**Predicting Future Study Success: How A Bayesian Approach can be of Help to
Analyze and Interpret Admission Test Scores**

Jorge N. Tendeiro

A. Susan M. Niessen

Daniela R. Crisan

Rob R. Meijer

Paper written for the Law School Admission Council

Jorge N. Tendeiro, A. Susan M. Niessen, Daniela Crisan, and Rob R. Meijer

Department Psychometrics and Statistics, Faculty of Behavioral and Social Sciences,
University of Groningen.

Correspondence concerning this article should be addressed to Jorge N. Tendeiro

Department of Psychometrics and Statistics, Faculty of Behavioral and Social
Sciences, University of Groningen, Grote Kruisstraat 2/1, 9712 TS Groningen, The
Netherlands. Email: j.n.tendeiro@rug.nl

Executive Summary

The aim of this study was twofold: First, we investigated whether scores on an admission test lead to similar predictions in future study success when administered in a proctored- and an unproctored setting. Second, we explored how Bayesian modeling can be of help in interpreting admission-testing data. Results showed that the mode of administration of an admission test did not result in different models for predicting study success and that Bayesian modeling provide a very useful- and easy-to-interpret framework for predicting the probability of future study success.

Arguably the most important aim of admission testing is the prediction of future academic success. Academic success is typically operationalized as GPA or study progress, but can also include leadership or citizenship (e.g., Stemler, 2012; Sternberg, 2010). In order to accept those students with the highest academic potential, students are admitted to college or graduate programs based on admission criteria such as scores on admission tests and other possible predictors such as high school performance (in the case of undergraduate admissions), undergraduate performance (in the case of graduate school admissions), biodata (such as life and work experience), personal statements, recommendations, and interviews (Clinedinst & Patel, 2018). Since access to higher education programs is an important determinant of later life outcomes, such as income, attitudes, and political behavior (Lemann, 1999, p. 6), it is important that admission procedures consist of fair and valid instruments and procedures.

The widespread use of computers further allows for more varied forms of assessment, which makes admission testing even more difficult. Testing at a distance is now more common, although it does raise questions concerning the validity of the test results. Dishonest testing behavior (e.g., cheating) is more difficult to control in unproctored, online, tests. Furthermore, the security of test items is also potentially jeopardized, which may contribute to inflated test scores. Hence, it is crucial to ascertain that test takers who are assessed at a distance (i.e., unproctored) are not advantaged over test takers who are assessed in a proctored environment. In this study we investigate whether proctored and unproctored tests may lead to different test results, and to differences in prediction, which is of major importance in admission testing. If unproctored test-takers engage in cheating, we would expect that their academic performance is overpredicted, that is, they perform less well academically

than we would expect based on their admission test scores. We study differential prediction between unproctored and proctored test using real admission test data. Specifically, we compare scores across the two groups by means of the moderated multiple regression model proposed by Lautenschlager and Mendoza (1986), under both the frequentist and the Bayesian paradigm. Our goal is to investigate whether differential prediction of first year GPA exists between the unproctored and proctored groups of applicants. Finally, in our last study we use a Bayesian model that uses prior information from earlier years to investigate how we can quantify the probability of success in a future study based on admission test scores.

The overall aim of this research is to contribute to the current knowledge with respect to: (a) the extent to which candidates' scores administered in a proctored setting differ from those administered in an unproctored setting, and (b) the performance of Bayesian methods in the context of admission testing and prediction, and how they can supplement the information obtained using frequentist methods. In particular, we explore how Bayesian methods can be used to obtain information about future study success based on admission test scores, with a particular emphasis on deriving applicant-based prediction information.

Proctored versus Unproctored Testing

Both in high-stakes personnel-selection and educational admission procedures, tests and questionnaires are sometimes administered online in an unproctored way. For administering noncognitive- or character-based measures, several studies showed that the differences between proctored and unproctored administrations were minimal (e.g., Chuah, Drasgow, & Roberts, 2006). Administering cognitive measures online in an unproctored mode is less common, and if this is done, often a second shorter

version of the test is administered to selected candidates in a proctored setting (e.g., Makransky & Glas, 2011). Furthermore, research results with respect to the similarity of proctored and unproctored administered tests have been mixed. For example, Alessio et al. (2017) found large differences between unproctored and proctored scores, whereas Beck (2014) did not. The different results may be explained by the different control techniques designed to minimize cheating, such as administering items randomly, or preventing candidates from inspecting earlier administered items in the test administration. In large scale admission testing, there has not been much research comparing proctored and unproctored test scores. This research seems timely because universities are using both proctored and unproctored exams in admission testing.

Frequentist versus Bayesian Approach

In prediction research, there are several reasons to supplement the frequentist approach with the Bayesian approach (e.g., Kruschke, Aguinis, & Joo, 2012). First, the Bayesian approach offers many possibilities for hypothesis testing and parameter estimation (e.g., Gelman et al., 2014; Kruschke, Aguinis, & Joo, 2012). Based on the frequentist approach, one usually computes a p -value, which is the probability of observing the data at hand or more extreme, given that the model under consideration holds, that is, one considers conditional probabilities of the type $p(\text{data} \mid \text{model})$. However, in most cases, we are actually interested in the plausibility of certain hypotheses, given the observed data (after all, the data are just a vehicle for us to learn about the phenomenon that we are hypothesizing about). In other words, we are interested in the reversed conditional probability $p(\text{model} \mid \text{data})$. Questions of this type cannot be answered directly based on the frequentist approach, because in model

parameters are considered fixed under the frequentist paradigm and hence not subject to the laws of probability. This is unlike the Bayesian approach, which takes into account model uncertainty based on the (fixed) observed data. So, for example in the context of hypotheses testing, the Bayesian approach allows quantifying which of the competing hypotheses is more likely in light of the data, that is, we do consider $p(\text{model} \mid \text{data})$ explicitly. Concerning parameter estimation, the Bayesian paradigm also provides direct answers in terms of the ranges of the most probable values for a particular parameter. The frequentist confidence interval fails at this because its stochastic nature lies within the process utilized to compute it (under hypothetical repeated sampling) and not in the confidence interval itself. Therefore, a 95% confidence interval does not imply that the true parameter lies within the two numerical bounds with probability .95. Simply, if computed many times under similar sampling conditions (and assuming all assumptions hold), one expects that 95% of the intervals computed in this way will contain the unknown parameter of interest. Arguably, this is a property of little practical value in most instances: Researchers want to learn from the only data they observed instead of relying on an imaginary infinite sampling procedure to justify the range of numbers they found. In contrast, a Bayesian credible interval (BCI) does provides a direct answer. That is, a 95% BCI does imply that there is .95 probability that the unknown parameter lies between the two estimated bounds (based on the stipulated prior and model). So, BCIs can be interpreted as the most probable values of a parameter given the data (Kruschke & Liddell, 2018).

