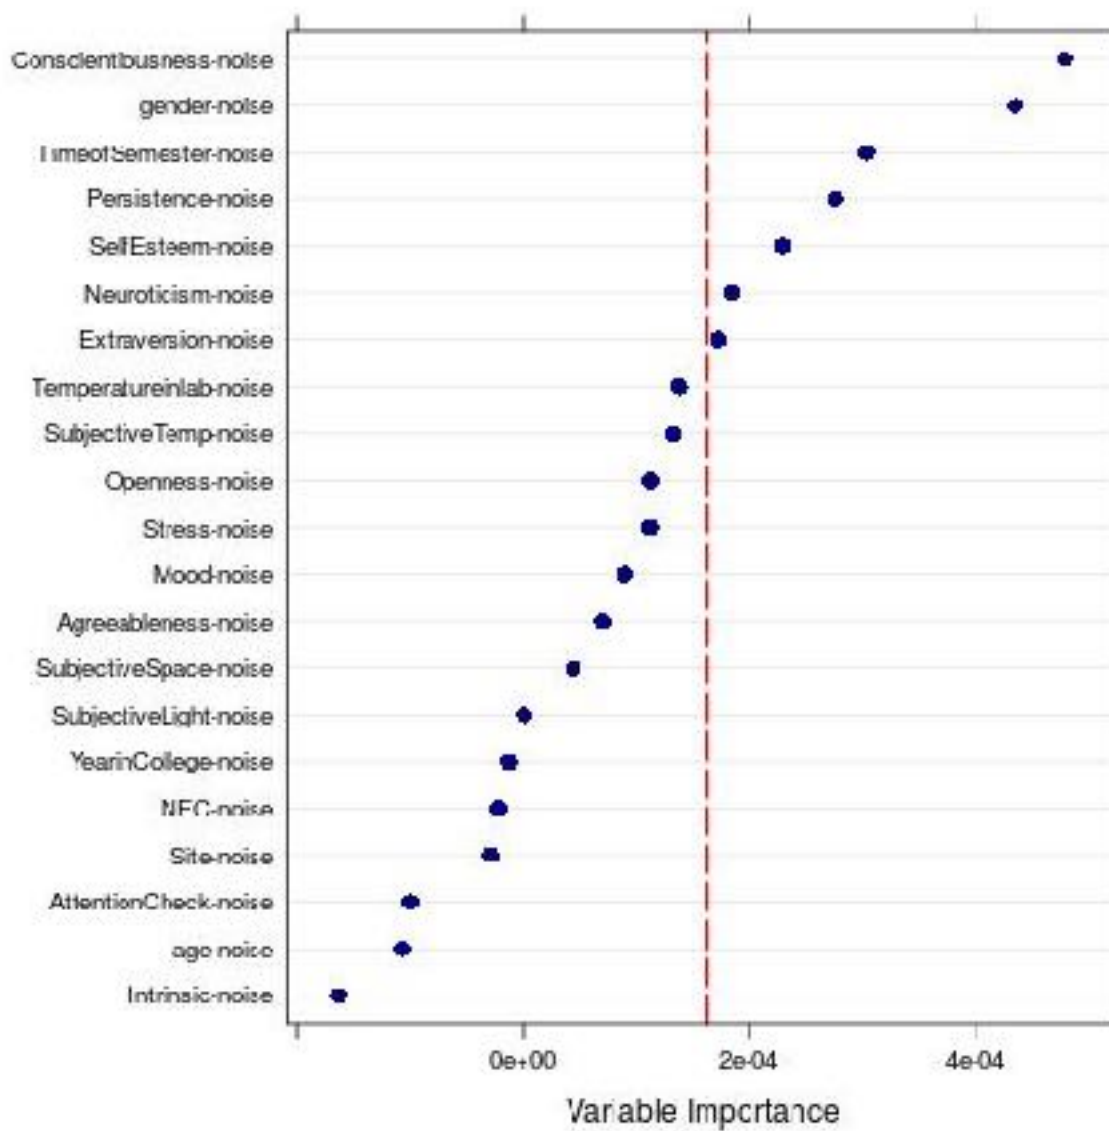
Figure 5. Permutation variable importance example obtained from our random noise dataset (seed=1, mtry=4, trees=1k). The data from the signal (Stroop Effect) is not real data, but replaced by random noise.
