Construct Validity of Multiple Mini Interviews –**Investigating the Role of Stations, Skills, and Raters using Bayesian G-theory**

Simon M. Breil¹, Boris Forthmann¹, Anike Hertel-Waszak¹, Helmut Ahrens¹, Britta Brouwer¹,
Eva Schönefeld¹, Bernhard Marschall¹, & Mitja D. Back¹

¹University of Münster, Germany

in press, Medical Teacher

This is an unedited manuscript that has been accepted for publication. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form.

We embrace the values of openness and transparency in science

(www.researchtransparency.org), we therefore published all raw data necessary to reproduce the reported results and provide scripts for all data analyses reported in this manuscript osf.io/d8qmu

This work was supported by the Bundesministerium für Bildung und Forschung (Federal Ministry of Education and Research, Germany; project number: 01GK1801A)

Correspondence concerning this article should be addressed to Simon Breil

(simon.breil@wwu.de, University of Münster, Institute for Psychology, Psychological Assessment and Personality Psychology, Fliednerstr. 21, 48149 Münster, Germany).

Notes on Contributors

Simon M. Breil, M. Sc. is a researcher at the department of psychology, psychological assessment and personality psychology at the University of Münster.

Dr. Boris Forthmann is a researcher at the department of psychology in education, assessment and evaluation in schools at the University of Münster.

Dr. Anike Hertel-Waszak is a researcher at the Institute for education and academic affairs at the medical department of the University of Münster.

Dr. Helmut Ahrens is a researcher at the Institute for education and academic affairs at the medical department of the University of Münster.

Dr. Britta Brouwer is a researcher at the Institute for education and academic affairs at the medical department of the University of Münster.

Dr. Eva Schönefeld is a researcher at the Institute for education and academic affairs at the medical department of the University of Münster.

Dr. Bernhard Marschall is a professor at the Institute for education and academic affairs at the medical department of the University of Münster.

Dr. Mitja Back is professor for psychological assessment and personality psychology at the psychological department of the University of Münster

Declaration of Interest

The authors report no declarations of interest.

Construct Validity of Multiple Mini-Interviews –**Investigating the Role of Stations, Skills, and Raters using Bayesian G-theory**

Abstract

One popular procedure in the medical student selection process are multiple mini-interviews (MMIs), which are designed to assess social skills (e.g., empathy) by means of brief interview and role-play stations. However, it remains unclear whether MMIs reliably measure desired social skills or rather general performance differences that do not depend on specific social skills. Here, we provide a detailed investigation into the construct validity of MMIs, including the identification and quantification of performance facets (social skill-specific performance, station-specific performance, general performance) and their relations with other selection measures. We used data from three MMI samples ($N = 376$ applicants, 144 raters) that included six interview and role-play stations and multiple assessed social skills. Bayesian generalizability analyses show that, the largest amount of reliable MMI variance was accounted for by station-specific and general performance differences between applicants. Furthermore, there were low or no correlations with other selection measures. Our findings suggest that MMI ratings are less social skill-specific than originally conceptualized and are due more to general performance differences (across and within-stations). Future research should focus on the development of skill-specific MMI stations and on behavioral analyses on the extents to which performance differences are based on desirable skills versus undesired aspects.

Keywords: Multiple mini-interviews, medical student selection, construct validity, dimensionality, generalizability theory

Practice Points

- When analyzing the contribution of different variance sources in MMI performance ratings, it is important to consider all relevant facets (i.e., applicants, raters, dimensions, stations, samples) for unbiased estimations.
- The largest amount of reliable MMI variance seems to stem from station-specific performance differences between applicants.
- There is little evidence for the reliable assessment of individual differences in specific social skills.
- For the future development of MMIs it might be beneficial to focus on one social skill per station that is assessed at multiple stations.

Introduction

Selecting students into medical school is a high-stakes exercise for all parties involved (e.g., Hughes 2002; Patterson et al. 2012). Career and self-fulfillment for applicants largely depend on whether or not they are accepted, whereas as a society, we hope the best-suited individuals will become our future medical doctors. Thus, there has been much discussion about how optimal selection decisions can be achieved (e.g., Albanese et al. 2003; Patterson et al. 2016). Besides using the generally accepted grade point average (GPA) as a selection aid, researchers and practitioners have been calling for the inclusion of social skills (also called “non-cognitive” abilities) as variables in the selection process (Albanese et al. 2003; Stern et al. 2003). In the multiple mini-interview (MMI), one commonly used procedure to measure social skills in medical education, applicants face several short exercises (i.e., stations; e.g., interviews, role-plays; Eva and Rosenfeld et al. 2004; Ziv et al. 2008; Knorr and Hissbach 2014). This procedure is very similar to assessment centers (ACs) used for personnel selection, in which applicants are also faced with several different exercises that often simulate work situations. After completing each station, applicants are typically evaluated on predefined social skill dimensions, for example communication skills, empathy, ethical reasoning, or the ability to handle stress (e.g., Eva and Rosenfeld et al. 2004; Ziv et al. 2008). Furthermore, overall performance at one station is also often assessed with a global judgment.