Second, the Bayesian approach does not capitalize on issues such as dependence on unobserved data, subjective stopping data collection rules (i.e., continuing data collection until a certain result is achieved), multiple testing, and lack

of support towards the null hypothesis (e.g., Dienes, 2016; Gelman et al., 2014; Rouder, 2014; Tendeiro & Kiers, 2018; Wagenmakers, 2007).

A third reason that in particular applies to differential prediction studies is that there are some shortcomings of the classical step-down regression analysis (Lautenschlager & Mendoza, 1986) to investigate differential prediction (e.g., Aguinis et al., 2010). Tests for slope differences tend to be underpowered, even in large samples, and tests for intercept differences tend to have inflated Type I errors (Aguinis et al., 2010). However, many authors have argued that slope differences are nonexistent in most cases or that a test was unbiased based on a lack of statistically significant slope or intercept differences (e.g., Hough et al., 2003). This is not allowed in frequentist statistics: Lack of significance does not equal evidence for the null model. There have been suggestions to overcome these problems (Aguinis et al., 2010), but most suggestions are difficult to implement, especially when slope differences are present (Berry, 2015). A Bayesian approach does not solve all these problems, but inconclusive (i.e., not statistically significant) results can be distinguished from evidence in favor of the null hypothesis of no differential prediction.

A final advantage of the Bayesian approach is that it allows the use of prior information. This is an interesting feature in predictive validity studies. For example, Anthony, Dalessandro, and Trierweiler (2016) discussed in a validation study of the LSAT: “A primary purpose of conducting validity studies for most schools is to obtain the best possible prediction weights so that they can be applied to the application credentials of the subsequent year’s applicant pool to aid in the decision process.” That is, data from past experience is used to make future predictions. When results from the predictive validity studies are used in this way, the most relevant

question to ask is: How well do the equations from previous first-year classes predict the performance of future first-year classes?” In the present study we will use the information from the previous years as prior information for future years (“today’s posterior is tomorrow’s prior”; Lindley, 1972).

Method

Data

Data were used that were described and analyzed in Niessen, Meijer, and Tendeiro (2016, 2018a, 2018b). However, in these studies data were not analyzed with the aim of investigating the similarity between proctored and unproctored conditions and the use of a frequentist and Bayesian approach for predicting student performance.

Three samples of students who applied and enrolled to a psychology undergraduate program at a Dutch university were used (2013, 2014, and 2015). All participants completed a curriculum-sampling test as part of the admission procedure¹. A curriculum-sampling test is designed to mimic (part of) an academic program (de Visser et al., 2016; Niessen et al., 2018a). For the test used in this study, applicants had to study two chapters of a book on *Introduction to Psychology*, and take an exam about the material. This approach yielded high predictive validity (uncorrected $r = .46$ for first year GPA, Niessen et al., 2018a). None of the applicants were rejected based on the admission procedure, because the final number of potential enrollees did not exceed the number of available places. However, this was not known at the time the admission tests were administered, so the applicants perceived the test

¹ The admission procedure also consisted of an English-reading comprehension test and a math test in 2013 and 2014, and of a math test and a test about material provided through a video-lecture in 2015.

as high-stakes. Some applicants did choose not to enroll voluntarily, however. Students followed the psychology program in Dutch or English, with similar content. Most students who chose to follow the program in English were international students. More information on the admission procedure and the academic program can be found in Niessen et al. (2016; 2018a).

International applicants living far away or Dutch students who were abroad during the administration of the admission tests were allowed to take the tests online (unproctored), while the other applicants had to come to the university to take the admission test (proctored).

Cohort 1. The first sample consisted of the 638 students who applied to the program and enrolled in 2013. Seventy percent were female and the mean age was $M = 20$ ($SD = 2.0$). The Dutch program was followed by 43% of the students and, in the entire sample, 53% had a non-Dutch nationality. Of all students, 10% ($n = 62$) took the test unproctored.

Cohort 2. The second sample consisted of the 635 students who applied to the program and enrolled in 2014. Sixty-six percent were female and the mean age was $M = 20$ ($SD = 1.7$). The Dutch program was followed by 42% of the students and, in the entire sample, 55% had a non-Dutch nationality. Of all students, 13% ($n = 83$) took the test unproctored.

Cohort 3. The third sample consisted of the 531 students who applied to the program and enrolled in 2015. Seventy percent were female and the mean age was $M = 20$ ($SD = 2.0$). The Dutch program was followed by 38% of the students and, in the entire sample, 57% had a non-Dutch nationality. Of all students, 11% ($n = 60$) took the test unproctored.

Measures

Admission test. The admission test was a curriculum-sampling test (denoted CSTESt throughout), which was designed to mimic the first course in the program: *Introduction to Psychology*. The applicants had to study two chapters of the book used in this course. On the test day they took an exam about the chapters, constructed by a course instructor. In each cohort the exams consisted of different multiple-choice items (40 items in 2013 and 2014, 39 items in 2015); the estimated reliability of the tests was $\alpha = .81$ in 2013, $\alpha = .82$ in 2014, and $\alpha = .76$ in 2015.

