Title: Creating Carelessness: A comparative analysis of common techniques for the simulation of careless responder data

Authors:

Paul G. Curran
Assistant Professor of Psychology
Michigan State University
East Lansing, MI
paulcurranatwork@gmail.com

Alexander J. Denison
Graduate Student of Psychology
University of South Florida
Tampa, FL
adenison@mail.usf.edu

NOTE: This is a working draft of this paper under review. Before citing this paper please check with the authors to see if a newer version exists or if the publication status of this paper has changed.

Abstract

It is an accepted fact in survey research that not all participants will respond to items with the thoughtful introspection required to produce a valid response. When participants respond without sufficient effort their responses are considered to be careless, and these responses represent error. Many methods exist for the detection of these individuals (Huang, Curran, Keeney, Poposki, & Deshon, 2012; Johnson, 2005; Meade & Craig, 2012), and several techniques exist for testing their effectiveness. These techniques often involve generating careless responses through some process, then attempting to detect those known cases in otherwise normal data. One method to produce these data is through the simulation of data with varying degrees of randomness. Despite the common use of this technique, we know little about how it actually maps onto real world data. The purpose of this paper is to compare simulated data with real world data on commonly used careless response metrics. Results suggest that care should be applied when simulating data, and that decisions researchers make when generating this data can have large effects on the apparent effectiveness of these metrics. Despite these potential limitations, it appears that with proper use and continued research simulation techniques can still be quite valuable.

Introduction

Collection of self-report data in psychological research carries the assumption that individuals respond to items in a thoughtful way. Unfortunately, we know that this is not always the case, and the study of those individuals who respond carelessly or without sufficient effort has become a growing topic of research (Huang, Curran, Keeney, Poposki, & Deshon, 2012; Meade, & Craig, 2012). This type of behavior is similar but distinct from behaviors such as cheating or faking, as those behaviors are content responsive - individuals are actively changing their responses based on the content of the items. Unlike these behaviors, carelessness is content non-responsive, as participants respond without regard for the content of the items.

A number of methods exist to potentially detect careless responders, (Curran, 2016; DeSimone, Harms, & Desimone, 2015; Johnson, 2005; Meade & Craig, 2012), however an important challenge that this field of work faces is the circular nature of validating these detection methods. To show that a method can detect careless responders, researchers must know which participants are responding carelessly. However, if we had a way to know which participants were responding carelessly (aside from the methods being validated), we wouldn't need these methods in the first place!

Several methodologies have been used to address this problem, from instructing participants to respond as if they were being careless (e.g., Huang et al., 2012) to simulating careless data through a number of different methods (DeSimone, DeSimone, Harms, & Wood, 2018; Meade, & Craig, 2012). The focus of this paper is on the second technique, the use of simulated data to validate and compare methods of careless response detection. While a useful technique, simulation of careless data suffers from the same circular problem as above: if we

don't have a good understanding of what a careless responder looks like, how can we simulate data that looks like a careless responder?

The goal of this paper is threefold: 1) to examine current methods of simulation used for the purposes of validating metrics of careless response, 2) to evaluate boundary conditions of these techniques and identify relevant researcher degrees of freedom (e.g., Fried, & Flake, 2018; Wicherts, Veldkamp, Augusteijn, Bakker, van Aert, & van Assen, 2016), and 3) to provide guidance for the future use of simulations for validation and comparison purposes.

What does careless responding look like?

This question is core to the argument of this paper, and is a question that has been asked in different ways in many other papers on careless responding (e.g., Johnson, 2005; Huang, Curran, Keeney, Poposki, & Deshon, 2012; Meade & Craig, 2012). Carelessness does not always manifest as the same pattern of response. One individual may respond to each question at random, while the next may use the exact same response option for every question. These two responses are fundamentally distinct, and at the extremes are only infinitesimally overlapping. For example, even on a 20-item inventory with only five response options, the odds of drawing 20 consecutive numbers from a random uniform distribution has a probability of around 0.000000000000005.¹ A participant selecting the same response over and over again will produce this same pattern with little challenge at all, and a simulation technique instructed to pick the same response over and over will produce this pattern every time.

This duality in how people carelessly respond is reflected in part by the distinctions between a number of the methods of careless response detection. For instance, long-string

¹ The first of the 20 items can be any of the options; each following question has a $\frac{1}{5}$ chance of being the same value. Thus, the odds of all 20 items randomly displaying the same response is equal to $1 * .20^{19}$

analysis (i.e., counting the number of identical responses in a row) is exceptionally well suited for picking up the above pattern of 20 identical responses, but more or less blind to fully random data. Depending on the characteristics of the scale, a technique like odd-even correlations (used to assess inconsistency in responses over items) is well suited to pick up on random data, but more or less blind to the 20 identical responses. Data simulated in these two different ways (drawing from a uniform random distribution vs. a set of identical responses) would thus offer conflicting evidence on the utility of these two techniques in the detection of careless responses.

It is important to note that not all of the methods of careless response detection lend themselves to data simulation. Response time, for instance, is not dependent on the pattern of responses, and consensus appears to be that little would be learned about the efficacy of detection methods from simulating response times in a vacuum. Techniques like bogus items/infrequency items (Huang, Bowling, Liu, & Li, 2014), instructed response items (Curran, 2016), and instructional manipulation checks (Oppenheimer, Meyvis, & Davidenko, 2009) tend not to be evaluated with simulation either, as the behaviors that drive them are potentially less understood than basic careless responding.

Before describing methods of simulating response patterns, then, it is worth describing what these simulations are being used to evaluate; that is, the ways we detect carelessness that lend themselves to simulation studies.

Methods for the detection of careless responders

Many methods exist for the detection of careless response patterns (Curran, 2016; Desimone, Harms, & Desimone, 2015; Johnson, 2005; Meade & Craig, 2012). Score variance, or invariance, is the main driver by which individuals are identified as careless responders (Huang, Curran, Keeney, Poposki, & Deshon, 2012; Meade & Craig, 2012). Detection is most commonly

accomplished through the use of established cut scores that divide a sample into normal and careless components (Johnson, 2005). However, as described, detection methods for careless responding can look quite different, mainly due to the fact that careless responders do not always look the same.

Meade and Craig (2012) proposed two such conceptualizations, where careless responders either respond in an overly inconsistent way (i.e., randomly), or an overly consistent way (i.e., responding “5” to all items, or employing some patterned set of responses). Thus, there are a number of different techniques of detection that are often considered and described, and each potentially covers a different element of the manifestation of carelessness.

Long-string analysis (LS). Long-string analysis (Johnson, 2005) is a technique by which scores reflect the longest consecutive string of the same response option for each participant. LS scores range from 1 (which reflects no consecutive identical responses) to k , where k is the number of all items on the scale (a score of k reflects a participant choosing the same response for every item). This only reflects continuous strings of identical response, so a change of one response at the midpoint of a scale potentially halves this score. This score can also be divided by k to make a proportion - this is useful for comparison across scales, but results in a linear transformation over participants when k is constant. This is the only detection technique examined in this paper that explicitly detects overly consistent responses.

Within odd-even correlation (WOE) and resampled internal reliability (RIR). The within odd-even correlation is a technique by which two scale scores are generated for each unidimensional scale in a survey, one based on the odd items in that scale and one based on the even items in that scale (Jackson, 1977; Johnson, 2005). These two sets of scores (one set for odd

items, one set for even items) are then correlated within each individual to produce an odd-even correlation. For a more detailed description of this calculation, see Curran (2016).

Resampled internal reliability is a variant of the odd-even correlation that relies on randomly sampled halves of each scale rather than simply using the odd and even halves (Curran, 2016). This correlation can be generated using the same method described above, and resampled using bootstrapping to calculate and then average as many samples as desired. For our purposes we will bootstrap this score with 1000 iterations.