There is some evidence for the predictive validity of MMIs judgments, (Eva et al. 2009; Knorr et al. 2018) but what exactly MMIs measure has yet to be concretely identified (Knorr and Hissbach 2014). Specifically, there has been a debate about whether MMIs measure one construct (i.e., a general factor: some applicants perform better than others, regardless of dimensions or stations), one construct per station (i.e., exercise approach: some applicants perform better at some stations than others, regardless of dimensions), or one construct per social skill dimension (i.e., dimensional approach: some applicants perform

better on some dimensions than others, regardless of stations). Previous findings have been inconclusive (Knorr and Hissbach 2014; Rees et al. 2016). This state of research is grounded in universal ambiguity and disagreement about constructs assessed in ACs in general (e.g., Bowler and Woehr 2006; Putka and Hoffman 2013; Jackson et al. 2016). Here we aimed to contribute to this debate and provide a detailed investigation into the construct validity of MMIs. This investigation of constructs validity includes the assessments of dimensionality (i.e., identification and quantification of performance facets assessed with MMIs, e.g., estimating the amount of variance that is attributable to social skill-specific performance differences) and discriminant validity (i.e., relations with other selection measures, e.g., GPA). In doing so, we move beyond previous research by applying a more comprehensive analytical approach (i.e., Bayesian generalizability analysis) that is suited to fully consider the complexity of the data structure.

Dimensionality

Regarding the number and content of factors assessed with MMIs, previous analyses have been mostly exploratory with contradictory results: A factor analysis by Lemay and colleagues (2007) identified one factor for each station, Hecker et al. (2009) performed multiple factor analyses for each dimension and reported one, two, and three-factor solutions, and Oliver and colleagues (2014) found support for a two-factor dimension-based solution. Recently, however, factor analytical methods to investigate the construct validity of AC-like procedures have been criticized (e.g., Putka et al. 2011; Putka and Hoffman 2013): They have difficulty handling complex designs (i.e., “ill-structured” designs in which applicants and raters are neither fully crossed nor nested, which is common in ACs/MMIs), they oversimplify underlying variance components (e.g., inconsistencies across raters are not controlled for), and they confound reliable and unreliable sources of variance. This makes it impossible to disentangle specific exercise, dimension, or general effects (Jackson et al. 2016). Using a generalizability theory (G-theory; Brennan 2001) framework (i.e., the

decomposition of observed variance components due to different facets and their interactions; see Table 1 for an overview of variance components in the present study) and controlling for complex designs and unwanted variance, both Jackson et al. (2016) and Putka and Hoffmann (2013) found that most variance in ACs is explained by exercises and general performance facets and much less by behavioral dimensions. Whereas a variety of MMI studies (e.g., Eva and Reiter et al. 2004; Eva and Rosenfeld et al. 2004; Roberts et al. 2008; Eva et al. 2009; Roberts et al. 2010; Hecker and Violato 2011; Eva and Macala 2014; Oliver et al. 2014; Sebok et al. 2014) used G-theory methods, they did not control for complex designs or did not include all possible variance components (e.g., rater-related effects, dimension-related effects, sample effects). Failure to model such potential sources of (error) variance may result in misleading and biased estimates (Zaidi et al. 2017). Thus, to thoroughly study the dimensionality of MMI performance ratings, our aim here is to transfer the methods used in recent AC research (e.g., Bayesian generalizability approaches) to MMIs, which enables us to include all possible variance components and their interaction effects.

[Table 1 near here]

Discriminant validity

Regarding relations with other (selection) measures, MMI scores have shown low or no correlations with traditional admissions tools (e.g., interview, GPA, cognitive ability; Eva and Rosenfeld et al. 2004; Hecker et al. 2009; Rees et al. 2016) emotional intelligence (Yen et al. 2011), or most of the Big Five (e.g., Kulasegaram et al. 2010; Oliver et al. 2014). Furthermore, scores from MMIs including only interview stations were strongly correlated ($r = .59$) with scores from MMIs including different kinds of exercises (i.e., interviews, role-plays, group tasks; Gafni et al. 2012). This shows that while it is unclear what MMIs measure, different MMIs with different exercises seem to somehow converge. Most previous research has focused on relations between overall MMI scores and other measures and did not investigate the impact of different performance facets in MMIs. However, if we assume, for

example, that MMIs measures distinct factors for each station, we need to consider the correlation between results for the individual stations and other measures. The previous focus on correlations with overall MMI scores is only warranted if we reliably measure one general factor. Therefore, we aimed to identify the importance of the performance facets and then to correlate these facets with other selection measures to gain more differentiated insights about discriminant validity.