First year GPA. The first year GPA (denoted FYGPA throughout) is based on the grades on all courses the students took within the psychology program, given on a scale from 1 to 10, with a 6 or higher representing a pass. There are no elective courses in the first year, so all students took the same courses. However, some students chose not to participate in some courses. Most courses require literature study and are accompanied by weekly lectures, and end with a multiple-choice exam. When students fail the first attempt, they can take a resit exam. The first year GPA was based on the course grades after resits.

Analyses

The analyses were divided into three parts. First, we quantified the difference in test scores between students who took the test unproctored or proctored, using both frequentist and Bayesian methods; second, we investigated differential prediction for the unproctored and proctored groups, using both frequentist and Bayesian methods; and third, we investigated the predictive power of the CSTESt using Bayesian updating procedures. All analyses were conducted in R (R Core Team, 2018). Bayesian estimation was done by means of rstan (Stan Development Team, 2018), which is the R interface of the Stan platform for Bayesian modeling (Carpenter et al., 2017). Bayes factors for hypothesis testing were computed via the BayesFactor R

package (Morey & Rouder, 2018). The `ggplot2` package (Wickham, 2016) was used to produce all plots (in particular, function `stat_density()` was used to display smoothed kernel density estimates of posterior distributions of continuous parameters).

Score differences based on unproctored or proctored testing. The goal of these analyses was to fully characterize the differences in test scores between students who took the tests unproctored or proctored. First, we described the distributions of CSTEEST and FYGPA across both groups, for each of the three cohorts taking the admission tests. Differences between groups for each variable (per cohort) were quantified by raw group mean differences and by Cohen's d . By means of Bayesian statistics, posterior distributions for group mean differences and for Cohen's d were estimated, although we focus on inferences for Cohen's d for simplicity. We used the robust estimation procedure for two groups discussed by Kruschke (2015, Chapter 16); specific details concerning the model, priors, and sampling specifications can be found in Appendix A. Using this Bayesian approach enabled us to make probabilistic statements about the difference between the unproctored and the proctored groups. Specifically, we could answer questions of the type "What is the probability that d is small (i.e., between -0.2 and 0.2 by Cohen's standards)?" Furthermore, we compared the cutoff scores between both groups based on the curriculum-sampling test scores, at various selection ratios.

Differential prediction based on unproctored or proctored testing. We investigated whether the admission test was equally predictive of first year GPA for both the unproctored and the proctored groups. To do this, we used the step-down hierarchical multiple regression algorithm discussed by Lautenschlager and Mendoza (1986). This algorithm is based on comparing nested regression models by means of

the F test for the difference in R^2 between the models. There are four models of interest (notation: PROC = dichotomous variable (0 = unproctored, 1 = proctored)):

$$\text{FYGPA} = b_0 + b_1\text{CSTEST} \quad (1)$$

$$\text{FYGPA} = b_0 + b_1\text{CSTEST} + b_2\text{PROC} + b_3\text{CSTEST} \times \text{PROC} \quad (2)$$

$$\text{FYGPA} = b_0 + b_1\text{CSTEST} + b_3\text{CSTEST} \times \text{PROC} \quad (3)$$

$$\text{FYGPA} = b_0 + b_1\text{CSTEST} + b_2\text{PROC} \quad (4)$$

This algorithm consists of three main steps:

- Step 1: Compare Models 1 and 2. In case the F test is statistically significant:
Infer prediction bias and proceed to Step 2, otherwise stop.
- Step 2: Compare Models 4 and 2. In case the F test is statistically significant:
Infer regression slope differences and proceed to Step 3a, otherwise do not
infer regression slope differences and proceed to Step 3b.
- Step 3a: Compare Models 3 and 2. In case the F test is statistically significant:
Infer both intercept and slope differences, otherwise infer only slope
differences.
or
Step 3b: Compare Models 4 and 1. In case the F test is statistically significant:
Infer only intercept differences, otherwise infer no intercept nor slope
differences (a very unlikely outcome according to Lautenschlager & Mendoza,
1986).

Step 1 offers an omnibus test of differential prediction. If Model 2 does not explain a significantly larger proportion of variance of FYGPA than Model 1, then the algorithm stops and the next steps are not taken. Observe that failing to reject the null hypothesis in this context implies that there was not have enough evidence in the data favoring Model 2 over Model 1. This does *not* imply that Model 1 is more likely,

however, because no conclusions can be drawn from a rejection failure in null hypothesis significance testing (NHST). In other words, we cannot ascertain that no differential prediction exists by means of not rejecting Model 1 over Model 2. Similar limitations exist if one fails to find a significant result for Step 2 (test for difference in slopes) and for Step 3 (test for difference in intercepts).

Bayes factors do allow quantifying the evidence in the data *for* either hypothesis that we compare, and in this sense Bayes factors are symmetric, unlike NHST. Therefore, we can compare the likelihood of the observed data under either model under consideration. We applied the same algorithm as Lautenschlager and Mendoza (1986, discussed above) but replaced the F tests by Bayes factors to perform model comparison. In order to decide whether the evidence for the alternative was strong enough, it is necessary to set up a threshold for the Bayes factor. We followed the guidelines from Jeffreys (1961) and used a value of 10, indicating “strong” evidence in favor of the alternative, as the minimum level of evidence required to reject the null model. The default Bayes factor in the function `regressionBF()` from the `BayesFactor` package was used for these analyses.

Predictive power. In situations where admission testing is routinely used, information from previous test administrations may be used to improve the predictive power of the model used to fit the data. In the context of our data, this implies using the data from previous administrations (i.e., from the first two cohorts) to improve the regression model that relates curriculum-sampling test scores with first year GPA. This model can then be used for applicants in the third cohort to predict their first year GPA. The Bayesian framework naturally allows for this updating mechanism, because posterior distributions for model parameters that are estimated after one test administration can become the prior distributions in the coming year’s model fit

analysis. In other words, the estimation of the model parameters is updated by the data from previous test administrations. Therefore, we estimated the Bayesian regression model that predicts first year GPA from the curriculum-sampling test scores based on the data from the first two cohorts. This model was then used to predict the first year GPA for all applicants in the third cohort, using their admission test scores as the model input. Because a Bayesian regression model was used, we were able to draw probabilistic information for each applicant based on their observed test score, which is a strong advantage of using these types of models. We used the robust single regression estimation procedure of Kruschke (2015, Chapter 17); see Appendix B for more details concerning the model, priors, and sampling specifications.