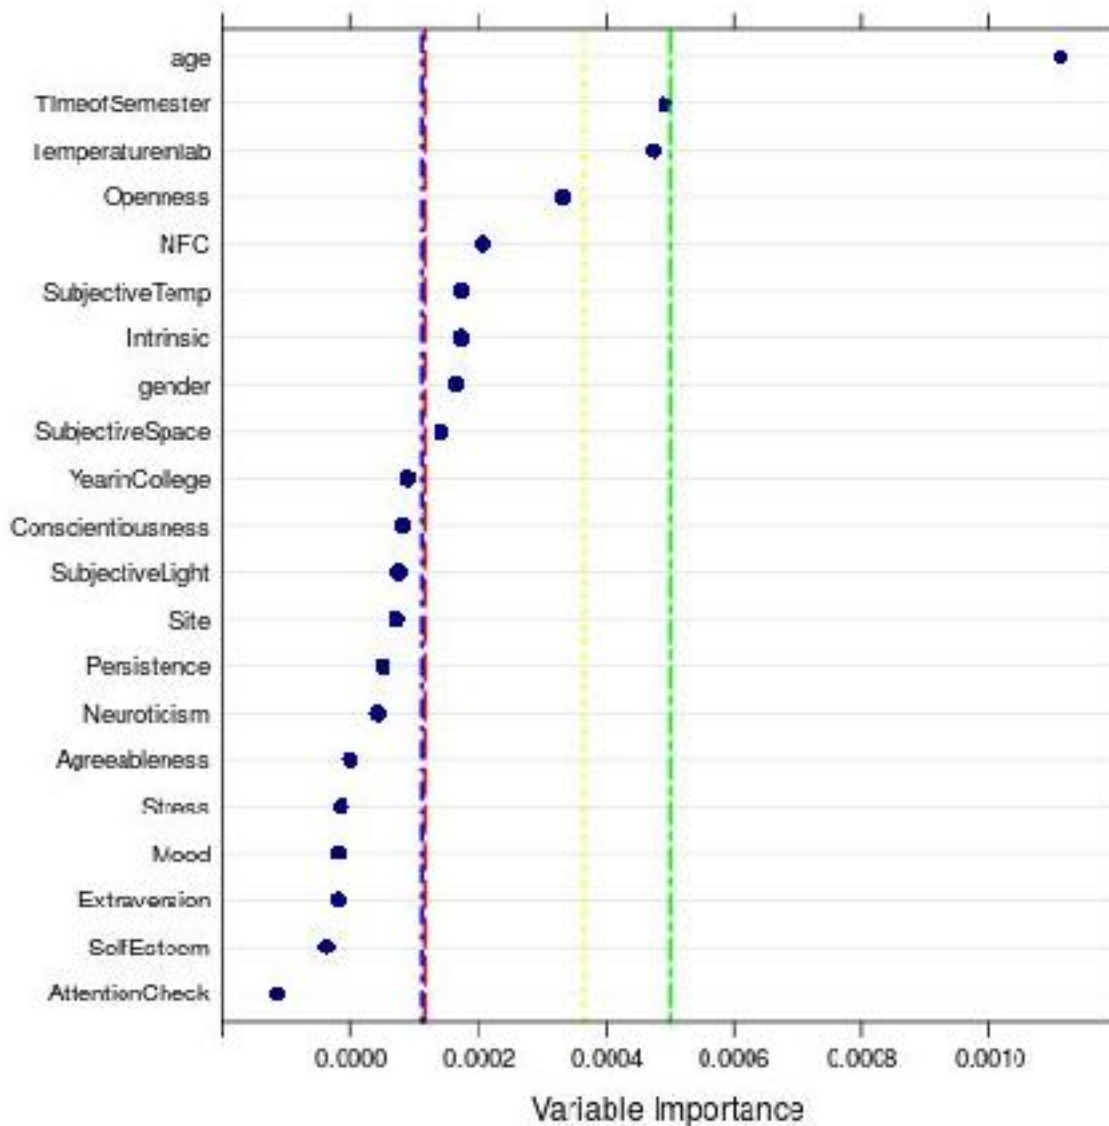
Figure 6. Permutation variable importance obtained from original analyses in the full dataset (including mTurkers; seed=4, mtry=4, trees=10k). The red line is Strobl et al.'s (2008) original threshold. The blue line is the mean of the mean of the permutation variable importance of our four random noise random forests. The yellow line is the mean of the max of the permutation variable importance of our four random noise random forests. The green line is the max of the max of the permutation variable importance of our four random noise random forests. Figures for other seeds are reported on our project page on the Open Science Framework.