Present research

First, we identified the importance of performance facets assessed in MMIs. For this, we applied Bayesian G-theory to assess the amount of reliable variance that could be attributed to a general factor, to exercises, and to dimensions. Specifically, we investigated whether a distinction of a priori chosen dimensions is empirically warranted. In doing so, we controlled for unwanted variance (e.g., rater bias) as well as complex designs. Second, we investigated the relations between the resulting components and other selection measures (i.e., GPA, knowledge test, and application essay). To achieve these goals, we used data from three samples of applicants applying to medical school at a German University. All applicants were administered to MMIs and other selection procedures. The development of the MMIs was very similar to approaches employed at other universities (e.g., Eva and Rosenfeld et al. 2004; Ziv et al. 2008; Oliver et al. 2014; Knorr et al. 2018) and included role-play stations with professional actors and interview stations.

Methods

To support open science, we will report how we determined our sample size, all data exclusions, all manipulations, and all measures in the study. Furthermore, in an online supplement, we provide the data and R-codes necessary to reproduce the results (osf.io/d8qmu). The university's Institutional Review Board approved this study.

Participants

Applicants

Out of an initial 475 applicants who applied for medical school and took part in the MMIs, 376 applicants (age: $M = 19.20$, $SD = 1.85$; 267 female) gave informed consent for further data analyses (no further exclusions). Applicants were preselected on the basis of their GPA (3.6 or better),¹ and data were taken from three different selection samples (6 months apart).²

Rater

In each selection sample, 48 professional physicians (in teams of two raters) evaluated the applicants. Each rater team was assigned to one station (4 teams per station per sample) and evaluated 40 applicants on the corresponding dimensions.

Procedures

In all samples, participants provided their application essays and took a medical-scientific understanding test (MSU) before taking the MMIs. Out of the 10 possible stations, we included 6 stations that were consistent across samples and included multiple skill ratings in our analyses (see Table 2 for an overview of the stations we used and their corresponding dimensions).³ At each station, participants were allowed 90 seconds to read the respective mission task. After that, depending on the station at hand, participants either interacted with professional actors or watched a short video concerning a doctor-patient interaction and answered standardized questions. For the role-plays they were observed by raters through a one-way mirror whereas for the interview raters were in the same room. After each 5-minute station, participants switched to the next station, and raters were asked to judge participants on the corresponding dimensions. For the raters to develop a common understanding of the assessment scales, they received rater training the morning before the MMIs with a theoretical (e.g., rater bias; separating observation and evaluation) and practical background (appraising simulation candidates/videos and moderated discussion).

[Table 2 near here]

Measures

Judgments

Depending on the stations, raters had to judge applicants on different social skills (i.e., dimensions, see Table 2) that were chosen before the procedure took place. The ones under consideration included information handling, relationship building, emotional management, and emotional understanding⁴ (cf. Lievens 2013; Libbrecht et al. 2014). All dimensions were rated on a 6-point Likert scale ranging from 0 to 5. The anchor specifications were based on aspects of the individual station (e.g., accident station, information handling: applicant acquired all necessary information regarding the accident; relationship building: applicant encouraged the actor to communicate) and simultaneously constructed to complement the anchor descriptions in the other stations. In addition, a global overall judgment was assessed at all stations (cf. Eva and Reiter et al. 2004; Eva and Rosenfeld et al. 2004) concerning the question of whether one could imagine this applicant in her/his practical year ranging from 0 (*not very imaginable*) to 5 (*very imaginable*).

Further selection measures

Further selection measures included GPA, the MSU, and an application letter. As pre-selection was based on the applicant's GPA, the participants exhibited a restricted GPA range (3.6 or better). The MSU, a computer test, measures basic competencies in absorbing, analyzing, and interpreting complex scientific and medical information (Cronbach's alpha = .82). Regarding the application letter, applicants were asked to describe their motivation for the course and place of study, previous experiences, and achievements. On the basis of a comprehensive assessment scheme (e.g., ratings of social competence, experience, reasoning, formal aspects), each letter was then evaluated by five physicians ($ICC(1, k) = .58$).

Analytic strategy

To quantify the roles of different performance facets, we estimated variance components on the basis of the post-exercise dimension ratings (excluding the global judgment, which was not a desired social skill) provided by raters. Similar to Putka and

Hoffman (2013), we first estimated single- and average-rater reliability based on Putka et al.'s (2008) $G(q,k)$ coefficient. We then used a G-theory approach to decompose different variance components. Another option that is sometimes used to control for specific design components within MMIs (e.g., rater effects, exercise effects), is the Multi-faceted Rasch Model (MFRM; see Sebok et al. 2014 for an application). MFRM is somehow similar to G-theory, but its focus lies on the estimation of corrected individual applicant, exercise, or rater scores (Iramaneerat et al. 2008). Another difference of G-theory, compared to MFRM, is the possibility to extend it to acknowledge complex (i.e., ill-structured) measurement designs (see Putka et al. 2008), which are typical for MMIs. Thus, to analyze an MMI procedure as a whole and to investigate the relative impact of different aspects, G-theory is the more common and straightforward approach.

