**Pairwise comparison using a Bayesian selection algorithm:
Efficient holistic measurement**

Elise A. V. Cromptvoets^{1,2}, Anton A. Béguin³, Klaas Sijtsma¹

¹Tilburg University, the Netherlands

²Cito, the Netherlands

³International Baccalaureate, United Kingdom

Corresponding Author:

Elise Cromptvoets. E-mail: e.a.v.cromptvoets@uvt.nl. Phone: 0031134664328. Institutional address Tilburg University: P.O. Box 90153, 5000 LE Tilburg, The Netherlands.

Abstract

The method of pairwise comparison has been used in a wide range of contexts. In educational measurement, many pairwise comparisons are required for reliable measurement when the comparisons to be performed are selected using the commonly-used semi-random selection algorithm (SSA). We proposed a Bayesian selection algorithm (BSA) to obtain smaller standard errors of parameter estimates and higher reliability compared with the SSA, and we evaluated the performance of these algorithms in a simulation study. We conclude that 1) the BSA should be preferred to the SSA, 2) the number of comparisons required for reliable measurement depends on the object variance, and 3) the Scale Separation Reliability (SSR) may systematically overestimate reliability even when the SSA is used.

Keywords: adaptive measurement, Bayesian estimation, comparative judgment, holistic measurement, pairwise comparison

Pairwise comparison using a Bayesian selection algorithm: Efficient holistic measurement

Imagine that your son or daughter has made five drawings, and you want to keep the most creative one as a souvenir. You most likely compare the drawings in pairs and decide which drawing to get rid of, until one drawing is left. In this situation, usually without realizing it, you used pairwise comparison to rank the drawings with respect to creativity. This is just one of many situations where people use pairwise comparisons in their daily life. Pairwise comparison can be used for many different attributes, and creativity is just one (e.g., Jones, Swan, & Pollitt, 2015; Newhouse, 2014; Steedle & Ferrara, 2016; Van Daal et al., 2016). Comparing multiple objects in pairs can create a ranking of more than two objects (Thurstone, 1927). Because people perform pairwise comparisons or at least some form of comparative judgment on a daily basis (Laming, 2004) and usually do this unconsciously, pairwise comparison is highly intuitive, and each comparison can be performed in a short amount of time.

The method of pairwise comparison has been used in a wide range of contexts. Examples include chess or sports games where players or teams are ranked based on one to one comparisons, marketing research where several brands of a product are ranked based on pairwise comparisons of the brands performed by consumers (Cattelan, 2012), and education where student assignments are ranked based on pairwise comparison of these assignments performed by assessors or peers (e.g. Jones & Alcock, 2013; Lesterhuis et al., 2017). Different contexts may require an adaptation of the general pairwise comparison methodology to obtain relevant results. For example, in chess ratings, changes in skill level between comparisons are allowed by assigning higher weights for more recent comparisons, and for some sports games, home advantage is taken into account.

Recently, in education, more attention has been given to the development and measurement of attributes like creativity and other twenty-first century skills. These attributes are complex in that they are difficult to disentangle into sub-aspects together defining the attributes. For creativity, this means that the ‘total’ creativity is more than a conglomerate of aspects such as use of colors and shapes. For this reason, they are difficult to measure validly using analytical scoring schemes such as rubrics or criteria lists (Van Daal et al., 2016), which are conventional measurement methods used in education. In contrast to these analytic measurement methods, which assume that an attribute is a list of sub aspects and evaluate each aspect separately, pairwise comparison is a holistic measurement method where an attribute is evaluated as a whole (Pollitt, 2012).

Unfortunately, using the ‘standard’ random pair selection, a large number of comparisons are required for reliable measurement. Making the comparison process adaptive was proposed to obtain reliable measurement using fewer comparisons (Pollitt, 2012). However, the initial adaptive algorithms did not sufficiently take the uncertainty of the object parameters into account, resulting in biased estimates (Bramley, 2015; Bramley & Vitello, 2018). Cromptvoets, Béguin, & Sijtsma (2020) proposed an adaptive selection algorithm that reduced standard errors of parameter estimates, but did not improve reliability or rank order accuracy compared with a commonly-used semi-random algorithm. In this study, we build on the research of Cromptvoets et al. (2020), proposing a new Bayesian selection algorithm (BSA) to obtain both smaller standard errors of parameter estimates and higher reliability compared with a commonly-used semi-random algorithm and testing it in a simulation study that builds on their simulation design and evaluation criteria.