Results

Unproctored versus proctored testing

Table 1 summarizes the mean and standard deviation of CSTEEST scores and FYGPA across the three cohorts and the two groups (proctored, unproctored). The differences between the test scores of the proctored and unproctored groups were very small for both variables. Figure 1 displays the densities of the two variables across the three cohorts; from this figure we conclude that there is a large overlap between the distributions across the two groups.

Because minimum scores are sometimes required in admission testing, it is important that these scores are similar for different subgroups of applicants. Therefore, we compared the quantiles of the distribution of CSTEEST for both the proctored and the unproctored groups. Figure 2 shows the quantiles for the first cohort (the plots were similar for the other two cohorts). Note that the quantiles were very similar for the proctored and unproctored group particularly for the right-end side of

the scale, which is the most relevant one for selection purposes. Thus, also according to this criterion, no relevant differences between groups were observed.

We can use inference to better gauge group differences. These differences can be further quantified by means of Cohen's d , for each variable in a cohort. Table 1 presents all d values together with their corresponding 95% confidence intervals. We also estimated posterior distributions for d and compared these Bayesian results (Table 2) to those from the frequentist analysis (Table 1). We first assumed normality of the scores within each group. However, inspection of the posterior predictive distributions revealed misfit in particular for variable FYGPA. As can be seen in Figure 1, the distributions for both CSTEEST and FYGPA are skewed to the left. The skew is more pronounced for FYGPA in cohorts 2 and 3 (skew values between -1.03 and -1.29). This deviation from normality was poorly captured by a Bayesian independent samples model based on the normal likelihood. As an example, in Figure 3 the posterior predictive distributions for FYGPA with the observed data in cohort 3 are provided. It can be seen that data sampled from the posterior predictive distribution fit the observed data better when the t likelihood is used instead of the normal likelihood. Therefore, we report the posterior summaries for Cohen's d based on the robust Student- t model (see Table 2). We would like to highlight this idea that the Bayesian paradigm makes the selection of a proper likelihood model explicit (since it needs to be input into the statistical model) and adjustments as done here are relatively straightforward. In this particular situation of two groups comparison, the implicit assumption of normality of scores within each group is paramount under the frequentist approach. Changing the statistical model and adjusting the model estimation procedure is not easy to implement and practitioners are simply often inclined to assume that normality holds "reasonably well."

In both the frequentist and the Bayesian approaches, the effect sizes for the differences in test scores between students who took the tests proctored or unproctored were small, with slightly higher test scores for students who took the test proctored. However and as previously explained, the reported confidence intervals and credible intervals should be interpreted differently albeit they look similar. A 95% confidence interval is a range of values which is computed by a procedure that, when all underlying assumptions are met, lead to an interval that covers the population d value in 95% of the cases across repeated sampling. However, a specific confidence interval is not probabilistic with respect to the parameter. Therefore, we cannot conclude that the population Cohen's d in the first cohort for CSTEEST lies within the interval $(-.12, .40)$ (see Table 1) with probability .95. In contrast, Bayesian credible intervals *are* stochastic, thus we may conclude that there is 95% probability that the population Cohen's d in cohort 1 for CSTEEST is between $-.18$ and $.44$ (see Table 2), for the priors and model considered. Credible intervals, therefore, allow a more direct interpretation than confidence intervals. Furthermore, by means of Bayesian statistics we have access to full posterior distributions of d . As an example, Figure 4 displays the posterior distribution of Cohen's d in cohort 2 for CSTEEST. The credible interval $(-.17, .30)$ is one outcome that can be derived from this distribution. For instance, we can also derive that the probability of the population Cohen's d being small, medium, or large (i.e., with absolute value $< .2$, between $.2$ and $.5$, or larger than $.5$) is .8677, .1322, and .0001, respectively. This thus tells us that the difference between the proctored and the unproctored groups is negligible.

Differential prediction based on unproctored versus proctored testing

One way to compare the validity of the tests scores of the proctored and unproctored groups is by comparing the corresponding regressions of FYGPA on

CSTEST. Ideally, the same prediction model applies to both groups. In Figure 5 we plotted the regression of FYGPA on CSTEST in each cohort and for each group. Note that there may be slope differences between the groups: The slope is consistently larger for the proctored group. As explained above, we used the step-down hierarchical multiple regression algorithm of Lautenschlager and Mendoza (1986) to investigate how serious these differences are. As shown in Table 3, the algorithm stopped after Step 1 was taken. Based on the frequentist procedure (first three columns), we conclude that there is little evidence in the data supporting rejecting Model 1, and therefore we have no evidence suggesting that there is differential prediction. We do not conclude, however, that there is no differential prediction: A failure to reject the null model cannot be interpreted as evidence *for* the null model. Thus, all we can say is that there is not enough evidence to reject the differential prediction hypothesis.

In this respect, the Bayesian analysis is more informative because it takes into account the likelihood of the data under the null model and the alternative model. NHST only considers the likelihood of the data under the null model. The last three columns in Table 3 show that the data are consistently more likely under Model 1 than under Model 2. For example, the Bayes factor of 52.80 in cohort 1 indicates that the observed data are over 50 times more likely under Model 1 than under Model 2. If the prior model odds equal 1, that is, if both models were considered equally likely beforehand, then the Bayes factor equals the posterior odds and we can derive that the posterior model probabilities are equal to $p(\text{Model 1}|D) = .981$ and $p(\text{Model 2}|D) = .019$. That is, given the data, Model 1 is much more likely than Model 2. This provides direct information concerning the lack of differential prediction between

both groups under consideration, that is, we have evidence *for* the null model, unlike under the frequentist paradigm.