Psychometric synonyms and antonyms (PS and PA) coefficient. Psychometric synonyms and psychometric antonyms (Johnson, 2005) are pairs of items identified as highly positively or negatively correlated through empirical analysis of the dataset. Their calculation is identical, with synonyms relying on positive relationships and antonyms relying on negative relationships. Once these pairs are identified, they are then examined within each individual by calculating a correlation between the array of the first item in each pair and the array of the second item in each pair. High positive correlations for synonyms and high negative correlations for antonyms reflect thoughtful response, lower or oppositely signed scores on each metric reflect carelessness.

Mahalanobis distance (MD). Mahalanobis distance is a multivariate outlier technique that identifies individuals whose responses are inconsistent in scores compared to the remainder of the sample (Meade & Craig, 2012). This functions in much the same way normal outlier analysis does, by detecting responses in the tail of a distribution. However, in the case of personality measures simply detecting responses that fall within the tails of a distribution is not sufficient, as non-careless responders can easily be in the tails of a distribution if they possess very high or low levels of a given trait. Mahalanobis Distance circumvent this issue by creating a multivariate space that maps the distribution of responses to all items on a survey. Individuals can then be

compared to this distribution space to test if their set of responses falls outside the normal range of responses for that dataset. For a more in depth discussion of this method, see Curran, (2016) or Meade & Craig (2012). For our purposes in this paper we will use Mahalanobis distance as calculated in Meade & Craig (2012), as other studies appear to use this particular reference as a standard. In keeping with their calculation, Mahalanobis distance will be calculated for each subscale, and then averaged across each subscale for each individual.

Person total correlation (PTC). Person total correlation is a technique based on the more well-known method of item-total correlation (Curran, 2016; Donlon & Fischer, 1968). Just as scores on items are correlated with the sum scores of all other items in an item-total correlation, scores from individuals are correlated with the sum scores of all other participants on those items in a person-total correlation. Higher scores on this correlation reflect response patterns that are psychometrically similar to other individuals over items.

Polytomous Guttman errors (PG). Guttman errors in dichotomously scored data are defined as situations where, when items are arranged in increasing difficulty, an individual scores correctly on the harder item of a consecutive pair after failing the easier item of that pair (Guttman, 1944, 1950). The extension into polytomous data follows this same basic concept, in that an individual who agrees more strongly to an item with a lower probability of agreement after agreeing less strongly to an item with higher probability of agreement.

This score can be generated as part of the process of fitting a graded response model to data, but this tends to require assumptions that are not commonly met in places where carelessness is a concern (e.g., low-stakes personality data). For simplicity, and to mirror what might be more common in careless responding research, we will calculate these errors consistent

with Curran (2016), in that ‘item difficulty’ will be established by mean scores on that item, rather than difficulty of that response option established by IRT analysis.

As identified, only one of these techniques explicitly evaluates overly consistent responses (long-string analysis)², while many of these techniques ostensibly evaluate inconsistency in response (e.g., odd-even correlation). That is not to say that these are mutually exclusive, as a set of responses that are all identical may score poorly on person-total correlation, for instance, depending on the characteristics of the items and the remainder of the sample. Correlations between these detection metrics are rarely strong, and may be weak or even negative when examining raw values (Maniaci & Rogge, 2014; Meade & Craig, 2012). This is not to say that one of these methods is ‘right’ and the other ‘wrong’ but rather that one may be identifying one behavioral manifestation of carelessness that another does not.

Because there are many potential techniques for the detection of carelessness, one of the main objectives of several papers has been to validate these techniques by examining how well each technique can separate careless data from otherwise good data. This validation is frequently done through the insertion of ‘bad’ data into an otherwise normal dataset, thereby allowing the calculation of the sensitivity (description) and specificity (description) of these methods. In some studies, this invalid data has been produced through instructing some proportion of participants to respond carelessly (e.g., Huang et al., 2012). Several studies in recent years (Curran & Kotrba, 2012; Desimone, Harms, & Desimone, 2015; Huang, Liu, & Bowling, 2014; Maniaci & Rogge, 2014; Meade & Craig, 2012) have also adopted a technique in which this invalid data is not collected from participants, but simulated.

² Other techniques may also detect these overly consistent patterns, but have limitations when compared to long-string analysis.

While simulating data is never simple, some data simulations are easier than others. With known item characteristics, for example, it is relatively easy to simulate likely responses for test data, or even personality scales. This is largely because scores on latent traits can be drawn from some distribution, then used to generate scores contingent on the standing for that latent trait. Important to this method is the link between item content and response; in simulating carelessness it is presumed that this important linkage is broken. Any standing on the latent trait no longer has any predictive power of response. Participants are not trying to pass a test or land a job, but rather to expend as little time or effort as necessary. As such, simulations rely on more assumptions from the researcher about what that response pattern should look like.

Simulating careless data

As not all patterns of careless responding are the same, and various methods are sensitive to different types of careless responding, it may not be surprising that methods for simulating this data also differ. Invalid responses may be simulated as fully random (e.g., Huang, et al., 2014), or simulated with some item to item dependence (e.g., Curran & Kotrba, 2012; Meade & Craig, 2012). These item to item dependencies may be fairly simple or fairly complex, but reflect the idea that even careless participants are likely to have some consistency in their response patterns.

Item-independent simulations. These methods of data simulation draw values from some random distribution (e.g., random, normal), but do so independently for each item, even a within person. This is the simplest way to simulate data for these purposes, and reflects a situation where individuals are responding more or less randomly on each item. As outlined earlier in this paper, this is unlikely to produce some types of data that are assumed to be careless, such as long strings of the same response.

Item-dependent simulations. Instead of each response within individuals being independent, it is also possible to simulate data that takes prior responses into account. An extreme case of this would be a participant who responds the same way to each item on a scale - each consecutive response is completely dependent on the prior response. There are a number of ways to simulate data in this way, including building a probabilistic model (e.g., Curran & Kotrba, 2012), or simply inserting cases where all responses are the same (Meade & Craig, 2012).

Researcher choices in simulating data

While simulating data may be relatively easy, the challenges in simulation often revolve around researcher choices and assumptions made in the process of simulation. To our knowledge, there has been no work in the current literature to determine how accurate common assumptions about careless response patterns are, or what effect changing some of these assumptions has on the final simulated data. Furthermore, there exists no standard for how these simulations are conducted, which results in discrepancies between the methods being used to simulate data. This introduces ample opportunity for what some have called researcher degrees of freedom to muddy the validity of the conclusions of these studies (Wicherts, Veldkamp, Augusteijn, Bakker, van Aert, & van Assen, 2016).

Often, these are seemingly arbitrary choices made during the design, collection, or analysis of experimental data that impact the results of the study in some way. This problem is of special relevance for simulated data as the researcher tailors their data to specific parameters they set. As such, there exist a number of difficulties with relying on simulation for the validation of careless response detection techniques.

One difficulty comes back to the question asked earlier - what do careless responses look like? Interestingly, any data collection at risk for careless responders is potentially able to offer an answer this question. Given what we know about the frequency of careless responders, it is highly likely that any normal low-stakes, self-report data collection should have some careless responders in it. By applying careless response detection metrics to these normal datasets, we can gain some measure of the properties of careless responses. If these metrics truly are capturing careless responders, scores from these careless responders should appear in the tails of these detection metrics. Given this, any simulated data that is purported to mirror patterns of careless responding should have a high probability of being within the tails of these metrics.

If a given simulation method does not fall within the tail of a detection metric, researchers may conclude that this detection metric is an ineffective method to detect careless responders. Take the example of long-string analysis used in this paper several times so far. If a researcher inserts a proportion of cases whose responses are all drawn from a uniform random distribution and attempts to use long-string analysis to detect these responses, their conclusion should be that long-string analysis has no utility as a tool for the detection of careless responses as uniform random responses fall nowhere close to the top-end of this distribution.