This article is organized as follows. First, we describe the parameter estimation procedure using the Bayesian framework. Second, we present the BSA, which was based on the parameter estimation procedure. Third, we describe the set-up of our main simulation study and the evaluation criteria we used, followed by a description of the results. Then, we describe an explorative study that addresses a remarkable result with respect to reliability found in the main study. The article ends with a discussion.

Parameter Estimation

Thurstone's (1927) probit model for pairwise comparisons was used to create a scale of objects based on comparison data. Let N be the number of objects, and let i and j ($i, j = 1, \dots, N$) be object indices. Thurstone's model defines the probability that object i is preferred to object j , denoted by $i > j$, in a paired comparison by means of

$$P(i > j | \theta_i, \theta_j) = \Phi(\theta_i - \theta_j). \quad (1)$$

where θ_i and θ_j are the parameters of objects i and j , respectively, and $\Phi()$ denotes the standard normal cumulative distribution.

We estimated the object parameters using Markov chain Monte Carlo (MCMC) sampling. The MCMC chains were constructed using the Gibbs sampler (Gelfand & Smith, 1990). We adjusted the data augmented Gibbs sampling procedure presented by Albert (1992) for estimating the two-parameter normal ogive model (2-PNO; Lawley, 1943; Lord, 1952)

$$P(i > j | \theta_i, \alpha_j, \delta_j) = \Phi(\alpha_j \theta_i - \delta_j). \quad (2)$$

to estimate the parameters of Thurstone's model. Removing the discrimination parameter α_j from this model yields the one-parameter normal ogive model (1-PNO). Subsequently replacing the location parameter of item j , δ_j , with the (location) parameter of object j , θ_j , produces Thurstone's model. It may be noted that whereas the normal ogive models contain both structural

and incidental parameters (see Baker & Kim, 2004, for a description of structural and incidental parameters), the distinction between structural and incidental parameters is less clear for Thurstone's model, because both model parameters say something about the same type of entity, namely objects. The distinction between structural and incidental parameters is not necessary for our purposes.

Albert (1992) augmented the data $\mathbf{X} = (X_{11}, \dots, X_{nk})$ from n persons on k items with latent parameters $\mathbf{Z} = (Z_{11}, \dots, Z_{nk})$, representing the continuous distribution underlying the dichotomous outcome X_{ij} , where Z_{ij} follows a normal distribution that is truncated at $Z_{ij} = 0$. The Z_{ij} -distribution includes only positive values if person i answered item j correctly, and the Z_{ij} -distribution includes only negative values otherwise. For pairwise comparisons, data were augmented by the latent parameters $\mathbf{Z} = (Z_1, \dots, Z_C)$, where C denotes the number of comparisons. Let c ($c = 1, \dots, C$) be a comparison index, let X_c be the binary outcome of comparison c between object i and object j , where $X_c = 1$ if object i is preferred to object j and $X_c = 0$ otherwise, and let $\mathbf{X}$ be the data vector containing all X_c . Then, the conditional distribution of Z_c can be written as

$$Z_c | \boldsymbol{\theta}, \mathbf{X} \sim \begin{cases} N(\theta_i - \theta_j, 1) \text{ truncated left at } 0 \text{ for } X_c = 1 (Z_c > 0) \\ N(\theta_i - \theta_j, 1) \text{ truncated right at } 0 \text{ for } X_c = 0 (Z_c \leq 0) \end{cases} \quad (3)$$

Examples of conditional Z_c -distributions are displayed in Figure 1.

To implement the Gibbs sampler, the parameter vector was divided into two components, the object parameters $\boldsymbol{\theta}$ and the data augmented parameters $\mathbf{Z}$. The parameters were sampled in an iterative chain. In every iteration, each component was sampled from its conditional distribution given sampled values of the other component. This sampling scheme was repeated until the sampled values formed stable posterior distributions.

Figure 1. Two examples of truncated conditional Z_c -distributions for $(\theta_i - \theta_j) = 1$. Dashed line for right-truncated distribution (if $X_c = 0$) and dotted line for left-truncated distribution (if $X_c = 1$).