Thus we conclude that there seems to be no differential prediction and we combined the scores of the proctored and unproctored students in our predictive power analyses discussed below.

Predictive power and the probability of success in a future study program

We fit a Bayesian robust regression model to the data from cohorts 1 and 2 ($n = 1,273$) to estimate a model to predict FYGPA from CSTEEST in cohort 3. Thus, we estimated posterior distributions for the intercept, slope, and standard deviation of the residuals. One benefit of using the Bayesian framework is that we can update the model estimates when new data arrives. This can be visualized by comparing the posterior distributions based on cohort 1 data only ($n = 638$) with the posterior distributions based on cohorts 1 and 2 data. Figure 6 shows that the posterior distributions are narrower for the combined data of cohorts 1 and 2 than using only cohort 1 data, which implies greater estimation precision.

Inspecting the posterior predictive distribution results from the regression model estimated from the data of the first two cohorts, we observe that the model provides a relative good fit. Figure 7 provides the posterior predictive distribution of FYGPA with the observed scores. The left panel of Figure 7 is based on the same data used to fit the model, so the (good) fit may be regarded as artificially good. However, in our current setting we happen to have available the FYGPA from cohort 3 (this will not always be the case in practice, which is why prediction is performed). We were therefore able to compare the posterior predictive distribution estimated from the data from cohorts 1 and 2 with the out-of-sample data from cohort 3 (right panel of Figure 7). Note that there is a reasonable fit, but that the left tail of the distribution of

observed FYGPA scores is not particularly well captured by the posterior predictive distribution. Figure 8 displays the posterior mean (dashed line) and 95% prediction bands (solid lines) computed based on the regression model estimated from cohorts 1 and 2 and on the observed CSTEEST scores in cohort 3. Figure 8 also displays 200 regression lines based on intercept and slope parameters randomly drawn from the posterior joint distribution. These regression lines are relatively close to each other because the standard deviations of the posterior distribution of the intercept and slope are small due to the large sample size of the calibration sample (see Figure 6). The prediction band is rather wide because the estimated standard deviation of the residuals is relatively large (see Figure 6; the 95% credible interval of σ is (.85, .98)). However, the usefulness of administering the CSTEEST (and admission testing, in general) can be illustrated by inspecting the posterior predictive distributions for different CSTEEST scores.

By means of Figure 8 we can derive the posterior prediction intervals conditional on an applicant's number-correct CSTEEST score. The posterior predictive distribution is shown in Figure 9 for four different number-correct scores on the CSTEEST (10, 20, 25, and 35). We can now compute the probability of any range of FYGPA of interest. As an example, in Figure 9 we provide shaded areas for the predicted probability of FYGPA at or above the passing threshold (5.5 in the Dutch grading system). A student with a number-correct score of only 10 (out of 40) at the CSTEEST admission test has a predicted .15 probability of attaining a FYGPA at or above 5.5. This probability increases when the CSTEEST scores increases. The probabilities are equal to .52, .73, and .94 for CSTEEST scores equal to 20, 25, and 35, respectively. This relation between admission test scores and later performance indicators is extremely useful both for applicants and admissions committee members.

It provides a direct answer to the question what the probabilities are that a candidate will succeed in a study. Moreover, these probabilities can also assist in determining a sensible cutoff score for the admission test. If, for example, we want to optimize selection and admit only those candidates with a predicted probability of success, that is $FYGPA \geq 5.5$, of at least .70, then the advised minimum entrance score equals $CSTEST = 25$ (see Figure 10).

Discussion

The first question we addressed in this study was whether proctored and unproctored data from an admission test resulted in similar scores and similar predictions of first year GPA. The answer to this question was affirmative. This does not imply that candidates did not cheat in the unproctored condition. Theoretically it is possible that the candidates in the unproctored condition had lower ability than the candidates in the proctored condition and that they raised their scores through cheating to the level of the candidates in the proctored condition. We think, however, that this is unlikely because in general “unproctored candidates” did not perform less well during the study than “proctored ones”, given the same test scores.

The second question we addressed was how Bayesian statistics can be of help in the analysis of admission test data. We illustrated that in this context Bayesian modeling can be used to incorporate prior information in the decision process and we discussed that posterior predictive distributions can be of help to interpret admission test scores in terms of the probability of success in the later study.

The models we discussed and figures such as Figures 8 and 9 can be used to provide meaningful individual feedback to test users such as applicants and admission officers. These figures also show the uncertainty in the model; the prediction bands are quite wide, even though they are based on predictors with predictive validity that

is considered high (uncorrected $r = .46$, Niessen et al., 2018a) in the context of predicting human performance. A first response to this amount of uncertainty may be that the predictive accuracy of the models should be improved, for example by adding additional predictors, in order for the models to be useful in practice. However, in spite of decades of research efforts, predictors or combinations of predictors that yield substantially higher predictive validity² are not currently available. In high-stakes testing, only the combination of prior GPA and admission test scores yield somewhat higher validities than admission tests alone. The highest predictive validities found in operational admission testing based on this combination are about $r = .60$, after correcting for range restriction and criterion unreliability (e.g., Anthony et al., 2016; Kuncel & Hezlett, 2007; Zwick, 2017). So, improving the predictive accuracy of academic performance in admission procedures is very challenging, and one plausible reason for the imperfect results may be that the remaining variance in future performance may be hard, or even impossible to predict (Dawes, 1979; Zwick, 2017). However, the fact that predictors of academic performance do not possess near-perfect predictive accuracy does not mean that they cannot be not useful in practice. Utility models show that, depending on factors such as the base rate (the percentage of suitable candidates in the applicant pool) and the selection ratio (the percentage of students that will be admitted), such predictors can significantly increase the performance level of admitted students at the aggregate level (Naylor & Shine, 1965; Taylor & Russell, 1939).