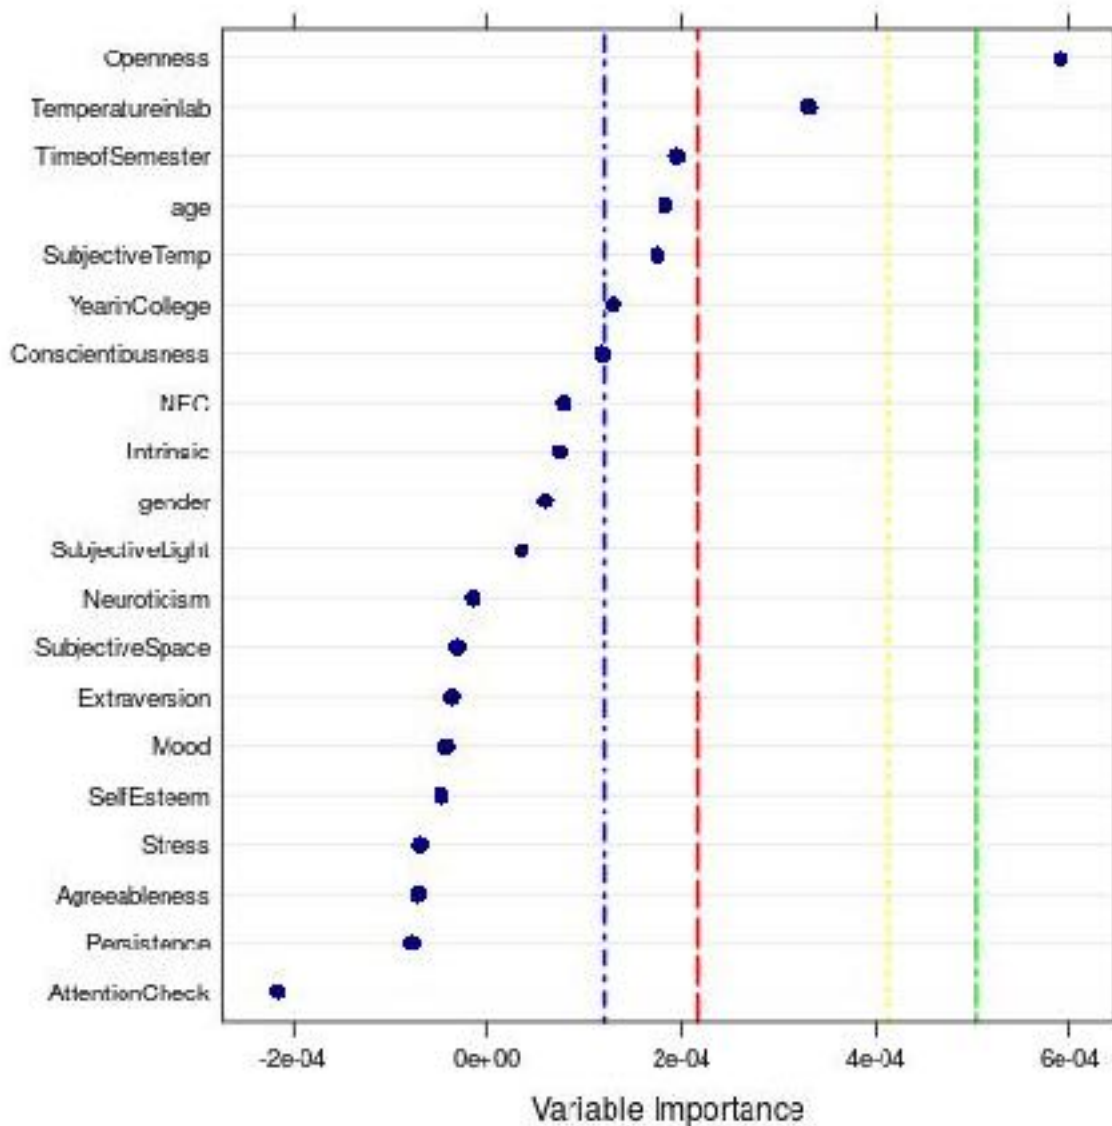
Figure 7. Permutation variable importance obtained from original analyses excluding mTurkers(seed=4, mtry=4, trees=10k). The red line is Strobl et al.'s (2008) original threshold. The blue line is the mean of the mean of the permutation variable importance of our four random noise random forests. The yellow line is the mean of the max of the permutation variable importance of our four random noise random forests. The green line is the max of the max of the permutation variable importance of our four random noise random forests (but in this case is "out of bounds". Figures for other seeds are reported on our project page on the Open Science Framework.