As an ingredient for the posterior distribution, we derived the likelihood function and the posterior distribution for Thurstone's model from the likelihood function and the posterior distribution of the 2-PNO as described by Albert (1992). The likelihood function of the 2-PNO can be written as

$$\theta_i \sim N\left(\frac{1}{\sum_j^K \alpha_j^2} \times \sum_{j=1}^K \alpha_j (Z_{ij} + \delta_j), \frac{1}{\sum_j^K \alpha_j^2}\right) \quad (4)$$

where K is the number of items. First, we fixed the α_j parameter to 1 to obtain the likelihood of the 1-PNO. This substitution implies that $\sum_j^K \alpha_j^2 = K$. Second, we replace the location parameter of item j , δ_j , with the (location) parameter of object j , θ_j . In addition, the 'test' now consists of comparisons instead of items, so we replace the number of items in the test with the number of comparisons including object i , which can be different for different objects i . Let C_i be the number of performed comparisons including object i , and let c_i ($c_i = 1, \dots, C_i$) be the index for comparisons including object i . The likelihood function for θ_i for Thurstone's model is

$$\theta_i \sim N\left(\frac{1}{C_i} \times \sum_{c_i=1}^{C_i} (Z_{c_i} + \theta_{j_{c_i}}), \frac{1}{C_i}\right). \quad (5)$$

When a normal prior distribution is used for θ_i with mean μ_i and variance σ_i^2 , multiplying this prior with the likelihood gives the conditional posterior distribution of θ_i ,

$$\theta_i | \boldsymbol{\theta}, \mathbf{Z}, \mathbf{X} \sim N \left(\frac{1}{C_i + 1/\sigma_i^2} \times \left(\sum_{c_i=1}^{C_i} (Z_{c_i} + \theta_{j_{c_i}}) + \frac{\mu_i}{\sigma_i^2} \right), \frac{1}{C_i + 1/\sigma_i^2} \right) \quad (6)$$

$$\theta_i | \boldsymbol{\theta}, \mathbf{Z}, \mathbf{X} \propto \phi(\theta_i; \mu_i, \sigma_i^2) \prod_{c_i=1}^{C_i} \phi(Z_{c_i}; \theta_i - \theta_{j_{c_i}}, 1)$$

following Equations (27) and (25) respectively from Albert (1992). In this study, we used a standard normal prior, $\theta_i \sim N(0, 1)$, for all objects, so the posterior reduces to

$$\theta_i | \boldsymbol{\theta}, \mathbf{Z}, \mathbf{X} \sim N \left(\frac{1}{C_i + 1} \times \left(\sum_{c_i=1}^{C_i} (Z_{c_i} + \theta_{j_{c_i}}) \right), \frac{1}{C_i + 1} \right) \quad (7)$$

$$\theta_i | \boldsymbol{\theta}, \mathbf{Z}, \mathbf{X} \propto \phi(\theta_i; 0, 1) \prod_{c_i=1}^{C_i} \phi(Z_{c_i}; \theta_i - \theta_{j_{c_i}}, 1).$$

We used the marginal posterior means as object parameter estimates, and we used the standard deviations of the marginal posterior distributions as standard errors of the object parameter estimates. Using MCMC sampling, we computed the marginal posterior mean for each θ_i as the mean of the sampled θ_i -values in the MCMC chain after burn-in. Similarly, we computed the standard deviation of the marginal posterior density as the standard deviation of the sampled θ_i -values in the MCMC chain after burn-in. We used a burn-in of 100 and a chain length after burn-in of 5000. Visual inspection of several MCMC chains showed that the chains indeed converged within 100 iterations.

Bayesian Selection Algorithm (BSA)

The goal of the BSA is to select the pair of objects for the rater to compare that contains most information about the rank order, given the results of previous comparisons. Most information can be obtained for objects that have the smallest distance on the attribute scale given the available objects. Therefore, in each step, we aimed at selecting the two objects that are closest together on the attribute scale. In practice, we need to be aware of the risk of chance capitalization based on the previous results. For this reason, we took the uncertainty of parameter estimates into account. We used a sampling-based method to incorporate the uncertainty of the parameters during data collection.

First, a sample of θ is drawn from the conditional posterior distribution from an MCMC iteration after burn-in. Second, for all possible pairs, the differences between the θ -values ($\theta_i - \theta_j$) was computed. If a comparison was already performed by every rater, this comparison was removed to avoid that a rater would compare the same objects twice. The smallest difference ($\theta_i - \theta_j$) corresponds with the smallest distance between the (sampled) object parameter values on the attribute scale, which should provide the most information. After the two selected objects were compared, the data and comparison counts were updated and the selection algorithm steps were repeated until the predefined stopping criterion was reached. This Bayesian algorithm automatically takes the uncertainty of the parameters into account, because parameter values are sampled from the conditional posterior distributions.