No one has done this (to our knowledge), perhaps because the result is so mathematically tautological as to be thoroughly uninteresting. This extreme case, however, highlights how assumptions made in the simulation of data for the purposes of validating or comparing metrics of careless response are potentially related to the apparent efficacy of those metrics in the detection of that simulated data. Simulating a particular type of data that uses a particular set of assumptions may thus under- or over-power a method of detection in a way that is not indicative of real data.

Put another way, data that is simulated to represent careless responses should itself have a reasonable chance of being within the range of responses of any otherwise normal dataset. This does not necessarily mean that the exact set of responses should be present, but rather that simulated responses should have similar properties to at least a portion of the real data. For example, if a researcher simulates invalid data to test of the effectiveness of psychometric antonyms that have values well outside the range of values ever found for psychometric antonyms in a normal study, then that test will be drastically overpowered. In the same way, if values on a metric are essentially overlapping the normal distribution of scores, then a test of that metric may be underpowered.

Despite the increased use of simulation as a tool for testing the validity of careless response detection, there has been little to no work on examining and evaluating the properties of these simulated datasets. This is somewhat troubling, as these simulation studies are often used, at least partially, to provide validation for methods of detecting careless responses, or comparisons between methods to show that one technique is better than another.

The purpose of this paper is to explore the distributions of scores on common careless responding metrics in both normal data and data generated through commonly used simulation techniques. An important question to answer about these techniques is how closely they resemble real response patterns, and how the assumptions underlying these simulations change this overlap. If there is little resemblance between real and simulated data, then there is no reason to assume that detection methods that perform well on simulated datasets will have similar performance on real data.

Method

Participants

Two sets of human subjects data were collected for this study. These data were collected from two sources:

- 1) One sample of 148 individuals were collected from the psychology research pool of a small liberal arts college.
- 2) One sample of 149 individuals were collected from Amazon's Mechanical Turk (MTurk).

All data from human subjects were collected online using the Qualtrics survey platform. Participants in the college sample received course credit for their participation, while participants in the MTurk sample received \$0.75 for their participation.

Before simulating data, the two sets of human-subjects data were compared to check for any substantial differences in carelessness metrics that would prevent these cases from being considered as a group rather than as two different samples. While there were occasional differences in the location of these distributions, the majority of this appears to be in cases of rare outliers (e.g., an individual who responded the same way to 40 of the 50 items). This lack of consistent differences and existence of these outliers makes a case for treating these two samples as a larger 297 subject sample, and simulating data to match this number.

A number of datasets were simulated to roughly match this sample. All simulated data was created by simulating 50 responses for each 'person' to mirror the structure of the real data. This was implemented for 297 cases for each simulation methodology. While all data are drawn from some random distribution, choices in how that data is drawn produce quite different outcomes. These simulation methods are described below.

Uniform. For this simulation, simulated participants' responses were drawn from a random uniform distribution of integers from the scale minimum (1) to the scale maximum (7). Each response was selected independently of every other response, and the data generated by this process should represent purely random, fully content non-responsive responding.

Normal. Building on the idea that responses for any given scale may have scores closer to the center of the scale rather than the endpoints, data simulated for this condition was drawn from a random normal distribution with a mean of 4.0 (the midpoint of the scale) and a standard deviation of 1.0. Responses were again fully independent across items.

Normal Inverted. While a case could be made that individuals cluster toward the midpoint of a scale, an alternative case could be made that individuals cluster at the endpoints. This method draws simulated participants' responses from an inverted random normal distribution, a bimodal distribution with highest frequencies at the top and bottom of the scale.

Empirical. Simulating responses using any of the techniques above requires making assumptions about what normal responses on a scale might look like. Because we have human subjects data, we can use that data to alter these assumptions to fit a real-world sample. For this simulation, responses were drawn from a response distribution with probabilities matching that of the human subjects data, as described above. This created a simulated dataset with roughly the same mean and standard deviation as the human subjects data, but without any of the interdependence and covariance present in that data. Simulated participants' responses were still drawn from a random distribution, simply one that matched the characteristics of the human subjects sample (for reference, these probabilities, from responses 1 to 7, are: .02, .07, .11, .09, .23, .34, .14).

Invariant. As mentioned earlier, drawing a string of the same number over and over from a random normal or random uniform distribution is exceptionally rare. However, we know that carelessness can sometimes manifest as long strings of the same response (this is why long-string analysis exists). For this simulation, simulated participants' responses to the first item of the scale was drawn from a random uniform distribution, and that response was replicated for each of the 49 following questions. This produced completely invariant data from a within-subjects standpoint, though variance still exists between subjects.

Patterned 50-50, Patterned 50-75, and Patterned 75-75. To attempt to build dependence across items that does not represent fully invariant responses, a probabilistic model was built to generate responses for each simulated individual over the course of 50 items. First, a simulated individuals' first response was generated by drawing one number from a random uniform distribution. For the next response, each simulated individual had a 25% chance to move down a scale point (i.e., from 3 to 2), a 25% chance to move up a scale point (i.e., from 3 to 4), and a 50% chance to remain on the same response.

This works perfectly well as long as these responses don't move to the endpoints of the scale. That is, a simulated response of 7 can't 'move up' to 8 on the next item if responses are being generated on a 7-point scale. Two ways to deal with endpoints were used in order to illustrate how a minor decision like this can impact the final dataset. In the first, this 25% was grouped with the 25% chance to move down, creating a 50% chance to move away from an endpoint (up from 1, down from 7) and a 50% chance to stay on that endpoint (Patterned 50-50). In the second, this 25% was grouped with the 50% chance to stay, creating a 75% to stay on that endpoint and a 25% chance to move away (Patterned 50-75).

Another way that a researcher could make decisions regarding this type of patterning would be to change the first value of this set. That is, instead of a 50% chance to stay on any given value, a simulation could use a 75% chance to stay on any given value (including the endpoints). This simulation (Patterned 75-75) provides a contrast to the 50-50 patterned data that should have more invariance from this simple change in odds.

Within-Person Normal (WP Normal). To further attempt to build some habitual responding at the individual level, scores were also simulated using a normal distribution within a simulated individual. To start, each simulated individual had a mean value drawn from a random uniform distribution. This mean was then used for that individual with a standard deviation of 1.0, as above. Fifty responses were then drawn from a random normal distribution for each individual. This mimics a situation where an individual may choose a part of the scale that they hover around for the majority of their responses.

This, as with the patterned responses, creates a decision point regarding how to handle out-of-bounds responses. For the random normal responses at the full sample level this wasn't an issue, as the mean was the midpoint of the scale. However, an individual with a mean of 6 or 7 will regularly generate scores of 8 or 9. Three methods were used to handle these out-of-bounds values.

In the first method, scores simulated out of bounds (e.g., 8, -1) were simply replaced with their nearest in-bounds neighbor. That is, scores of 8 or 9 would be replaced by 7, and scores of 0 or -1 would be replaced with 1 (*Within-Person Normal, Cut; WP Normal (C)*). In the second method, these out-of-bounds values were mirrored back in bounds across the mean. That is, if an individual had a mean of 6, an 8 would be changed from 2 points above the mean to be 2 points below the mean (4), but if an individual had a mean of 7, that same 8 would be changed to be 1

point below the mean (6) (*Within-Person Normal, Mirror; WP Normal (M)*). Finally, in the third method, any out-of-bounds values were simply resampled until they produced in-bounds values (*Within-Person Normal, Resampled; WP Normal (R)*).

Measures

Items. The big five personality traits (Openness to Experience, Conscientiousness, Extraversion, Agreeableness, Neuroticism) were measured using the 50 positively worded items (10 items for each scale) from the 100-item IPIP representation (Goldberg, 1992) of Costa and McCrae's (1992) NEO-PI-R markers for the big five factor structure. As the point of this study is careless response metrics, not personality traits, these scales scores will not be used or reported in this study. These items are simply present to allow for the calculation of careless response metrics, and represent commonly used items in this type of research.