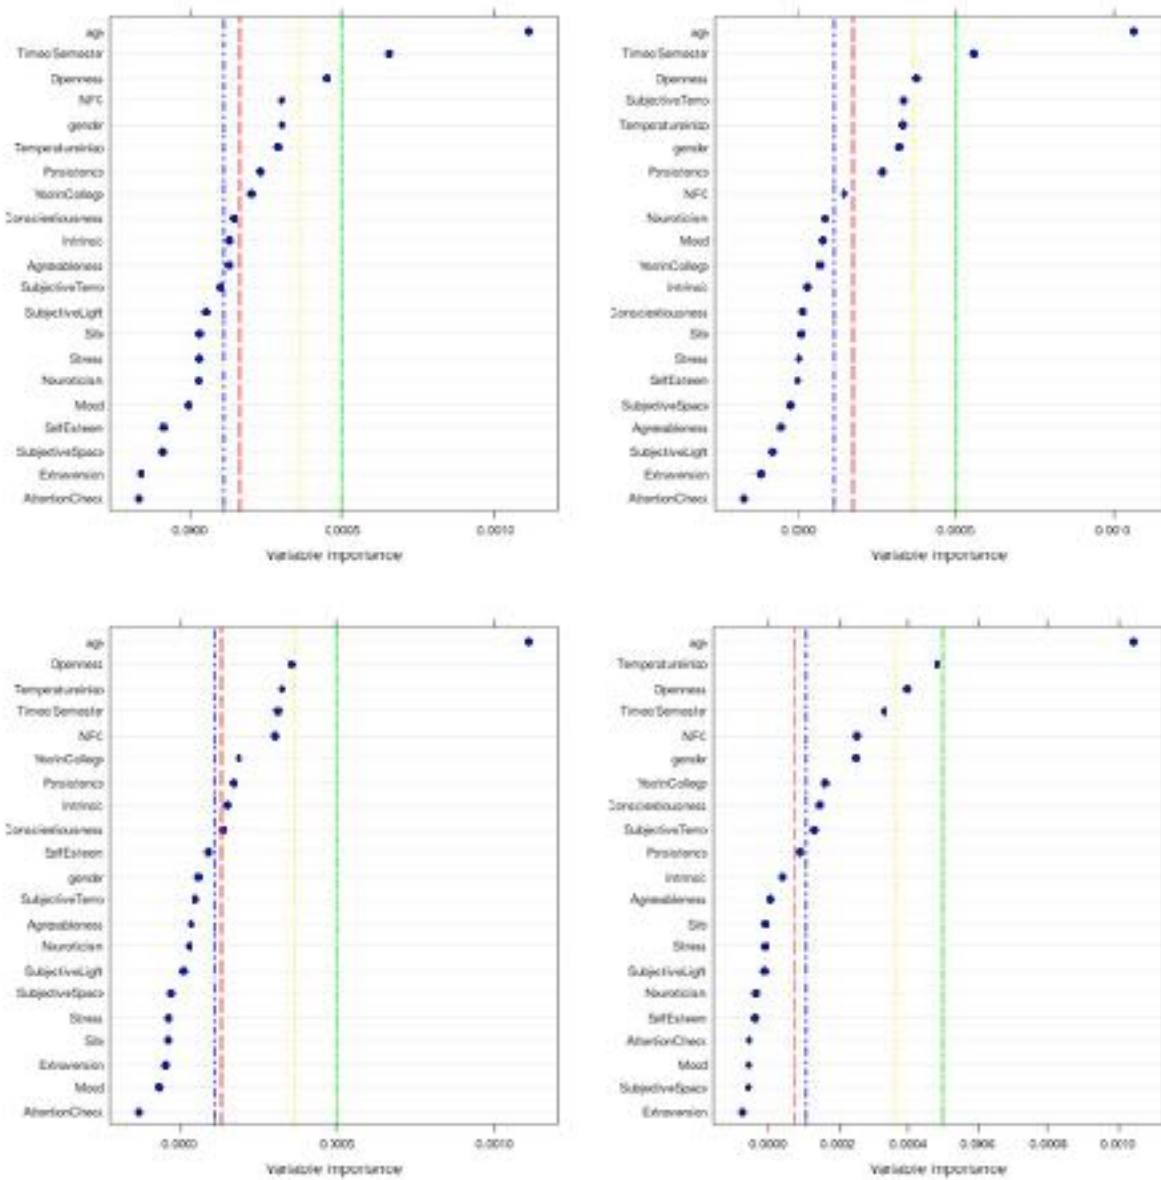
Figure 8. A side-by-side comparison of four different seeds (with mTurkers, all 1k trees); on the basis of visual inspection of these four graphs we suspect that the random noise blue line will outperform Strobl et al.'s (2008) threshold (red line) in terms of consistency.
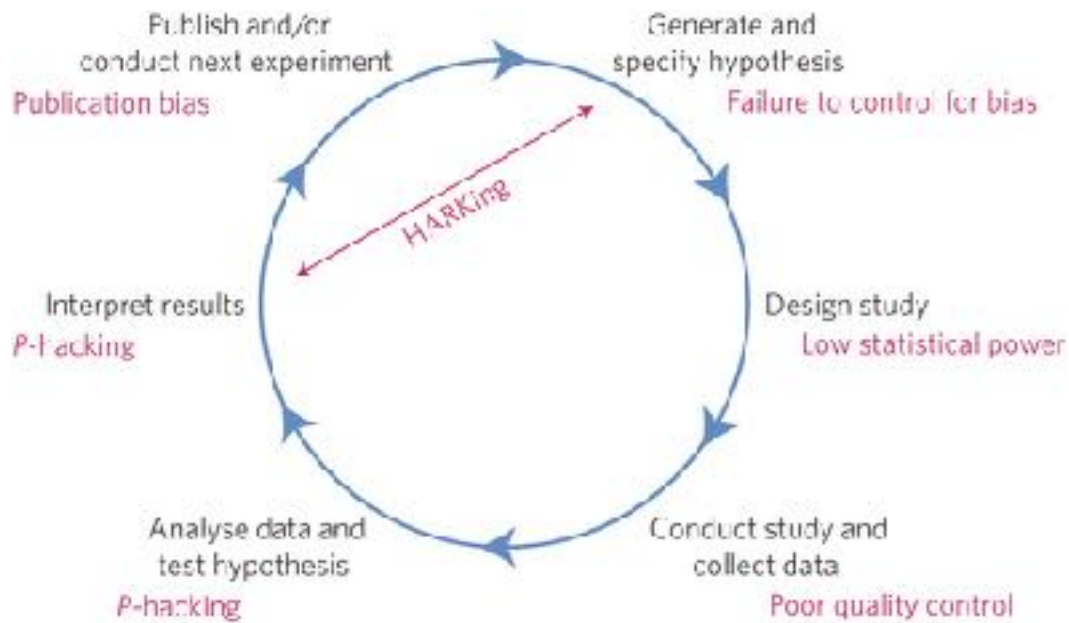
Figure 9. "Threats to reproducible science" (Taken from Munafò et al., 2017). Munafò et al. (2017) provide an idealized version of the empirical cycle (in black) with various threats (in red). P-hacking and small samples play a role throughout the empirical cycle. Supervised machine learning primarily play a role in the generation and specification of hypotheses. They can help improve the ability to generate and specify hypotheses in a data-driven manner, thereby increasing the ability to control for bias (and reducing the chances that results become p-hacked).