Main study

Method

Design. We evaluated a completely crossed design with five different factors, producing 180 conditions in total. First, we varied the selection algorithm, using the BSA and a baseline algorithm to which the results of the newly developed algorithm were compared. In the baseline

algorithm, which is the semi-random selection algorithm (SSA), a pair of objects is randomly selected with the constraint that the objects in the pair were not previously compared to each other. After the two selected objects were compared, the selection algorithm was repeated until the predefined stopping criterion was reached. In the final step, the object parameter estimates and corresponding standard errors were computed using Gibbs sampling.

Second, we varied the average number of comparisons per object, C . Previous studies provide varying recommendations for the number of comparisons per object required for reliable measurement. The recommendations range from $C = 9$ (Pollitt, 2012), to $C = 12$ (Verhavert et al., 2019), to $C = 20$ based on average reliability and $C = 30$ based on the lower bound of a 68% confidence interval for reliability equal to .80 (Cromptvoets et al., 2020). Furthermore, a meta-analysis of Verhavert et al. (2019) suggested that 10 to 14 comparisons are required for a reliability equal to .70, but 26 to 37 comparisons for reliability of .90. In this study, we used $C = \{10, 15, 20, 25, 30\}$.

Third, we varied the standard deviation of the objects to be compared, using $SD = \{0.5, 1, 3\}$. Previous studies have found varying standard deviations of the object parameters. For example, Van Daal et al. (2017) reported standard deviations for different samples ranging from 1.28 to 1.58, while Pollitt (2012) reported a standard deviation of 3.85. Although the standard deviation reported in Pollitt's study is likely overestimated due to the use of adaptive pairwise comparisons, we wanted to investigate the influence of such standard deviation values of the object parameters on the evaluation criteria.

Fourth, we varied the number of raters, using two, three, and five raters. We assumed that all raters performed an equal number of comparisons. For example, if 20 objects are each compared ten times, two raters would each perform 50 comparisons by the following line of

reasoning. Each comparison involves two objects. The total number of comparisons to be performed then equals $20 \times \frac{10}{2} = 100$. Dividing this number of comparisons equally over the two raters means that the raters would each perform $\frac{100}{2} = 50$ comparisons.

Fifth, we varied rater agreement as either perfectly model compliant, where the comparisons of all raters are based on the preference probabilities in the generating model, or imperfectly model compliant, where one rater does not comply perfectly with the generating model. Although perfect model compliance for all raters is unlikely, we included this condition because it facilitates comparison with previous research (Crompvoets et al., 2020). We simulated less than perfect model compliance by adding a noise component to the true values of the latent variable for one of the raters: $\theta_{noisy} = \sqrt{0.2} \cdot noise + (1 - \sqrt{0.2}) \cdot \theta$, where the noise is sampled from a standard normal distribution for each θ . In this equation, 80 percent of the variance of θ is true variation among the objects and 20 percent of the variance of θ is noise. In short, we call this the ‘noisy rater’ condition.

For each of the 2 (Algorithm) x 5 (Number of Comparisons per Object) x 3 (SD of Objects) x 3 (Number of Ratets) x 2 (Rater Agreement) = 180 design conditions, we drew 30 object parameters from the standard normal distribution. Because the mean of the object parameter estimates is constrained to zero for model identification, we rescaled the object parameters to have a sample mean of zero as well. Subsequently, the probabilities that the objects are rated higher on the latent variable scale than other objects were computed by inserting their true parameters in Thurstone’s model (Equation 1).

In each design condition, for each comparison, two objects were selected based on the selection algorithm. The comparison of the two objects was simulated by comparing a random value between 0 and 1 to the probability that a rater prefers object i to object j . Object i was

chosen if the random value was smaller than the probability value, and object j was chosen otherwise. These steps were repeated until the maximum number of comparisons in the condition was reached. After reaching the maximum number of comparisons, object parameter estimates and standard errors were computed. Lastly, we computed the parameter uncertainty and reliability of the scale. This procedure, starting with drawing 30 object parameters in all conditions, was repeated 400 times. The number of repetitions was based on a previous simulation study (Crompvoets et al., 2020). For both data simulation and data analysis, Thurstone’s probit model was used. The R code of this study can be found in the supplementary material.

Evaluation criteria. We used the uncertainty of the parameters and the reliability of the scale as evaluation criteria. These are the same evaluation criteria as Crompvoets et al. (2020). However, rank order accuracy showed the same trend as reliability in the study of Crompvoets et al. (2020), which we expect to be the same in this study, so we only evaluated reliability and not rank order accuracy. We expect the BSA developed in this study to result in similar values as the current standard (an SSA algorithm) for these outcomes but using fewer comparisons because that was the purpose of developing the BSA.