Results

One of the main questions of this work is how well simulated careless responder data overlaps with the distributions of human subjects' data on each of the careless response detection metrics examined. The distributions of human subjects and simulated data should neither fully overlap nor be fully distinct. When examining human subjects on these metrics, the individuals we think are responding carelessly should end up in the tail of the distribution for scores on a given metric. As such, simulated data should also generally be concentrated towards the tail end of that metric's distribution of scores, because that is where we would assume human careless responders to be located.

An initial representation of this overlap can be seen in Table 1, which displays the proportion of simulated data that falls into the last decile (90th-100th percentile) of the human subjects data for each metric. This use of a 10% cut is by no means a firmly known value, but

rather represents a best guess based on prior best estimates of the modal average (Curran, 2016). It is also meant to simply act as a standardized value across all metrics so that scores can be more easily compared. This avoids potential issues with the estimation of cut scores, as trying to estimate cut scores for these metrics creates a circular chain of analysis logic; wherein the metrics being examined are precisely those that were used to set the benchmarks for that examination. In addition, this circumvents the problem that one detection metric is favored simply because better cut scores have been established for it.

When examining Table 1, the values represented in each row of this table are a way to conceptualize how accurately the data generated from a given simulation method maps to natural careless data. In many cases represented in this table, one can see that the simulated data does in fact overlap well with the last decile of human subjects data, but there are also many where it does not. While these numbers are a way to summarize this information, they are not the entire story, as the distributions are more complex than simply what proportion falls into that last decile.

To gain a more complete picture of both the human subjects and simulated distributions, Figures 1-8 give a visual representation of the distribution of human subjects data and a type of simulated data across each metric. The red bars in these figures represent the last decile of the human subjects data, which are the human subjects we believe might be careless. This means a simulation method that is performing well would ideally fill the space between these two bars. Additionally, Supplemental Table 1 provides the proportion of simulated data that overlaps with the entire human subjects distribution.

Values for all simulation methods and careless detection metrics were computed, but for simplicity those data simulated to be invariant (e.g., all responses identical) were not included in

our figures. This is due to many metrics failing to be able to compute scores without variance, and of the clear ease with which long string analysis can identify this data. These graphs are thus mostly empty, with the exception of those techniques where what is present is uninteresting.

While the above graphs are useful in gaining a general idea of the distributions of these simulation methods, there is an important distinction between the methods of careless response detection in that some are sample-specific, and some are not. Given a response from a single individual, long-string analysis, ISD, within odd-even correlations, and RIR can be calculated on that individual. This individual response can be placed into any sample, measured on the same items, and the scores on those metrics *will not change*. These metrics produce fully sample-independent scores, and no other cases need be referenced in their computation. A single observation can be scored in complete isolation of any other observation.

Given the same singular response, and in absence of any other observations, none of the other metrics (psychometric synonyms/antonyms, polytomous Guttman errors, person-total correlations, and Mahalanobis distance) can be calculated. Each requires a sample for the calculation of at least part of the resulting score. Beyond this, the same individual response string, inserted into two different and distinct samples, will produce two different scores on these metrics. These are sample-dependent metrics, and other response strings in the sample have an impact on any resulting scores for that individual case- a single observation cannot be scored in isolation of other observations.

Mixing Simulated and Human Subjects Data

So far, the figures and graphs have examined a situation where scores were calculated on the human subjects data and simulated data in full isolation. While this has low fidelity for these sample-dependent metrics relative to practical examples, the authors of this paper believe it still

represents an important starting place for additional examination. Because these sample-dependent metrics are, simply put, sample dependent, inserting cases into a single sample is a specific case representation of that data that potentially obscures the more general case. This more general case, presented above, allows researchers to better understand what these distributions look like in the absence of sample-variable human subjects data. In addition, for sample-independent metrics, the scores should not be impacted by the sample.

However, it would be a disservice to the reader not to examine what happens when these two samples (human subjects and simulated) are mixed together. Discussed earlier, a common technique used in the evaluation of careless response detection metrics is to inject simulated data into human subjects data in order to test the resulting sensitivity (true positive rate) and specificity (true negative rate) of those metrics. Thus, each simulated dataset was introduced as a mix into the human subjects data, at two levels: 10% mixing and 20% mixing.

To build this mixed dataset while eliminating the potential spuriousness of replacing cases, the full dataset was simply increased in size by tenfold. In the case of 10% mixing this included nine identical versions of the human subjects data and one instance of a simulated dataset; in the case of 20% mixing the ratio was 8:2. As no statistical tests are being used, only proportions, this will not have any impact on power or magnitude of our results. Proportions, like those that drive sensitivity and specificity, are not impacted by simple multiplicative increase of sample size. The main potential limitation of this methodology is that each datapoint can exist in the data multiple times. While this is a potential limitation, we believe that other alternatives to this method also have limitations that outweigh this one. Curious readers can refer to Supplemental Table 2 to see similarities between sensitivity and specificity across the duplicated

dataset and one where the simulated dataset was reduced and the human subjects dataset was not modified.

In each 10% mix dataset, on each metric of detection, the top 10% (or bottom, depending on whether higher or lower values represent carelessness) of cases were flagged as careless. In the 20% mix, this was the top (or bottom) 20%. The reason for this raw proportional cut is one of standardization. Many, but not all, of these metrics have published (and sometimes competing, sometimes vague) recommendations for cut scores on those metrics. Some are sample-independent (e.g., correlations below 0.0), some are sample-dependent (e.g., 2SD above the mean). There is no dispute that some of these cut scores are more robustly researched than others, and, beyond that, that some are simply more likely to be accurate than others. Using a raw proportion for this analysis again holds this potential confound constant across metrics, and gives each an opportunity to simply catch the 10% or 20% of worst cases that we know have been inserted.

It should also be noted that, as above, we assume these human subjects data contain careless responders. It is a broad limitation of this type of research that not all papers address, but should always be considered. This will potentially decrease both sensitivity and specificity, overall. While it is unknown how much of an impact this will have, our assumption is that it is a negligible impact at this point in understanding this phenomenon. Even if this effect is non-negligible, there is no reason to believe that this error would differentially impact only some metrics and not others. A potential solution would be to attempt to clean the human subjects data, but doing so would again involve the application of the very metrics that are being examined. The same argument as earlier applies.

Table 2 presents the sensitivity and specificity for each of these metrics of careless response detection on each of our mixed datasets at both 10% and 20% mix levels. In the 10% mix, the worst 10% of cases on each metric were identified as careless, and in the 20% mix the worst 20% of cases were. Sensitivity can range from 0 to 1.0, reflecting a metric capturing some proportion of simulated data in that mix. However, because of these proportional cuts, rather than cuts driven by other techniques, specificity is constrained to a value of .88-1.0 for the 10% mix, and .75-1.0 for the 20% mix. That is, even if no simulated data is captured by a technique, only 10% (or 20%) of the total data can be *incorrectly* removed. As the human subjects data represent 9/10 and 8/10 of the total cases in both cases, the lowest possible specificity is found by $1.0 - \frac{1/10}{9/10}$ and $1.0 - \frac{2/10}{8/10}$ in these cases, or more simply 8/9 and 6/8 (.88 and .75).

Despite these theoretical floor values, a detail-oriented reader will recognize that specificity occasionally drops below these values. This is due to the practical inability to take percentile-based cuts from a dataset where individuals can share the same value on a measure. In a simple case, consider a variable with ten observations. If the top three of those observations are all the value '5', a cut to exclude the top decile of that variable either fails to select any cases or cuts all three. In this analysis, we took the stance of cutting all three.