In admission testing, and in selection research in general, several authors have called upon providing methods that can help to better communicate research findings (e.g., Bridgeman, Burton, & Cline, 2009). In this respect the use of the posterior

² And that are legal and considered acceptable, see Zwick (2017, p. 86).

predictive distribution is an interesting tool. This distribution can be used to link admission test scores to, for example FYGPA, as we illustrated above. The probability of later study success given a particular score admission test score is for example very useful in matching procedures. In matching procedures in higher education, candidates are tested for diagnostic purposes to provide information about their match between the study and/or university. Posterior predictive distribution can then be used to show their probability of success given their obtained scores on the “matching” test. This is much more informative than communicating that a test has a correlation of .4 with a criterion measure, for example. In particular in a time where there seems to be a deep suspicion about standardized testing, communicating test results to different stakeholders such as teachers, parents, and future students is very important. The Bayesian procedure discussed in this study can be of help to improve the communication with stakeholders.

References

- Aguinis, H., Culpepper, S. A., & Pierce, C. A. (2010). Revival of test bias research in preemployment testing. *Journal of Applied Psychology, 95*, 648-680.
doi:10.1037/a0018714
- Allessio, H. M., Malay, N., Maurer, K., Bailer, A. J., & Rubin, B. (2017). Examining the effect of proctoring on online testing. *Online learning, 21*, 146-161.
- Anthony, L. C., Dalessandro, S. P., & Trierweiler, T. J. (2016). *Predictive Validity of the LSAT: A National Summary of the 2013 and 2014 LSAT Correlation Studies* (LSAT Technical Report 16-01).
- Beck, V. (2014). Testing a model to predict online cheating: Much ado about nothing. *Active learning in higher education, 15*, 65-75.

- Berry, C. M. (2015). Differential validity and differential prediction of cognitive ability tests: Understanding test bias in the employment context. *Annual Review of Organizational Psychology and Organizational Behavior*, 2, 435–463. doi: 10.1146/annurev-orgpsych-032414-111256
- Bridgeman, B., Burton, N. , & Cline, F. (2009). A note on presenting what predictive validity numbers mean. *Applied Measurement in Education*, 22, 109-119. doi:10.1080/08957340902754577
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76. doi: 10.18637/jss.v076.i01
- Chuah, S. C., Drasgow, F., & Roberts, B. W. (2006). Personality assessment: Does the medium matter? No. *Journal of Research in Personality*, 40, 359-376.
- Clinedinst, M., & Patel, P. (2018). *State of college admission 2018*. Arlington, Virginia: National Association for College Admission Counseling. Retrieved from: https://www.nacacnet.org/globalassets/documents/publications/research/2018_soca/soca18.pdf
- Dawes, R.M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, 34, 571-582. doi:10.1037/0003-066X.34.7.571
- De Visser, M., Fluit, C., Fransen, J., Latijnhouwers, M., Cohen-Schotanus, J., & Laan, R. (2016). The effect of curriculum sample selection for medical school. *Advances in health Science Education*, 22, 43-56. Doi:10.1007/s10459-016-9681-x

- Dienes, Z. (2016). How Bayes factors change scientific practice. *Journal of Mathematical Psychology*, 72, 78–89.
- Gabri, J., & Mahr, T. (2018). bayesplot: Plotting for Bayesian Models. R package version 1.6.0. <https://CRAN.R-project.org/package=bayesplot>
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian Data Analysis* (3rd ed.). Boca Raton, FL: CRC Press.
- Hough, L. M., Oswald, F. L., & Ployhart, R. E. (2001). Determinants, detection and amelioration of adverse impact in personnel selection procedures: Issues, evidence and lessons learned. *International Journal of Selection and Assessment*, 9, 152-194. doi:10.1111/1468-2389.00171
- Jeffreys, H. (1961). *Theory of probability* (3rd ed.). Oxford: Oxford University Press.
- Kruschke, J. (2015). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan*. San Diego, CA: Elsevier, Inc.
- Kruschke, J. K., & Liddell, T. M. (2018). Bayesian data analysis for newcomers. *Psychonomic Bulletin & Review*, 25, 155–177.
- Kruschke, J. K., Aguinis, H., & Joo, H. (2012). The time has come: Bayesian methods for data analysis in the organizational sciences. *Organizational Research Methods*, 15, 722-752. doi: 10.1177/1094428112457829.
- Kuncel, N. R., & Hezlett, S. A. (2007). Standardized tests predict graduate students' success. *Science*, 315, 1080-1081. doi:10.1126/science.1136618
- Lautenschlager, G. J., & Mendoza, J. L. (1986). A step-down hierarchical multiple regression analysis for examining hypotheses about test bias in prediction. *Applied Psychological Measurement*, 10, 133-139.
- Lemann, N. (1999). *The big test: The secret history of the American meritocracy*. New York: Farrar, Straus & Giroux.

- Lindley, D. V. (1972). *Bayesian statistics, a review*. Philadelphia, PA: SIAM.
- Makransky, G., & Glas, C. A. W. (2011). Unproctored internet test verification: Using adaptive conformation testing. *Organizational Research Methods, 14*, 608-630.
- Naylor, J. C., & Shine, L. C. (1965). A table for determining the increase in mean criterion score obtained by using a selection device. *Journal of Industrial Psychology, 3*, 33-42.
- Niessen, A. S. M., Meijer, R. R., & Tendeiro, J. N. (2016). Predicting performance in higher education using proximal predictors. *PLoS ONE, 11*(4), e0153663. doi:10.1371/journal.pone.0153663
- Niessen, A. S. M., Meijer, R. R., Tendeiro, J. N. (2018a) Admission testing for higher education: A multi-cohort study on the validity of high-fidelity curriculum-sampling tests. *PLoS ONE 13*(6): e0198746. doi:10.1371/journal.pone.0198746
- Niessen, A. S. M., Meijer, R. R., & Tendeiro, J. N. (2018b). *Gender-Based Differential Prediction by Curriculum Samples for College Admissions*. Paper submitted for publication.
- Morey, R. D. & Rouder, J. N. (2018). BayesFactor: Computation of Bayes Factors for Common Designs. R package version 0.9.12-4.2. <https://CRAN.R-project.org/package=BayesFactor>
- 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/>.
- Rouder, J. N. (2014). Optional stopping: No problem for Bayesians. *Psychonomic Bulletin & Review, 21*, 301–308.