Uncertainty of parameters. We evaluated the uncertainty of the parameters using the standard deviation of the posterior distributions of the object parameters. We expected that parameter uncertainty was smaller for larger numbers of comparisons C , and because the number of unique comparisons grows multiplicatively with the number of objects via the multiplication factor $N(N - 1)/2$, we expected this effect to be larger for larger numbers of objects. In addition, we expected that parameter uncertainty was larger using the SSA than the BSA.

Reliability. We used two measures of reliability. First, we used the squared correlation between the object parameters used in the generating model and the object parameter estimates based on the data, which we refer to as the benchmark reliability. Let θ be the object parameter in the generating model and let $\hat{\theta}$ be the object parameter estimate. The benchmark reliability can then be computed as

$$\rho_{\hat{\theta}\theta} = \text{cor}(\theta, \hat{\theta})^2. \quad (8)$$

Second, we used the commonly used Scale Separation Reliability (SSR) estimate (Bramley, 2015). Let $S^2(\theta)$ be the estimated variance of the object parameters in the generating model, and let $SE(\hat{\theta})$ be the estimated variance of the estimated object parameters. Further, let $MSE[SE(\hat{\theta})]$ be the mean squared error of the standard errors corresponding to the object parameter estimates. The SSR is computed as follows:

$$SSR = \frac{S^2(\theta)}{S^2(\hat{\theta})}, \quad (9)$$

where

$$S^2(\theta) = S^2(\hat{\theta}) - MSE[SE(\hat{\theta})],$$

that is, the observed variance minus an error term (Bramley, 2015).

The number of comparisons was expected to increase reliability by decreasing the standard errors of the object parameter estimates. We also expected the reliability to be higher using the BSA rather than using the SSA.

Results

The use of different numbers of raters showed only small differences in the uncertainty of the parameters. Thus, we will only display the results for two raters. For two raters and an SD = 3 for the object parameters in the generating model, Figure 2 displays means of standard errors

Figure 2. Standard errors of object parameter estimates for both selection algorithms, for two raters, $SD = 3$, for different numbers of comparisons, and with and without a 'noisy rater'. Means of standard errors (dots) for object parameter estimates within ranges of 0.5 units on the latent variable scale.

(dots) for ML estimates within ranges of 0.5 units on the latent variable scale. The columns of Figure 2 show that more comparisons per object resulted in lower standard errors for both algorithms. The rows of Figure 2 show that compared to the noisy rater condition (bottom row), there were small differences in parameter uncertainty when there was perfect model compliance (top row). Mostly, parameter uncertainty was greater at both ends of the scale in the noisy rater condition, but the order of magnitude was quite small. These results indicate that the accuracy of the comparison may have little influence on the uncertainty of the parameters, especially compared to the influence of the number of comparisons. In other words, adding more comparisons can decrease parameter uncertainty, even if the quality of the comparisons may be suboptimal, that is, if the comparison results do not comply with the generating model.

The most interesting result with respect to parameter uncertainty was that the BSA resulted in lower or similar parameter uncertainty compared to the SSA. This gain in parameter certainty was mostly found in the middle of the latent variable scale, and was largest for a $SD = 3$ (Figures 2 and 3). For $SD = 0.5$ and $SD = 1$, Figure 3 shows for 20 comparisons (i.e., $C = 20$), per object that the BSA can result in higher parameter uncertainty than the SSA at the ends of the scale if, but the parameter uncertainty was lower for the BSA than for the SSA in the middle of the latent variable scale. That is, when the BSA was used rather than the SSA, parameter uncertainty shifts from the middle to both ends of the scale. When averaged over the full latent variable scale, the BSA results in similar parameter uncertainty as the SSA for $SD = 0.5$, slightly lower parameter uncertainty for $SD = 1$, and substantially lower parameter uncertainty for $SD = 3$. In general, parameter uncertainty was lower for smaller SDs. The explanation is that comparisons contain more statistical information when the objects are closer together on the latent variable scale.

Figure 3. Standard errors of object parameter estimates for both selection algorithms for two raters and $C = 20$. Means of standard errors (dots) for object parameter estimates within ranges of 0.5 units on the latent variable scale.

An increasing number of comparisons led to an increase in both benchmark reliability and the SSR (Figure 4). Another trend across conditions was that both benchmark reliability and SSR increased with larger standard deviation of the latent variable values. In the noisy rater conditions, benchmark reliability was lower than in the perfect model compliance conditions, but the SSR was similar. When the BSA was used, the SSR was always higher than when the SSA was used. For differences in benchmark reliability between the two algorithms, the results differed per true SD condition. For $SD = 0.5$, the benchmark reliability was similar or slightly lower when the BSA was used compared to the benchmark reliability when the SSA was used.