Missingness is also a decision point here, further highlighting the range of decisions that a researcher needs to make when using these techniques. In situations of invariance, or even local invariance, many of these metrics produce no value whatsoever. While this is not explicitly discussed in any prior literature that we could find, the general practice among researchers in this field would be to ignore cases where values cannot be calculated. That is, those cases are not removed, but retained (potentially in the hope that they are better distinguished by a different metric that does produce a score). In this analysis, cases where a score could not be calculated

were considered to not be flagged as careless, and the denominator of sensitivity and specificity calculations was not changed to reflect missingness due to this point.

Discussion

The goal of this paper was threefold: 1) to examine current methods of simulation used for the purposes of validating metrics of careless response, 2) to evaluate boundary conditions of these techniques and identify relevant researcher degrees of freedom, and 3) to provide guidance for the use of simulations for validation and comparison purposes in the future.

As a part of this, we expected that different simulation methods would produce simulated ‘careless data’ that looked quite different, both across and within careless responding detection techniques. The extension of this is that choices by the researcher about how to simulate these data can have consequences regarding the apparent strength or weakness of that metric. This result is quite evident.

The results of long-string analysis on non-mixed data (Figure 1, Table 1) are perhaps some of the clearest cut, due to the relatively straightforward computation of the metric. As is mathematically obvious, no random process without interdependence between items comes close to producing values that might be flagged by this technique. The only simulation method that gets close is patterned data, where each response is linked to the last with some chance to go up or down the scale. Small choices made by the researcher can have large impact here; the difference between the three patterned datasets is noticeable. The patterned 50-50 data is more or less overlapping the human subjects data, the 50-75 starts to become more distinct, and the 75-75 is starting to look reasonably distinct. This speaks to the second point of this paper, as even this seemingly trivial change *within* one type of simulation is enough to drive changes in outcomes regarding the strength of this metric.

This second (50-75) and third (75-75) method are the only simulation methods for long-string analysis that produces more than 5% of data falling into the last decile range of human subjects data (50-75 has 21%, 75-75 has 55%). In fact, the invariant data does not even fall into this range, as no participant chose the same response to all 50 questions (the highest was 40). Using invariant simulated data here then seems to be overpowered, and a strawman for long-string to catch far too easily and effectively. The same is largely true for the mixed datasets (Table 2) and patterned data. The data simulated with lowest probability to stay on any point (50-50) looks much like any of the other simulated data, with sensitivity near zero and specificity at or below the theoretical floor value. However, the data simulated with highest probability to stay on any point (75-75) is considerably better, with a sensitivity of .43 and a specificity of .92. This is a shift from detecting almost nothing to being second best among all techniques (only person-total correlations does better here, with ST/SF of .58/.95). Thus, small researcher decisions like this can have seemingly large impacts on the conclusions one would draw about the effectiveness of these techniques.

Long string analysis suffers from a simple problem – it's not good at picking up random data. This is not a new result, nor is it mathematically surprising. That said, any researcher trying to compare the usefulness of these techniques by inserting data simulated from a random distribution would find long string to perform poorly relative to other metrics. While this is true, it obscures a different point, that long string analysis is unparalleled in picking up invariant responses where other metrics may not even be able to produce a score. While this seems obvious, long string analysis and invariant data are simply the tip of this iceberg that is clearly visible to us - not all of these results are so obvious or so easy to consider.

Both the within odd-even correlation and its resampled extension, resampled internal reliability (Table 1, Figures 2 and 3), also have a somewhat straightforward story on the non-mixed dataset. Fully item-independent simulated data of any sort (e.g., uniform, normal, inverted normal) produced distributions on these metrics that are roughly centered on zero, demonstrating a lack of covariance between these items within person. This is not unexpected, as it is precisely the argument for using these metrics. The distribution of RIR also has a smaller standard deviation than within odd-even correlations, reflecting that RIR is a distribution drawn from sampling and averaging within odd-even correlations (i.e., RIR is a sampling distribution of WOE).

Contrasting fully item-independent data, results change drastically when even a simple pattern of covariance is built between items. Each of these patterned distributions are largely indistinguishable from human subjects data, with the exception that these simulated data seem to actually appear slightly *better* than the human subjects data. This again speaks to the second point of this paper on decisions in simulating data. A researcher testing either of these metrics using the insertion of simulated data could drive large differences in the efficacy of these methods through their decisions regarding how to simulate data. Psychometric synonyms (tables 1 and 2, Figure 4) follows much the same pattern, as it is also built around these correlations.

The same holds true for the mixed data (Table 2). In this case, to a researcher employing data simulated with a uniform random distribution (or normal, or inverted normal, etc.), these metrics look quite good. In the 10% mix the specificity is relatively stable around .95, and sensitivity is in the range of .50-.60. Overall, this shows positive support for this method, if this was how data were simulated to test it. However, a researcher using data simulated with a pattern, as in our simulated data, would instead find dismal results for these methods: sensitivity

near 0.01-0.02 and specificity very near or at the lowest possible value (.89). In this case a researcher might opt against using these techniques, or cautioning against their use in the future, all depending on what data were being used to test the quality of this metric. Psychometric synonyms again have a similar pattern in the mixed data (Table 2).

Psychometric antonyms (Table 1 & 2, Supplementary Figure 1) provide a bit of a tangent, as the inclusion of their calculation is not directly related to the hypotheses of this paper, but rather to illustrate the potential boundary conditions of this work and others. These data do not contain negatively worded items, and all items on each scale are keyed in the same direction at the raw item level. As such, psychometric antonyms, which finds and utilizes the largest negative correlations in a dataset, was left to operate on what minimal cross-scale negative correlations it could find. Those correlations were lower than what would be considered reasonable to generate these scores, and as such the resulting scores from this metric look largely indistinguishable between randomly simulated data and human subjects data. In both cases these scores are largely being driven by randomness.

Guttman errors and ISD (Table 1, Figures 5 and 6) are an unexpected pairing in terms of somewhat similar results in the non-mixed data. Specifically, both are quite good at picking up on uniform random data and inverse normal data (though ISD particularly shines here). With the exception of empirically simulated data (which Guttman errors does well with), both are relatively weak with every other simulation technique. Both are notably weak regarding invariant data, as these are both metrics which are not limited on the requirement of variance to produce scores. Thus, invariant data scores a true zero for both of these metrics. Other results make mathematical sense regarding ISD, particularly in the contrast between uniform data and normal data. One practical difference between these two techniques is a reduction of standard deviation

of the resulting within-person distributions (between uniform and normal random data), exactly what ISD is measuring.

When examined in the mixed dataset (Table 2), polytomous Guttman errors show reasonable values for uniform, normal inverted, and empirically simulated data, but are weaker in most other situations. ISD shows excellent results when simulated data are drawn from an inverted normal distribution (ST/SF: .84/.99), quite good results when data are drawn from a uniform distribution (ST/SF: .44/.95), and reasonable results when data are patterned in the 50-50 case (ST/SF: .25/.93) or 50-75 case (ST/SF: .30/.93). In all other cases, sensitivity drops below .15 and specificity drops to at or near the floor value.

Person-total correlations, examined in the non-mixed data (Table 1, Figure 7), appear to work very well in all item-independent simulated data, and even perhaps marginally well in those cases where items were not fully independent (patterned data). Contrasting this, Mahalanobis distance (Table 1, Figure 8) looks, technically speaking, somewhat weird. With the exception of uniform random data, which more or less overlaps human subjects data, none of the remaining simulation methods produce distributions that map to the correct tail of the human subject distributions. This highlights the complexity of this topic which led to the creation and examination of our mixed dataset. Of the sample-dependent metrics, it is reasonable to expect person-total correlation and Mahalanobis distance to change the most from our non-mixed data to mixed data.