- Stan Development Team. 2018. Stan Modeling Language Users Guide and Reference Manual, Version 2.18.0. <http://mc-stan.org>
- Stemler, S. E. (2012). What should university admissions tests predict? *Educational Psychologist*, 47, 5-17. doi:10.1080/00461520.2011.611444
- Sternberg, R. J. (2010). *College admissions for the 21st century*. Cambridge, MA: Harvard University Press.
- Taylor, H. C., & Russell, J. T. (1939). The relationship of validity coefficients to the practical effectiveness of tests in selection: Discussion and tables. *Journal of Applied Psychology*, 23, 565-578. doi:10.1037/h0057079
- Tendeiro, J., & Kiers, H. (2018, November 25). A review of issues about NHBT. Retrieved from <https://osf.io/jmwk6>
- Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14, 779–804.
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. New York, NY: Springer-Verlag.
- Zwick, R. (2017). *Who gets in?: Strategies for fair and effective college admissions*. Cambridge, MA: Harvard University Press.

Tables

Table 1

Sample size, mean, and standard deviation of CSTEEST and FYGPA, across the three cohorts (years 1, 2, and 3) and the two groups (proctored, unproctored), and Cohen's d with 95% confidence intervals for group differences.

		Year 1 (n = 638)			Year 2 (n = 635)			Year 3 (n = 531)		
		N	Mean	SD	N	Mean	SD	N	Mean	SD
CSTEEST	Proctored	576	29.80	5.07	552	29.94	5.45	471	29.25	4.61
	Unproctored	62	29.08	5.92	83	29.63	5.44	60	28.75	5.58
		$d = .14$ (-.12, .40)			$d = .06$ (-.17, .29)			$d = .11$ (-.16, .38)		
FYGPA	Proctored	576	6.63	1.29	552	6.44	1.34	471	6.63	1.23
	Unproctored	62	6.55	1.40	83	6.46	1.44	60	6.68	1.40
		$d = .06$ (-.20, .32)			$d = -.01$ (-.24, .22)			$d = -.04$ (-.31, .23)		

Table 2

Posterior mean and 95% credible interval for Cohen's d across the three cohorts (years 1, 2, and 3).

		Year 1		Year 2		Year 3	
		d	95% BCI	d	95% BCI	d	95% BCI
CSTEEST		.13	(-.18, .44)	.06	(-.17, .30)	.09	(-.23, .43)
FYGPA		.12	(-.24, .51)	-.03	(-.30, .25)	-.07	(-.45, .32)

Note. BCI = Bayesian credible interval.

Table 3

Results from Step 1 of the differential prediction algorithm from Lautenschlager and Mendoza (1986). See text for details.

	Frequentist			Bayesian		
	F (df1, df2)	p	ΔR^2	BF ₁₂	$p(\text{Model 1} D)$	$p(\text{Model 2} D)$
Year 1	1.005 (2, 634)	.37	.002	52.80	.981	.019
Year 2	2.044 (2, 631)	.13	.005	16.12	.942	.058
Year 3	0.743 (2, 527)	.48	.002	46.50	.979	.021

Figure 1

Figure 1. Density of CSTE (left column) and FYGPA (right column) for each group (proctored, unproctored). The top, middle, and bottom panels concern the three cohorts (years 1, 2, and 3), respectively. The vertical lines represent group mean scores. All group means are very close to each other.

Figure 2

Figure 2. Quantiles of CSTEtest in the first cohort for the proctored group against the unproctored group. The quantiles are closely aligned along the identity line in the upper range of the scale, which is the part of the scale more relevant for student selection.

Figure 3

Figure 3. Two-hundred draws from the posterior predictive distribution of FYGPA (gray lines), with the density of the observed FYGPA scores superimposed (thick line), in the third cohort. The top panel is based on the normal likelihood, the bottom panel is based on the Student-t likelihood (more robust to the violations to normality of the observed data). The model based on the Student-t likelihood fits better.

Figure 4

Figure 4. Posterior distribution of Cohen's d in cohort 2 for CSTEST. The shaded area corresponds to the central 95% density under the curve, which determines the bounds of the 95% Bayesian credible interval (BCI).

Figure 5

Figure 5. Regression of FYGPA on CSTEST across groups (unproctored, proctored), per year. The solid line corresponds to the proctored group, the dashed line corresponds to the unproctored group. The scattered points were jittered in order to improve visibility.

Figure 6

Figure 6. Posterior distributions of the intercept (top-left), slope (top-right), and standard deviation of the residuals (bottom) for the simple linear regression model predicting FYGPA from CSTEEST. The solid line is based on data from the first two cohorts, the dashed line is based on data from cohort 1 only. The posterior distributions get narrower as data accumulates. The displayed intervals are 95% Bayesian credible intervals and correspond to the central 95% area under the corresponding curve (shaded gray).

Figure 7

Figure 7. Two-hundred draws from the posterior predictive distribution of FYGPA (gray lines), with the density of the observed FYGPA scores superimposed (thick line). The left panel is based on the same data used to fit the model (i.e., from cohorts 1 and 2). The data from the right panel are from cohort 3, hence they were not used to fit the model.

Figure 8

Figure 8. Data from cohort 3, with a 95% posterior prediction band (solid lines) and the posterior mean predicted score (diagonal dashed line). The gray lines around the posterior mean predicted line are two-hundred regression lines with parameters randomly drawn from the corresponding posterior distributions. The vertical line at $\text{CSTEST} = 25$ includes the prediction interval at this particular value (from $\text{FYGPA}=3.7$ through $\text{FYGPA}=8.4$).