For our G-theory analyses, we specified applicants, raters, dimensions, exercises, samples, and all their possible interactions as crossed random factors in a Bayesian linear random effects model with the R package brms (Bürkner 2017). In contrast to classic random or mixed effects models (i.e., non-Bayesian approaches), Bayesian approaches overcome potential problems with variance component estimation when estimates are close to zero (Chung et al. 2013). Whereas for simple designs (i.e., designs with a simple nesting structure and only few variance components) results are typically similar across different estimation methods, Bayesian analyses have the advantage of being applicable to very complex data structures with a high number of variance components (see Jackson et al. 2016 for a more detailed discussion on the advantages of Bayesian estimation methods). All estimation parameters were chosen analogous to the approach by Jackson et al., (2016).⁵ Please note that raters were nested within exercises (as is common with MMIs) so we were not able to separate rater effects from rater within exercises effects.

To analyze the relations of performance facets with other selection procedures, we calculated individual estimates (predicted values) of each of the variance components

(involving reliable variance). Thus, for each applicant we calculated an individual estimate on the general factor, on each station, on all dimensions, and on all station - dimension combinations. These estimates were then correlated with other selection measures. We also explored the relations of performance facets with an aggregate of the overall global score across (and within) exercises. For data, R-code, and the equation for the G-theory model see osf.io/d8qmu.

Results

Variance attribution

For descriptive statistics and reliability estimates, see Tables 2 and 3. All average-rater reliability estimates per exercise and dimensions indicated good reliability ($M = .66$), and the single-rater reliability estimates were still acceptable according to commonly used cut-offs ($M = .48$; Cicchetti 2001). In addition, Table 4 (Column 2) shows the decomposition of the 15 variance components divided into reliable variance, unreliable variance, and components that do not contribute to between-candidate variance (e.g., 5% of total variance was attributable to rater main effects).⁶ The mentioned other components refer to general exercise, dimension, and sample effects that are independent of applicants or raters (e.g., some exercises are harder than others, some samples received better ratings). In Column 3, the variance percentages were rescaled by not including these other components (e.g., 7% of the between-applicant variance, not including other components, was attributable to rater main effects). Column 4 shows the percentages within the reliable components as well as within the unreliable components (e.g., 13% of the unreliable variance was attributable to rater main effects).

The overall single-rater reliability for any given ratings was .48. That is, 48% of the between-applicant variance could be attributed to reliable differences between candidates. The largest amount of the reliable variance was accounted for by an applicant-exercise interaction (55%), which points to multidimensionality due to stations (individual differences in exercise-specific performances across dimensions; exercise approach). Another 21% of the reliable

variation was explained by applicant main effects (individual differences in overall performance across dimensions and exercises; general factor). A subject-dimension interaction (individual differences in dimension-specific performances across exercises; dimensional approach) accounted for only 4% of the reliable variance. Specific dimension-exercise combinations accounted for the remaining reliable variance (20%).

[Table 3 and 4 near here]

Relations with other selection measures

Table 5 presents the correlations between the individual estimates of each variance component and the other selection measures. Overall, the results showed evidence for discriminant validity because there were at best weak correlations (maximum : $r = .17$). The aggregated global judgment (obtained as the last question for each exercise) was strongly correlated with the general factor, whereas the global judgments within stations were strongly correlated with the individual estimates from the respective stations.

[Table 5 near here]

Discussion

In the current work, we assessed the construct validity of MMIs. For this, we analyzed dimensionality (i.e., importance of different performance facets) and discriminant validity. Moving beyond previous research, we applied Bayesian G-theory analyses, which made it possible to directly compare all of the relevant MMI approaches (i.e., general, exercise, and dimensional approaches). This was done by decomposing the underlying variance components while controlling for unwanted variance (e.g., variance related to rater influences) and complex design. We tested these research questions in a state of the art MMI used for the selection of medical students at a German University, which includes different stations (interviews and role-plays), carefully developed behaviorally anchored dimensions, and trained judges.

Dimensionality: Relevant sources of MMI variance

Our results suggest that MMIs contain three main sources of reliable variance: station-related variance, general performance variance, and station-specific dimension-related variance. This provides evidence for station-specific factors (consistent individual differences in performance for specific stations across dimensions and raters) as well as for a general factor (consistent individual differences in performance across dimensions, raters, and stations). Contrary to the idea that MMIs can measure distinct social skills within stations (Pau et al. 2013; Oliver et al. 2014), we found little evidence for dimension-specific constructs. Thus, in our case, the desired dimensions (e.g., information handling, relationship building) could not be distinctively assessed across stations. Whereas a significant amount of variance could be attributed to dimension-exercise interaction effects, these ratings of dimensions depended on the station in which they were assessed (Jackson et al. 2016). These effects might also be exaggerated due to the nesting of raters in stations and also speak against cross-situational dimensions. Thus, to some degree, observers have the same ideas about how to rank-order applicants regarding a particular dimension at one station, but these ratings lack common ground across stations. The overall composition of reliable variance components is very similar to the results found in Assessment Center studies (Putka and Hoffman 2013; Jackson et al. 2016) and also mirrors the variance decomposition from an MMIs construct validation study by Oliver et al. (2014).