The results of person-total correlation on mixed data (Table 2) are those in weakest support of one of the themes of this paper, that changing simulation type can change the ability of that measure to detect simulated carelessness. With the exception of fully invariant data (which person-total correlation is mathematically blind to), the sensitivity for the 10% cut on all

other simulation types only ranges from .57 (on two of the patterned datasets) to .61 (on the uniform random dataset). The remaining seven simulation types all give the same sensitivity (.58). Again, discounting fully invariant data, the range of specificity is even smaller, ranging only from .95 (on all simulation datasets except uniform random) to .96 (on the uniform random dataset). While this potentially and unexpectedly shows strong support for the use of this technique, it still highlights the fact that even a technique like this would be blind to someone answering in a fully invariant way.

The results for Mahalanobis distance were strange in our first analysis, but in the mixed data the picture becomes a little clearer (Table 2). Mahalanobis distance is quite good when assessing data simulated from a uniform distribution (ST/SF: .55/.95 on the 10% cut). Mahalanobis distance is the only metric better than ISD (at least on the 10% cut) at detecting data simulated from an inverse normal distribution (ST/SF: .94/.99). Aside from these, however, it's merely passable on data simulated from empirical values (ST/SF: .29/.92), and relatively bad everywhere else. The next best performance is on data simulated within-person normal with endpoints cut (ST/SF: .10/.90), and on data simulated from a normal distribution, and on all patterned datasets it gives a sensitivity of zero.

If we, as researchers, wanted to put out a paper in support of Mahalanobis distance, we would simply need to make sure that the data we used to simulate carelessness was from a uniform random distribution. If a competing researcher wanted to put out a paper against the use of Mahalanobis distance, they could perhaps just as easily argue that their data should be drawn from a random normal distribution. These two papers would paint drastically different pictures of Mahalanobis distance, and the risk this paper is attempting to demonstrate is that both of these papers could be internally consistent, yet show completely opposite results.

This result, seen not just in Mahalanobis distance but on a number of other metrics as well, supports the core concern of this paper: that even subtle choices in simulation methods used to examine and evaluate careless response detection metrics can potentially confound the resulting efficacy of those metrics. Researchers using these simulation methods to examine careless responding need to consider how data are being simulated in their studies, and consider how they are potentially operationalizing carelessness through the decisions that they make.

The remaining point of this paper is to attempt to provide some guidance regarding these types of studies as well as the use of these metrics. The statement above, that researchers should consider these results when simulating data is perhaps the most important guidance that we can give. An extension of this would perhaps be to consider simulating a number of different ‘careless’ samples using different techniques in order to gauge the true strengths and weaknesses of a metrics, not just one or the other. This study is not explicitly about evaluating the usefulness of each metric, but there are still potentially valuable insights to be taken from the data in an exploratory fashion.

Metric evaluation

In general, no one simulation method seems to work in all cases. This is not entirely unexpected, given that different methods of careless response detection pick up on different characteristics of responses. For instance, within odd-even analysis is quite literally blind to invariant response patterns, while long-string analysis is very quick to find these cases. There are many pairings like this, and it speaks to the fact that there will likely never be a singular ‘silver bullet’ technique that captures all the different manifestations of carelessness in any given data. Instead, researchers tend to use a mix of different techniques, presumably, but rarely explicitly, to detect individuals with different manifestations of carelessness.

Each technique has blind spots, but if pressed for a technique with the fewest, the answer would seem to be person-total correlation. Unlike most techniques that have a split between the ability to detect random data vs patterned data, it does well on both. Its main blind spot is invariant data, for which scores are simply not calculated. Fortunately, there is another technique that perfectly fills this gap: long-string analysis. A pairing of long-string analysis, to catch those fully invariant cases, with person-total correlation, to catch everything else, seems to potentially be quite powerful. Future work may be well suited to examine this (and other) pairings of this sort.

To examine one additional point, this study, particularly the second element involving mixing data, is limited in some ways by knowing how much ‘careless’ data has been inserted. We never know this actual number, and some recommend applying overly conservative cut scores in order to prioritize specificity in every metric on the basis that each metric will catch the most egregious cases that metric is singularly good at. This question is somewhat outside of the main elements of this paper, but can at least be examined here as a way to potentially spark more work on this topic. Another table (Supplemental Table 3), in the same form as Table 2, was created to examine what would happen in the 10% mix data if only the worst 5% of data were removed. In essence, it answers the question of what happens when one prioritizes specificity over sensitivity. To those curious to this question, this supplemental table starts to provide a framework for an answer.

To reiterate the main takeaways of this paper, simulating careless responder data is not as simple as generating uniform random noise, or in fact generating any single type of data at all. There are many ways we can simulate careless responses, and the choices that we make as researchers in these simulations can have a potentially large impact on the conclusions one

draws. No simulation study of carelessness is complete without considering the different ways carelessness can be manifested, and therefore operationalized.

Working draft

References

- Costa, P. T., Jr., & McCrae, R. R. (1992). *Revised NEO Personality Inventory (Neo-PI-R) and NEO Five-Factor Inventory (NEO-FFI): Professional manual*. Odessa, FL: Psychological Assessment Resources.
- Curran, P., Kotrba, L., & Denison, D. (2010). Careless responding in surveys: Applying traditional techniques to organizational settings. Poster presented at the 25th Annual Conference of the Society for Industrial and Organizational Psychology, Atlanta, GA, April 8– 10, 2010.
- Curran, P., & Kotrba, L.M. (2012). The impacts of invalid responding: a simulation study. Presented at the 27th Annual Conference of the Society for Industrial and Organizational Psychology, San Diego, CA, April 26–28, 2012.
- Curran, P. G. (2016). Methods for the detection of carelessly invalid responses in survey data. *Journal of Experimental Social Psychology*, 66, 4–19. doi:10.1016/j.jesp.2015.07.006
- Curran, P.G. & Hauser K.A. (Revise & Resubmit). I'm paid biweekly, just not by leprechauns: Evaluating valid-but-incorrect response rates to attention check items. osf.io/y623k/
- DeSimone, J. A., DeSimone, A. J., Harms, P. D., & Wood, D. (2018). The Differential Impacts of Two Forms of Insufficient Effort Responding. *Applied Psychology*, 67(2), 309–338. doi:10.1111/apps.12117
- DeSimone, J. A., Harms, P. D., & DeSimone, A. J. (2015). Best practice recommendations for data screening. *Journal of Organizational Behavior*, 36(2), 171–181. doi:10.1002/job.1962

- Donlon, T. F., & Fischer, F. E. (1968). An Index of an Individual's Agreement with Group-Determined Item Difficulties. *Educational and Psychological Measurement*, 28(1), 105–113. doi:10.1177/001316446802800110
- Fried, E., & Flake, J. (n.d.). Measurement Matters – Association for Psychological Science. Retrieved August 22, 2018, from www.psychologicalscience.org/observer/measurement-matters
- Goldberg, L. R. (1992). The development of markers for the Big-Five factor structure. *Psychological Assessment*, 4, 26-42.
- Guttman, L. (1944). A basis for scaling qualitative data. *American Sociological Review*, 9, 139–150.
- Guttman, L. (1950). The basis for scalogram analysis. In S.A. Stouffer, L. Guttman, E.A. Suchman, P.F. Lazarsfeld, S.A. Star, & J.A. Claussen (Eds.), *Measurement and Prediction* (pp. 66–90). Princeton: Princeton University Press.
- Huang, J. L., Curran, P. G., Keeney, J., Poposki, E. M., & DeShon, R. P. (2012). Detecting and Deterring Insufficient Effort Responding to Surveys. *Journal of Business and Psychology*, 27(1), 99–114. doi:10.1007/s10869-011-9231-8
- Huang, J. L., Liu, M., & Bowling, N. A. (2015). Insufficient effort responding: Examining an insidious confound in survey data. *Journal of Applied Psychology; Washington*, 100(3), 828.
- Jackson, D.N. (1977). Jackson vocational interest survey manual. Port Huron, MI: Research Psychologists Press.