Figure 9

Figure 9. Posterior predictive distributions of FYGPA at four CSTEEST number-correct scores (10, 20, 25, and 35). The shaded areas are the posterior probabilities of FYGPA being at or above 5.5 (i.e., $p(\text{FYGPA} \geq 5.5 | \text{data})$), which is the minimum passing grade in the Dutch educational system. The probabilities are shown in each panel's title. The displayed intervals are 95% Bayesian credible intervals.

Figure 10

Figure 10. Posterior probability of FYGPA being at least 5.5 as a function of CTEST. The minimum CTEST score required such that $p(\text{FYGPA} \geq 5.5 | D)$ is equal to .70 is CTEST = 25.

Appendix A

Bayesian robust estimation procedure for two groups (Kruschke, 2015, Chapter 16)

Suppose we have data from two groups ($j = 1, 2$), where each group is normally distributed with mean μ_j and standard deviation σ_j . The goal is to quantify the difference between the population means, or to quantify the true effect size operationalized by Cohen's d .

Kruschke (2015) observed that the normal sampling model is not suitable when outliers or heavier tails exist. In such cases, the advice is to use the t distribution, which is a sampling model with heavier tails than the normal. Our own analyses corroborated Kruschke's theory (the posterior predictive distributions based on the normal model did not fit the data as well as the t student data model, as illustrated in Figure 3). Therefore, we used the t distribution as the data model (Kruschke, 2015, p. 468):

$$y_{i|j} \sim t(\nu, \mu_j, \sigma_j), \quad (A1)$$

where $y_{i|j}$ is the i -th score in group j ($i = 1, \dots, n_j$), ν is the number of degrees of freedom (Kruschke calls this the 'normality' parameter), μ_j is the location parameter, and σ_j is the scale parameter (it is *not* the standard deviation). The model has five parameters: Two location parameters, two scale parameters, and the normality parameter. Following Kruschke, we used the following prior distributions for the parameters:

$$\begin{aligned} \nu - 1 &\sim \text{Exponential}(1/29) \\ \mu_j &\sim \text{Normal}(\text{Mean}_y, SD_y \times 1000) \\ \sigma_j &\sim \text{Uniform}(SD_y/1000, SD_y \times 1000). \end{aligned} \quad (A2)$$

Here, Mean_y and SD_y denote the sample mean and standard deviation of the y values across both groups.

The model was written in Stan; the code is freely available at <https://osf.io/gec9m/>. Concerning the MCMC setup, the first 1000 iterations were discarded (burn-in). Four chains of 2,500 samples each were sampled (therefore, $k = 10,000$ samples per parameter were drawn from the joint posterior distribution). Thinning was used mostly to reduce issues with autocorrelations (only 1 in every 5 samples was saved). We performed extensive model convergence checks using the bayesplot R package (Gabry & Mahr, 2018), including looking at the following outputs: Parallel coordinates plot, trace plots, NUTS energy, Rhat, ESS, and autocorrelation. No convergence problems were identified.

The posterior distribution for the difference of means and for Cohen's d was computed by means of the 'generated quantities' block in the Stan model. After each of the $k = 10,000$ MCMC sampling steps, the following quantities were computed:

$$s_p = \sqrt{\frac{\text{diff}_{\text{means}}^2}{(n_1 - 1)\sigma_{1k}^2 + (n_2 - 1)\sigma_{2k}^2}} \quad (A3)$$

$$\text{Cohen}_d = \frac{\text{diff}_{\text{means}}}{s_p}.$$

where μ_{1k} , μ_{2k} , σ_{1k} , and σ_{2k} are the k -th MCMC step's estimates of μ_1 , μ_2 , σ_1 , and σ_2 , respectively.

Appendix B

Bayesian robust simple linear regression (Kruschke, 2015, Chapter 17)

Similarly to Appendix A, we opted to use the t distribution to model the error term in the regression model, as it provided more robust inferences in case outliers or heavy tails are present in the data. The data model is given as follows (Kruschke, 2015, p. 480):

$$y_i \sim t(\nu, \beta_0 + \beta_1 x_i, \sigma). \quad (B1)$$

This model has four parameters: The normality, intercept, slope, and scale parameters.

The following priors were used:

$$\begin{aligned} \nu - 1 &\sim \text{Exponential}\left(\frac{1}{29}\right) \\ \beta_0 &\sim \text{Normal}(0, SD_{\beta_0}) \\ \beta_1 &\sim \text{Normal}\left(0, \frac{SD_y}{SD_x} \times 10\right) \\ \sigma &\sim \text{Uniform}\left(\frac{SD_y}{1000}, SD_y \times 1000\right). \end{aligned} \quad (B2)$$

These priors specifications are the same as Kruschke's except for the value SD_{β_0} .

Kruschke used the value $SD_{\beta_0} = 10 \times \text{abs}\left(\text{Mean}_y \frac{SD_y}{SD_x}\right)$ because he reasoned that

$\text{abs}\left(\text{Mean}_y \frac{SD_y}{SD_x}\right)$ is the largest value that the intercept can attain for perfectly

correlated data. However, the general form of the intercept is given by $\beta_0 = \text{Mean}_y -$

$\beta_1 \text{Mean}_x$, with $\beta_1 = \text{cor}(x, y) \frac{SD_y}{SD_x}$. If x and y correlate perfectly (i.e., $\text{cor}(x, y) = \pm 1$)

then $\beta_0 = \text{Mean}_y \mp \text{Mean}_x \frac{SD_y}{SD_x}$. We took this into account and decided the rephrase

SD_{β_0} as follows:

$$SD_{\beta_0} = 10 \times \max\left(\text{Mean}_y - \text{Mean}_x \frac{SD_y}{SD_x}, \text{Mean}_y + \text{Mean}_x \frac{SD_y}{SD_x}\right). \quad (B3)$$

The model was written in Stan; see model is freely available at

<https://osf.io/gec9m/>. The same setup for the MCMC algorithm that was explained in

Appendix A also applies here (thus: Four chains, burn in the first 1000 iterations, 2,500 iterations per chain, thinning by saving one of every five values sampled). The same convergence checks identified in Appendix A were performed; no convergence problems were identified.