For $SD = 1$, the benchmark reliability was slightly higher for the BSA than the SSA. For $SD = 3$, the benchmark reliability was substantially higher for the BSA than the SSA. We conclude that BSA performed better than the SSA with respect to benchmark reliability when SD was moderate or high and similar when SD was low.

Figure 4. Benchmark reliability and estimated reliability for two raters.

In the perfect model compliance conditions, reliability was equal among the different numbers of raters. In the noisy rater conditions, the number of raters affected benchmark reliability, but barely affected the SSR (Figure 5). Benchmark reliability was higher when more raters performed comparisons. The explanation is that one rater who does not comply with the model has a smaller influence on the scale when there are more raters who do comply with the

Figure 5. Benchmark reliability and estimated reliability for imperfect model compliance.

Running head: PAIRWISE COMPARISON USING A BAYESIAN SELECTION ALGORITHM

model. The result that the SSR is not affected in the same way is surprising. Apparently, the number of comparisons is more important for the SSR than the accuracy of the comparisons. Lastly, a result with consequences beyond this study, is that except for the combinations of perfect model compliance and $SD = \{1, 3\}$, the SSR was higher than the benchmark reliability for both algorithms (Figure 6). This means that the SSR overestimated reliability when the true SD was low, but also when one rater did not perfectly comply with the model. A reliability estimate should not overestimate reliability due to bias (Sijtsma, 2009), but this requirement is violated by the SSR in these scenarios. It is especially striking that this is also the case for the SSA, which is usually viewed as the ‘safe route’ to using pairwise comparison such that the SSR can be trusted.

Figure 6. Bias of the SSR for two raters. The bias is computed as Benchmark reliability – SSR.

Exploratory study

To follow up on the reliability results, we investigated the estimated variance of the object parameter estimates, which appears as $S^2(\hat{\theta})$ in Equation (9). Crompvoets et al. (2020) showed that for their adaptive algorithm, the SSR overestimated reliability due to an overestimation of the variance of the object parameters. Figure 7 shows the average SD of object parameter estimates for two raters, one of which was a noisy rater. The figure shows that the variance was overestimated slightly for $SD = 0.5$ and grossly for $SD = 3$. Especially the pattern of $SD = 3$ indicates that the prior might have had substantial influence on the estimated variance of the latent attribute parameters. For this reason, we conducted an exploratory study using less informative priors to investigate the influence of the prior on the results.

Figure 7. Average SD of object parameter estimates for two raters and imperfect model compliance.

Method

In our exploratory study, we investigated four different prior standard deviations used in the estimation procedure. We studied the prior used in the main study, which is $\theta \sim N(0, 1)$, and

PAIRWISE COMPARISON USING A BAYESIAN SELECTION ALGORITHM

we also studied normal priors with larger variances: $\theta \sim N(0, 3)$, $\theta \sim N(0, 5)$, and $\theta \sim N(0, 10)$.

We used almost the same procedure as in the main study. However, we only used a selection of conditions that could give valuable information about prior sensitivity.

1. We used two raters. This condition shows the effect of noise most clearly, and we expect that the results for the other numbers of raters are less pronounced.
2. We used $C = \{10, 20, 30\}$, leaving out 15 and 25 comparisons. In addition, when the condition of 10 comparisons per object was finished, we kept those observations and sampled 10 new comparisons per object to complete the 20 comparisons per object condition, and we sampled 10 new comparisons per object to complete the 30 comparisons per object condition. By doing this, we used the same samples of object parameters and the same noisy raters in the noise conditions for all numbers of comparisons per object conditions, and we used the observations of the smaller numbers of comparisons per object conditions in the larger numbers of comparisons per object conditions. This process is more efficient, but the results are not independent. For this exploratory study, we deemed this process sufficient.
3. We used 50 repetitions, because we reasoned that this number would sufficiently show the influence of prior sensitivity.

Results

The results showed that the standard errors of the object parameter estimates were larger when the prior variance was wider (Figure 8). The standard errors of the object parameter estimates using prior SD = 3 were large at the ends of the scale reaching values of 1.5. The standard errors of the object parameter estimates using a prior SD of 5 were medium in the

PAIRWISE COMPARISON USING A BAYESIAN SELECTION ALGORITHM

middle of the scale with values between 0.5 and 1, but they were large at the ends of the scale reaching values of 2.5. The standard errors of the object parameter estimates using prior SD = 10 were excessive with values above 10. Thus, in order to obtain reasonable standard errors, a prior SD = 1 as we used in the main study is best. A prior SD up to a value of 3 can be deemed ‘acceptable’, but wider priors provide too little information to obtain reasonable standard errors. This result was stable across noise conditions and SD conditions.