- Johnson, J. A. (2005). Ascertaining the validity of individual protocols from Web-based personality inventories. *Journal of Research in Personality*, 39(1), 103–129. doi:10.1016/j.jrp.2004.09.009
- Maniaci, M. R., & Rogge, R. D. (2014). Caring about carelessness: Participant inattention and its effects on research. *Journal of Research in Personality*, 48, 61–83. doi:10.1016/j.jrp.2013.09.008
- Meade, A. W., & Craig, S. B. (2012). Identifying careless responses in survey data. *Psychological Methods*, 17(3), 437–455. doi.org.ezproxy.gvsu.edu/10.1037/a0028085
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4), 867–872. doi:10.1016/j.jesp.2009.03.009
- Wicherts, J. M., Veldkamp, C. L. S., Augusteijn, H. E. M., Bakker, M., van Aert, R. C. M., & van Assen, M. A. L. M. (2016). Degrees of Freedom in Planning, Running, Analyzing, and Reporting Psychological Studies: A Checklist to Avoid p-Hacking. *Frontiers in Psychology*, 7. doi:10.3389/fpsyg.2016.01832

Table 1

Proportion of simulated data captured by worst decile of human subjects scores.

Simulation	PTC	WOE	RIR	PS	PA	PG	ISD	LS	MD
Uniform	0.86	0.67	0.61	0.32	0.17	0.76	0.60	0.00	0.14
Empirical	0.91	0.66	0.65	0.34	0.12	0.64	0.14	0.00	0.00
Normal	0.89	0.65	0.64	0.38	0.12	0.41	0.00	0.00	0.00
Inverted	0.88	0.69	0.65	0.32	0.11	0.60	0.79	0.00	0.00
Patterned 50-50	0.35	0.01	0.02	0.00	0.12	0.35	0.21	0.01	0.00
Patterned 50-75	0.30	0.00	0.01	0.01	0.18	0.13	0.24	0.21	0.00
Patterned 75-75	0.29	0.01	0.02	0.00	0.08	0.07	0.11	0.54	0.00
WP Normal (C)	0.85	0.63	0.63	0.34	0.09	0.24	0.00	0.05	0.00
WP Normal (M)	0.88	0.61	0.61	0.49	0.11	0.28	0.00	0.00	0.00
WP Normal (R)	0.85	0.62	0.64	0.35	0.11	0.21	0.00	0.01	0.00

Notes: Proportion of simulated data present between the 90th and 100th percentile of human subjects' data.

Table 2 – Sensitivity and specificity of each dataset and metric using 10% and 20% cuts.

		PTC		WOE		RIR		PS		PA		PG		ISD		LS		MD	
Simulation		ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF
Uniform	10%	.61	.96	.58	.95	.57	.95	.41	.94	.25	.92	.76	.91	.44	.95	.00	.87	.55	.95
	20%	.76	.94	.68	.92	.70	.92	.57	.90	.32	.84	.79	.89	.57	.91	.00	.65	.52	.88
Normal	10%	.58	.95	.55	.95	.55	.95	.45	.94	.25	.92	.35	.87	.00	.90	.00	.87	.00	.89
	20%	.74	.94	.67	.92	.69	.92	.60	.88	.39	.85	.38	.84	.00	.77	.03	.65	.00	.75
Inverted	10%	.58	.95	.58	.95	.55	.95	.34	.93	.24	.92	.44	.93	.84	.99	.00	.87	.94	.99
	20%	.73	.93	.69	.92	.69	.92	.54	.89	.42	.86	.63	.85	.87	.99	.01	.65	.89	.97
Empirical	10%	.58	.95	.58	.95	.54	.95	.44	.94	.17	.91	.48	.93	.12	.91	.00	.87	.29	.92
	20%	.75	.94	.66	.91	.68	.92	.58	.90	.40	.86	.67	.83	.25	.84	.02	.65	.31	.83
Invariant	10%	.00	.90	.00	.90	.00	.90	.00	.88	.00	.90	.00	.88	.00	.90	1.0	1.0	.00	.89
	20%	.00	.80	.00	.80	.00	.80	.00	.84	.00	.80	.00	.64	.00	.77	1.0	1.0	.00	.75
Patterned 50-50	10%	.57	.95	.01	.89	.02	.89	.08	.90	.26	.92	.38	.89	.25	.93	.02	.87	.00	.89
	20%	.68	.92	.03	.76	.06	.76	.06	.77	.27	.82	.60	.78	.34	.86	.28	.65	.00	.75
Patterned 50-75	10%	.57	.95	.00	.89	.02	.89	.06	.90	.29	.93	.33	.88	.30	.93	.21	.90	.00	.89
	20%	.67	.92	.03	.75	.07	.77	.09	.78	.28	.82	.50	.78	.35	.86	.34	.83	.02	.75
Patterned 75-75	10%	.58	.95	.01	.89	.02	.89	.08	.89	.35	.93	.10	.89	.13	.91	.43	.92	.00	.89
	20%	.69	.92	.02	.75	.04	.76	.11	.79	.42	.86	.26	.75	.19	.82	.67	.87	.00	.75
WP Normal (C)	10%	.58	.95	.57	.95	.55	.95	.43	.94	.24	.92	.25	.88	.00	.90	.08	.87	.10	.90
	20%	.73	.93	.66	.92	.68	.92	.38	.85	.46	.87	.36	.78	.00	.77	.16	.79	.22	.80
WP Normal (M)	10%	.58	.95	.54	.95	.53	.95	.49	.94	.26	.92	.28	.88	.00	.90	.01	.87	.05	.89
	20%	.73	.93	.64	.91	.67	.92	.57	.88	.40	.85	.43	.78	.00	.77	.08	.65	.19	.80
WP Normal (R)	10%	.58	.95	.54	.95	.53	.95	.37	.94	.27	.92	.21	.88	.00	.90	.02	.87	.08	.90
	20%	.74	.93	.63	.91	.67	.92	.57	.88	.48	.87	.35	.78	.00	.77	.12	.65	.19	.79

Note: Proportion of simulated data correctly detected and real data correctly rejected for 10% and 20% mixed data sets. ST = Sensitivity; SF = Specificity

Figure 1. Long-string analysis of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.

Figure 2. Within odd-even analysis of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.

Figure 3. Resampled internal reliability analysis of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.

Figure 4. Psychometric synonym analysis of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.

Figure 5. Polytomous Guttman error analysis of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.

Figure 6. ISD analysis of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.

Figure 7. Person-total correlation analysis of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.

Figure 8. Mahalanobis distance analysis of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.

Supplemental Table 1 – Proportion of simulated data present within the entire human subjects distribution.

Simulation	PTC	WOE	RIR	PS	PA	PG	ISD	LS	MD
Uniform	0.98	0.96	0.84	0.98	0.98	0.87	0.60	1.00	1.00
Empirical	0.98	0.96	0.89	1.00	0.98	0.98	0.14	1.00	0.00
Normal	0.99	0.95	0.88	0.99	1.00	0.99	0.00	1.00	0.00
Inverted	0.98	0.95	0.89	0.99	0.99	0.99	0.79	1.00	1.00
Patterned 50-50	0.73	0.96	1.00	1.00	0.55	1.00	0.21	1.00	0.11
Patterned 50-75	0.69	0.95	0.98	1.00	0.80	1.00	0.24	1.00	0.04
Patterned 75-75	0.67	0.95	0.94	0.99	0.50	1.00	0.11	1.00	0.13
WP Normal (C)	0.99	0.95	0.88	0.91	0.94	1.00	0.00	1.00	0.08
WP Normal (M)	0.98	0.95	0.86	0.95	0.98	1.00	0.00	1.00	0.24
WP Normal (R)	0.98	0.97	0.88	0.95	0.98	1.00	0.00	1.00	0.13

Supplemental Table 2 – Sensitivity and specificity of normal and tiny datasets and metric using 10% cuts.