Validity: What do MMIs measure?

Whereas the interpretation of dimensional facets is straightforward, it is unclear what general and station-specific facets represent. They did not overlap with other selection tools (GPA, MSU, and application essay) but were almost perfectly related to the aggregated global judgments across/within stations. In accordance with this, two alternative (but not mutually exclusive) explanations can be drawn:

First, station-specific as well as overarching performance facets might reflect an undifferentiated mixture of desired social skills assessed within and across the stations. For

example, for a high rating of information handling in the accident station, applicants had to acquire all necessary information, whereas, for a high rating on relationship handling, applicants were to encourage the role-player to communicate. These behaviors are not necessarily independent, as performance differences for all social skills could be grounded in a common set of individual differences between applicants (e.g., extraversion, cognitive ability, ability to identify task criteria; Collins et al. 2003; Kleinmann et al. 2011; Judge et al. 2013). This could lead to consistent performances across social skills (and to some degree across stations).

Second, the agreement between raters could also be explained by a shared halo effect (Viswesvaran et al. 2005; Tavares et al. 2016). That is, on the basis of (seemingly) job-irrelevant aspects, raters' shared impressions might overshadow their subsequent judgments. These effects could be both station-specific and general. First impressions based on the way people move, talk, and gesture (e.g., Ambady and Rosenthal 1992), have been connected to subsequent ratings in employment interviews as well as assessment centers (Barrick et al. 2010; Ingold et al. 2018) and might contribute to undifferentiated performance judgments across social skills within stations. Appearance-related characteristics (e.g., candidates' attractiveness or professional appearance) can also influence interview ratings (Hosoda et al. 2003; Barrick et al. 2009) and might lead to shared halo effects across stations.

Limitations and implications

We designed the presented MMIs to be comparable to other approaches applied at other universities and the general size of variance components is similar to previous AC/MMIs studies involving alike procedures. Still, we are aware that MMIs are a method whose implementation varies heavily across settings (Eva and Macala 2014; Eva et al. 2018; Reiter and Eva 2018) and we do therefore not claim to answer the construct validity question for all MMIs at every institution. We do however believe, that this research offers first hints on what MMIs can and cannot measure and urge other researchers to replicate the

quantification of multiple performance facets at their institutions to compare facet influences when including different stations, different social skills and / or different raters. The general framework of using Bayesian generalizability methods to estimate the amount of reliable variance while controlling for unwanted variance is applicable to all MMI designs and implementations.

Future research should also try to disentangle the concrete behavioral differences that give rise to the station-specific and general facets. Through intensive video analyses, differences in anchor-specified behavior could be examined (e.g., How much did applicants encourage the role-player? How much information did they acquire?). Such systematic behavioral observations can be used to compare the structure of observable behavioral skills with the structure of social skills defined a priori. Specific behavioral measures could furthermore be linked with physicians' ratings to estimate the impact of desired behaviors. In a similar way, physicians' ratings could be connected to appearance-based ratings (e.g., attractiveness), first-impression ratings (e.g., voice), or ratings of other behaviors (e.g., extraverted behavior) to estimate the impact of undesired aspects. Putting all of these together would allow for further fine-grained insights into what is measured in MMIs.

Regarding the development of future MMIs, one potential option could be to focus on one social skill per station, which is assessed at not only one (e.g., Lemay et al. 2007), but at multiple stations. There are reliable differences in exercise-specific performances, and if each station is centered around only one social skill that is necessary to perform well in the respective station, differences in performance (aside from shared halos) would likely be attributable to this specific skill. If the goal is to measure applicants' empathy, for example, two stations in which the actor can be persuaded only through empathic behavior could be designed, and raters could judge empathic performance at the respective stations. Due to having this skill assessed at multiple stations one would then still be able to disentangle station and dimension effects.

Conclusion

What is measured in MMIs? A detailed investigation of the dimensionality and discriminant validity of MMI ratings revealed (a) substantial individual differences in overall and exercise-specific MMI performance, (b) little evidence for the reliable assessment of individual differences in specific social skills, (c) near independence between MMI performance facets and other selection procedures, and (d) strong associations between performance facets and global judgements. Future research should focus on the concrete behavioral cues and physical appearances that underlie MMI performance facets and that might be more or less related to the social skills that future physicians should have.