Figure 8. Standard errors of object parameter estimates for two raters, perfect model compliance, $SD(\theta) = 3$. Means of standard errors (dots) for object parameter estimates within ranges of 0.5 units on the latent variable scale.

In general, benchmark reliability decreased when the prior SD increased. As in the main study, when prior SD = 1, the SSR was lower than the benchmark reliability for objects SD = 1 and SD = 3, but not for objects SD = 0.5. When prior SD = 3 or prior SD = 5, the SSR was lower than the benchmark reliability for SD = 3, but not for SD = 0.5 or SD = 1 (Figure 9). When prior SD = 10, both SSR and benchmark reliability were unstable. Prior SD = 1 seems to

PAIRWISE COMPARISON USING A BAYESIAN SELECTION ALGORITHM

perform best. However, it is undesirable that the prior SD can determine whether the SSR behaves as a lower bound to reliability as it should.

Figure 9. Benchmark reliability and estimated reliability for two raters, perfect model compliance, and a prior SD equal to 3.

Discussion

In general, the BSA performed similar to the SSA or better with respect to both object parameter uncertainty and benchmark reliability. With respect to both parameter uncertainty and reliability, the BSA produced the largest benefits when the objects differed more on the latent attribute scale. For this reason, we recommend the BSA rather than the SSA when choosing a selection algorithm for data collection using pairwise comparison.

Consistent with Bramley and Vitello (2018), the SSR overestimated reliability in some scenarios. However, whereas Bramley and Vitello (2018) claim that this happens when an adaptive selection algorithm is used, we showed that the SSR can also overestimate reliability in some situations when a semi-random selection algorithm is used. It seems that both parameter uncertainty and the SSR are strongly influenced by the number of comparisons, but the accuracy

PAIRWISE COMPARISON USING A BAYESIAN SELECTION ALGORITHM

of the comparisons is not influenced. Hence, the SSR is not a trustworthy measure of reliability for pairwise comparison data. Future research should aim at finding another measure of reliability for pairwise comparison data that can be trusted.

The number of comparisons per object, C , that is required for sufficient reliability depends on the variance of the object parameters. The more the objects differ on the latent attribute scale, the smaller the required number of comparisons per object. When the standard deviation of the object parameters was $SD = 3$, (benchmark) reliability above .8 could already be achieved with $C = 10$, even when one of the raters did not comply perfectly with the generating model, meaning that he or she had an opinion different from the simulated truth. This result might look like an open door, but it is inconsistent with the study of Verhavert et al. (2019), who provided a general guideline for the number of comparisons required for a certain level of reliability irrespective of the variance of the objects. Of course, because they used real data, Verhavert et al (2019) did not know the true variance of the objects, so that they could not compute a benchmark reliability. In our study, the SSR was also dependent on the variance of the objects. Thus, the recommendations from Verhavert (2019) about the number of comparisons per object required for sufficient reliability should not be used blindly. One should realize that the variance of the objects in your situation might differ from the variance of the objects in the (average) situation Verhavert et al. (2019) studied, which could lead to a greater or smaller number of comparisons per object required for sufficient reliability. Lastly, we note that the recommendations from Verhavert et al. (2019) made should also be used with caution, because their results were based on the SSR reliability estimate, which we showed to be untrustworthy.

The benchmark reliability may show a pessimistic view of reality in the noisy rater condition, because the truth was simulated at the performance of the good rater(s). This means

PAIRWISE COMPARISON USING A BAYESIAN SELECTION ALGORITHM

that if the results from the good and noisy raters are averaged, which happens when computing the reliability, it is quite difficult to arrive at the truth. Therefore, the results could be less pessimistic if the truth was somewhere in the middle. The results are dependent on the way we chose to manipulate imperfect model compliance, that is, the noisy rater condition, and other choices could have been made. However, this would be true with any choice of manipulating model compliance, because there are many (incomparable) ways to vary the agreement among raters, let alone in relation to the truth, and it is unfeasible to include them all in one simulation study. Furthermore, the main goal of the study was exploring the performance of the BSA, and in a small study we also checked the results for $\theta_{noisy} = \sqrt{0.1} \cdot noise + (1 - \sqrt{0.1}) \cdot \theta$ and we found the same trend as in the main study but smaller. Therefore, we assume that the results for our model choice for imperfect model compliance are relevant for ‘imperfect model compliance’.