		PTC		WOE		RIR		PS		PA		PG		ISD		LS		MD	
Simulation		ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF
Uniform	10%	.61	.96	.58	.95	.57	.95	.41	.94	.25	.92	.76	.91	.44	.95	0.0	.87	.55	.95
	T10%	.61	.92	.55	.92	.67	.93	.33	.89	.21	.88	.70	.93	.45	.92	.03	.84	.48	.91
Normal	10%	.58	.95	.55	.95	.55	.95	.45	.94	.25	.92	.35	.87	0.0	.90	0.0	.87	0.0	.89
	T10%	.58	.92	.55	.92	.52	.91	.42	.90	.24	.88	.42	.82	.03	.87	.03	.84	0.0	.86
Inverted	10%	.58	.95	.58	.95	.55	.95	.34	.93	.24	.92	.44	.93	.84	.99	0.0	.87	.94	.99
	T10%	.64	.93	.58	.92	.64	.93	.39	.90	.15	.87	.52	.89	.85	.96	.03	.84	.91	.96
Empirical	10%	.58	.95	.58	.95	.55	.95	.44	.94	.17	.91	.48	.93	.12	.91	0.0	.87	.29	.92
	T10%	.61	.92	.64	.93	.73	.94	.42	.90	.12	.87	.52	.87	.06	.87	.03	.84	.30	.89
Invariant	10%	0.0	.90	0.0	.90	0.0	.90	0.0	.88	0.0	.90	0.0	.88	0.0	.90	1.0	1.0	0.0	.89
	T10%	0.0	.87	.03	.87	.03	.87	.09	.86	0.0	.87	0.0	.84	.03	.87	1.0	.97	0.0	.86
Patterned 50-50	10%	.57	.95	.01	.89	.02	.89	.08	.90	.26	.92	.38	.89	.25	.93	.02	.87	0.0	.89
	T10%	.52	.91	.03	.86	.09	.87	.27	.88	.36	.90	.36	.85	.27	.90	.06	.84	0.0	.86
Patterned 50-75	10%	.57	.95	0.0	.89	.02	.89	.06	.90	.29	.93	.33	.88	.30	.93	.21	.90	0.0	.89
	T10%	.52	.91	.03	.86	.03	.86	.06	.86	.18	.88	.39	.82	.30	.90	.15	.87	0.0	.86
Patterned 75-75	10%	.58	.95	.01	.89	.02	.89	.08	.89	.35	.93	.10	.89	.13	.91	.43	.92	0.0	.89
	T10%	.58	.92	.03	.86	.03	.86	.09	.87	.30	.89	.12	.83	.27	.90	.52	.89	0.0	.86
WP Normal (C)	10%	.58	.95	.57	.95	.54	.95	.43	.94	.24	.92	.25	.88	0.0	.90	.08	.87	.10	.90
	T10%	.55	.92	.55	.92	.64	.93	.48	.91	.30	.89	.18	.84	.03	.87	.09	.87	.15	.87
WP Normal (M)	10%	.58	.95	.54	.95	.52	.95	.49	.94	.26	.92	.28	.88	0.0	.90	.01	.87	.05	.89
	T10%	.55	.92	.58	.92	.55	.92	.58	.91	.30	.89	.15	.86	.03	.87	.06	.84	0.0	.86
WP Normal (R)	10%	.58	.95	.54	.95	.53	.95	.37	.94	.27	.92	.21	.88	0.0	.90	.02	.87	.08	.90
	T10%	.58	.92	.55	.92	.45	.91	.42	.89	.33	.89	.18	.86	.03	.87	.06	.84	.06	.86

Note: Proportion of simulated data correctly detected and real data correctly rejected for 10% mixed data sets. T10% is a dataset where human subjects data was not duplicated and only 33 simulated cases were used. ST = Sensitivity; SF = Specificity

Supplemental Table 3 – Sensitivity and specificity of each dataset and metric using 10% and 5% cuts.

		PTC		WOE		RIR		PS		PA		PG		ISD		LS		MD	
Simulation		ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF	ST	SF
Uniform	10%	.61	.96	.58	.95	.57	.95	.41	.94	.25	.92	.76	.91	.44	.95	0.0	.87	.55	.95
	5%	.32	.98	.38	.99	.38	.99	.25	.97	.13	.96	.42	.98	.22	.97	0.0	.94	.38	.99
Normal	10%	.58	.95	.55	.95	.55	.95	.45	.94	.25	.92	.35	.87	0.0	.90	0.0	.87	0.0	.89
	5%	.30	.98	.38	.99	.37	.99	.34	.98	.10	.96	.21	.94	0.0	.95	0.0	.94	0.0	.94
Inverted	10%	.58	.95	.58	.95	.55	.95	.34	.93	.24	.92	.44	.93	.84	.99	0.0	.87	.94	.99
	5%	.32	.98	.38	.99	.37	.99	.25	.97	.13	.96	.27	.96	.42	1.0	0.0	.94	.50	1.0
Empirical	10%	.58	.95	.58	.95	.55	.95	.44	.94	.17	.91	.48	.93	.12	.91	0.0	.87	.29	.92
	5%	.30	.98	.38	.99	.38	.99	.25	.97	.09	.96	.48	.93	.06	.96	0.0	.94	.16	.96
Invariant	10%	0.0	.90	0.0	.90	0.0	.90	0.0	.88	0.0	.90	0.0	.88	0.0	.90	1.0	1.0	0.0	.89
	5%	0.0	.95	0.0	.95	0.0	.95	0.0	.95	0.0	.95	0.0	.88	0.0	.95	1.0	1.0	0.0	.94
Patterned 50-50	10%	.57	.95	.01	.89	.02	.89	.08	.90	.26	.92	.38	.89	.25	.93	.02	.87	0.0	.89
	5%	.38	.99	0.0	.94	0.0	.94	.02	.94	.15	.96	.27	.94	.20	.97	0.0	.94	0.0	.94
Patterned 50-75	10%	.57	.95	0.0	.89	.02	.89	.06	.90	.29	.93	.33	.88	.30	.93	.21	.90	0.0	.89
	5%	.41	.99	0.0	.94	.01	.95	.02	.94	.20	.97	.22	.96	.22	.97	.08	.94	0.0	.94
Patterned 75-75	10%	.58	.95	.01	.89	.02	.89	.08	.89	.35	.93	.10	.89	.13	.91	.43	.92	0.0	.89
	5%	.38	.99	.01	.94	.02	.95	.05	.94	.21	.97	.10	.89	.10	.96	.17	.96	0.0	.94
WP Normal (C)	10%	.58	.95	.57	.95	.54	.95	.43	.94	.24	.92	.25	.88	0.0	.90	.08	.87	.10	.90
	5%	.29	.98	.38	.99	.36	.98	.34	.97	.10	.96	.12	.94	0.0	.95	.04	.94	.02	.95
WP Normal (M)	10%	.58	.95	.54	.95	.52	.95	.49	.94	.26	.92	.28	.88	0.0	.90	.01	.87	.05	.89
	5%	.31	.98	.36	.98	.37	.99	.33	.98	.10	.96	.12	.94	0.0	.95	0.0	.94	.01	.94
WP Normal (R)	10%	.58	.95	.54	.95	.53	.95	.37	.94	.27	.92	.21	.88	0.0	.90	.02	.87	.08	.90
	5%	.30	.98	.35	.98	.36	.98	.28	.98	.10	.96	.10	.94	0.0	.95	.01	.94	.01	.94

Note: Proportion of simulated data correctly detected and real data correctly rejected for 10% mixed data sets under 10% and 5% cut criteria. ST = Sensitivity; SF = Specificity

Supplementary Figure 1. Psychometric antonyms of simulated vs human subjects data. Note: Red bars represent 90th and 100th percentile of human subjects data.