References

- Albanese MA, Snow MH, Skochelak SE, Huggett KN, Farrell PM. 2003. Assessing personal qualities in medical school admissions. *Acad Med* 78(3):313–321.
- Ambady N, Rosenthal R. 1992. Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychol Bull* 111(2):256–274.
- Barrick MR, Shaffer JA, DeGrassi SW. 2009. What you see may not be what you get: relationships among self-presentation tactics and ratings of interview and job performance. *J Appl Psychol* 94(6):1394–1411.
- Barrick MR, Swider BW, Stewart GL. 2010. Initial evaluations in the interview: relationships with subsequent interviewer evaluations and employment offers. *J Appl Psychol* 95(6):1163–1172.
- Bowler MC, Woehr DJ. 2006. A meta-analytic evaluation of the impact of dimension and exercise factors on assessment center ratings. *J Appl Psychol* 91(5):1114–1124.
- Brennan RL. 2001. *Generalizability Theory*. New York: Springer.
- Bürkner P-C. 2017. brms: An R package for Bayesian multilevel models using Stan. *J Stat Soft* 80(1):1–28.
- Chung Y, Rabe-Hesketh S, Dorie V, Gelman A, Liu J. 2013. A nondegenerate penalized likelihood estimator for variance parameters in multilevel models. *Psychometrika* 78(4):685–709.
- Collins JM, Schmidt FL, Sanchez-Ku M, Thomas L, McDaniel MA, Le H. 2003. Can basic individual differences shed light on the construct meaning of assessment center evaluations? *Int J Sel Assess* 11(1):17–29.
- Eva KW, Macala C. 2014. Multiple mini-interview test characteristics: 'tis better to ask candidates to recall than to imagine. *Med Educ* 48(6):604–613.
- Eva KW, Macala C, Fleming B. 2018. Twelve tips for constructing a multiple mini-interview. *Med Teach*.

- Eva KW, Reiter HI, Rosenfeld J, Norman GR. 2004. The Relationship between interviewers' characteristics and ratings assigned during a multiple mini-interview. *Acad Med* 79(6):602–609.
- Eva KW, Reiter HI, Trinh K, Wasi P, Rosenfeld J, Norman GR. 2009. Predictive validity of the multiple mini-interview for selecting medical trainees. *Med Educ* 43(8):767–775.
- Eva KW, Rosenfeld J, Reiter HI, Norman GR. 2004. An admissions OSCE: The multiple mini-interview. *Med Educ* 38(3):314–326.
- Gafni N, Moshinsky A, Eisenberg O, Zeigler D, Ziv A. 2012. Reliability estimates: behavioural stations and questionnaires in medical school admissions. *Med Educ* 46(3):277–288.
- Hecker K, Donnon T, Fuentealba C, Hall D, Illanes O, Morck DW, Muelling C. 2009. Assessment of applicants to the veterinary curriculum using a multiple mini-interview method. *J Vet Med Educ* 36(2):166–173.
- Hecker K, Violato C. 2011. A generalizability analysis of a veterinary school Multiple Mini Interview: effect of number of interviewers, type of interviewers, and number of stations. *Teach Learn Med* 23(4):331–336.
- Hosoda M, Stone-Romero EF, Coats G. 2003. The effects of physical attractiveness on job-related outcomes: a meta-analysis of experimental studies. *Pers Psychol* 56(2):431–462.
- Hughes P. 2002. Can we improve on how we select medical students? *J Royal Soc Med* 95(1):18–22.
- Ingold PV, Dönni M, Lievens F. 2018. A dual-process theory perspective to better understand judgments in assessment centers: The role of initial impressions for dimension ratings and validity. *J Appl Psychol*.
- Iramaneerat C, Yudkowsky R, Myford CM, Downing SM. 2008. Quality control of an OSCE using generalizability theory and many-faceted Rasch measurement. *Advances in Health Sciences Education* 13(4):479–493.

- Jackson DJR, Michaelides G, Dewberry C, Kim Y-J. 2016. Everything that you have ever been told about assessment center ratings is confounded. *J Appl Psychol* 101(7):976–994.
- Judge TA, Rodell JB, Klinger RL, Simon LS, Crawford ER. 2013. Hierarchical representations of the five-factor model of personality in predicting job performance: integrating three organizing frameworks with two theoretical perspectives. *J Appl Psychol* 98(6):875–925.
- Kleinmann M, Ingold PV, Lievens F, Jansen A, Melchers KG, König CJ. 2011. A different look at why selection procedures work. *Organ Psychol Rev* 1(2):128–146.
- Knorr M, Hissbach J. 2014. Multiple mini-interviews: same concept, different approaches. *Med Educ* 48(12):1157–1175.
- Knorr M, Schwibbe A, Ehrhardt M, Lackamp J, Zimmermann S, Hampe W. 2018. Validity evidence for the Hamburg multiple mini-interview. *BMC Med Educ* 18(1):106.
- Kulasegaram K, Reiter HI, Wiesner W, Hackett RD, Norman GR. 2010. Non-association between Neo-5 personality tests and multiple mini-interview. *Adv Health Sci Educ* 15(3):415–423.
- Lemay J-F, Lockyer JM, Collin VT, Brownell AKW. 2007. Assessment of non-cognitive traits through the admissions multiple mini-interview. *Med Educ* 41(6):573–579.
- Libbrecht N, Lievens F, Carette B, Côté S. 2014. Emotional intelligence predicts success in medical school. *Emotion* 14(1):64–73.
- Lievens F. 2013. Adjusting medical school admission: assessing interpersonal skills using situational judgement tests. *Med Educ* 47(2):182–189.
- Oliver T, Hecker K, Hausdorf PA, Conlon P. 2014. Validating MMI scores: are we measuring multiple attributes? *Adv Health Sci Educ* 19(3):379–392.
- Patterson F, Knight A, Dowell J, Nicholson S, Cousans F, Cleland J. 2016. How effective are selection methods in medical education? A systematic review. *Med Educ* 50(1):36–60.