Our study has three takeaway messages. First, the BSA should be preferred to the SSA for data collection using pairwise comparison. Second, the number of comparisons required for reliable measurement depends on the variance of the objects. This result can be at one’s advantage or disadvantage depending on the object variance, but means that one should take object variance into account when deciding on the number of comparisons per object. Third, the SSR may systematically overestimate reliability even when a semi-random selection algorithm is used. Therefore, any statement about reliability using the SSR should be interpreted with caution, and additional research, momentarily in progress, should aim at finding an alternative to the SSR that does not systematically overestimate reliability.

References

- Albert, J. H. (1992). Bayesian estimation of normal ogive item response functions using Gibbs sampling. *Journal of Educational Statistics*, 17, 251-269.
- Baker, F.B. & Kim, S.H. (2004). Item Response Theory: Parameter estimation techniques. CRC Press, Boca Raton. doi: 10.1201/9781482276725
- Bramley, T. (2015). *Investigating the reliability of adaptive comparative judgement*. Cambridge Assessment Research Report. Cambridge, United Kingdom: Cambridge Assessment.
- Bramley, T. & Vitello, S. (2018). The effect of adaptivity on the reliability coefficient in adaptive comparative judgement. *Assessment in Education: Principles, Policy & Practice*, 2018, 1-16. doi: 10.1080/0969594X.2017.1418734.
- Cattelan, M. (2012). Models for paired comparison data: A review with emphasis on dependent data. *Statistical Science*, 27, 412-433. doi: 10.1214/12-STS396
- Crompvoets, E. A. V., Béguin, A. A., & Sijtsma, K. (2020). Adaptive pairwise comparison for educational measurement. *Journal of Educational and Behavioral Statistics*, 45, 316-338. doi: 10.3102/1076998619890589
- Gelfand, A. E., & Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* 85, 398-409.
- Jones, I., Swan, M., & Pollitt, A. (2015). Assessing mathematical problem solving using comparative judgement. *International Journal of Science and Mathematics Education*, 13, 151-177.
- Laming, D. (2004). *Human judgment: The eye of the beholder*. London, United Kingdom: Thomson Learning.

PAIRWISE COMPARISON USING A BAYESIAN SELECTION ALGORITHM

Lawley, D. N. (1943). On problems connected with item selection and test construction.

Proceedings of the Royal Society of Edinburgh, 61, 273-287.

Lesterhuis, M., Verhavert, S., Coertjens, L., Donche, V., & De Mayer, S. (2017). Comparative judgement as a promising alternative to score competences. In: *Innovative practices for higher education assessment and measurement*. Hershey, PA: IGI Global. doi: 10.4018/978-1-5225-0531-0.ch007.

Lord, F. M. (1952). A theory of test scores. *Psychometric Monograph* 7.

Newhouse, C. P. (2014). Using digital representations of practical production work for summative assessment. *Assessment in Education: Principles, Policy & Practice*, 21, 205–220. doi: 10.1080/0969594X.2013.868341.

Pollitt, A. (2012). The method of adaptive comparative judgement. *Assessment in Education: Principles, Policy & Practice*, 19, 281-300. doi: 10.1080/0969594X.0962012.0665354.

R Core Team (2018). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. URL: <https://www.R-project.org/>.

Sijtsma, K. (2009). On the use, the misuse, and the very limited usefulness of Cronbach's alpha. *Psychometrika*, 74, 107-120. doi: 10.1007/s11336-008-9101-0.

Steedle, J. T. & Ferrara, S. (2016). Evaluating comparative judgment as an approach to essay scoring. *Applied Measurement in Education*, 29, 211-223. doi: 10.1080/08957347.2016.1171769

Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review*, 34, 273–86.

Van Daal, T., Lesterhuis, M., Coertjens, L., Donche, V., & De Maeyer, S. (2016). Validity of comparative judgement to assess academic writing: Examining implications of its holistic

character and building on a shared consensus. *Assessment in Education: Principles, Policy & Practice*, 2016, 1-16. doi: 10.1080/0969594X.2016.1253542.

Verhavert, S., Bouwer, R., Donche, V., & De Mayer, S. (2019). A meta-analysis in the reliability of comparative judgement. *Assessment in Education: Principles, Policy and Practice*, 26, 541-562. doi: 10.1080/0969594X.2019.1602027