- Patterson F, Lievens F, Kerrin M, Zibarras L, Carette B. 2012. Designing selection systems for medicine: The importance of balancing predictive and political validity in high-stakes selection contexts. *Int J Sel Assess* 20(4):486–496.
- Pau A, Jeevaratnam K, Chen YS, Fall AA, Khoo C, Nadarajah VD. 2013. The Multiple Mini-Interview (MMI) for student selection in health professions training - a systematic review. *Med Teach* 35(12):1027–1041.
- Putka DJ, Hoffman BJ. 2013. Clarifying the contribution of assessee-, dimension-, exercise-, and assessor-related effects to reliable and unreliable variance in assessment center ratings. *J Appl Psychol* 98(1):114–133.
- Putka DJ, Lance CE, Le H, McCloy RA. 2011. A cautionary note on modeling multitrait-multirater data arising from ill-structured measurement designs. *Organ Res Methods* 14(3):503–529.
- Putka DJ, Le H, McCloy RA, Diaz T. 2008. Ill-structured measurement designs in organizational research: implications for estimating interrater reliability. *J Appl Psychol* 93(5):959–981.
- Rees EL, Hawarden AW, Dent G, Hays R, Bates J, Hassell AB. 2016. Evidence regarding the utility of multiple mini-interview (MMI) for selection to undergraduate health programs: A BEME systematic review: BEME Guide No. 37. *Med Teach* 38(5):443–455.
- Reiter H, Eva K. 2018. Vive la différence: The freedom and inherent responsibilities when designing and implementing multiple mini-interviews. *Acad Med* 93(7):969–971.
- Roberts C, Rothnie I, Zoanetti N, Crossley J. 2010. Should candidate scores be adjusted for interviewer stringency or leniency in the multiple mini-interview? *Med Educ* 44(7):690–698.
- Roberts C, Walton M, Rothnie I, Crossley J, Lyon P, Kumar K, Tiller D. 2008. Factors affecting the utility of the multiple mini-interview in selecting candidates for graduate-

- entry medical school. *Med Educ* 42(4):396–404.
- Sebok SS, Luu K, Klinger DA. 2014. Psychometric properties of the multiple mini-interview used for medical admissions: findings from generalizability and Rasch analyses. *Adv Health Sci Educ* 19(1):71–84.
- Stern DT, Wojtczak A, Schwarz MR. 2003. The assessment of Global Minimum Essential Requirements in medical education. *Med Teach* 25(6):589–595.
- Tavares W, Ginsburg S, Eva KW. 2016. Selecting and Simplifying: Rater Performance and Behavior When Considering Multiple Competencies. *Teach Learn Med* 28(1):41–51.
- Viswesvaran C, Schmidt FL, Ones DS. 2005. Is there a general factor in ratings of job performance? A meta-analytic framework for disentangling substantive and error influences. *J Appl Psychol* 90(1):108–131.
- Yen W, Hovey R, Hodwitz K, Zhang S. 2011. An exploration of the relationship between emotional intelligence (EI) and the Multiple Mini-Interview (MMI). *Adv Health Sci Educ* 16(1):59–67.
- Zaidi NLB, Swoboda CM, Kelcey BM, Manuel RS. 2017. Hidden item variance in multiple mini-interview scores. *Adv Health Sci Educ Theory Pract* 22(2):337–363.
- Ziv A, Rubin O, Moshinsky A, Gafni N, Kotler M, Dagan Y, Lichtenberg D, Mekori YA, Mittelman M. 2008. MOR: a simulation-based assessment centre for evaluating the personal and interpersonal qualities of medical school candidates. *Med Educ* 42(10):991–998.

Footnotes

¹This was translated from German high-school grades. A GPA of 3.6 translates into a very good high school diploma score of 1.4.

²We used data from these three samples because they included a large number of identical stations (six). Prior and succeeding samples used different stations and station-dimension combinations so that it was not feasible to merge these datasets. With an N of 376, we had sufficient power (.84) to detect small effect sizes ($r = .15$).

³Four stations were not included because one station was not identical across all three samples, two stations focused on a non-social skill (i.e., practical handling), and one station was built solely around the dimension ethical reasoning, which made it impossible to disentangle effects of stations from effects of dimensions.

⁴The ratings of emotional understanding were obtained via two items (concerning the emotional understanding of persons A and B after watching a video of a dyadic interaction).

⁵In contrast to Jackson et al., we did not include item effects. This was because three out of our four dimensions included one item only (thus confounding effects of items and dimensions). Furthermore, we used the default priors of brms (i.e., improper flat priors), which influence potential results as little as possible.

⁶Please note that the amount of (reliable) variance attributable to different components varied depending on the aggregation of ratings (e.g., aggregating ratings to one overall score, aggregating to dimensional scores, aggregating to exercise scores; Putka and Hoffman 2013). For this analysis, we focused on the decomposition of non-aggregated post-exercise dimension ratings. When aggregating to dimensional or overall scores, applicant main effects surpassed applicant-exercise interaction effects. However, applicant-dimension interaction effects were negligible, even when aggregating into dimensional scores. For all variance components of aggregated scores, see osf.io/d8qmu.

Tables

Table 1

Decomposition of observed variance in rater ratings

Variance component	Meanings
Reliable variance	
σ^2 applicants	individual differences in overall performance across dimensions, exercises, and raters (some applicants perform better than others, regardless of exercises or dimensions; <i>general factor</i>)
σ^2 applicants $\times$ dimension	individual differences in dimension-specific performances across exercises and raters (some applicants perform better on some dimensions than others, regardless of exercises; <i>dimensional approach</i>)
σ^2 applicants $\times$ exercise	individual differences in exercise-specific performances across dimensions and raters (some applicants perform better on some exercises than others, regardless of dimensions; <i>exercise approach</i>)
σ^2 applicants $\times$ dimension $\times$ exercise	individual differences in dimension-specific performance within specific exercises across raters (some applicants perform better on some dimension–exercise combinations than others)
Unreliable variance	
σ^2 applicants $\times$ rater:exercise	some applicants are rated higher by some raters (nested in exercises) than others, regardless of dimension
σ^2 applicants $\times$ dimension $\times$ rater:exercise + residual	some applicants are rated higher by some raters (nested in exercises) than others—but it depends on the dimension
σ^2 rater:exercise*	some raters (nested in exercises) tend to give more lenient/severe ratings than others, regardless of dimension
σ^2 rater:exercise $\times$ dimension*	some raters (nested in exercises) tend to give more lenient/severe ratings than others—but it depends on the dimension

Note. This table was adapted from Putka & Hoffman (2013) dimension = social skill; exercise = station. Please note that for designs usually applied in MMI research (in which raters are only assigned to one exercise), raters are nested within exercises (rater:exercise). Thus, variance components involving raters cannot be separated from variance components involving raters and exercises. In these cases (and in the present research) only 4 instead of 8 different unreliable variance components can be estimated. Furthermore, when including multiple samples, all applicants and raters are also nested within samples (as in the present research). *These variance components contribute to unreliable variance because applicants were not fully crossed with raters.

Table 2

Overview of stations, dimensions, and reliability estimates

	Information handling	Relationship building	Emotional understanding	Emotional management	Overall judgment
Station 1: Question an unwell person (role-play); with advance information the morning before the MMI	.59/.74	.46/.64			.48/.65
Station 2: Take care of a stranger after an accident (role-play)	.52/.68	.53/.70			.48/.65
Station 3: Persuade a patient (role-play)		.46/.63		.45/.63	.44/.61
Station 4: Observe a patient-physician interaction I (interview)			.50/.67	.49/.66	.43/.61
Station 5: Deliver bad news (interview in samples 1 & 2; role-play in sample 3)	.46/.63	.47/.65			.50/.67
Station 6: Observe a patient-physician interaction II. (interview)			.42/.60	.42/.59	.36/.53

Note. $N = 376$. Left: $G(q,1)$ = single-rater reliability: based on a single rater, proportion of between-applicant variance that is attributable to reliable differences between candidates (q -multiplier of .958). Right: $G(q,2)$ = based on average ratings, proportion of between-applicant variance that is attributable to reliable differences between candidates (q -multiplier of .458). In MMI studies, applicants are not fully crossed with raters. Thus, to some degree, rater main effects (e.g., leniency/severity) contribute to unreliable variance. Here, the q -multiplier is used to scale the contribution of rater main effects, depending on the amount of overlap between the sets of rater who rate each applicant (see Putka et al. 2008 for additional information). The results displayed here are calculated across all samples.

Table 3

Descriptive statistics for relevant variables

	<i>Range</i>	<i>Mean</i>	<i>Standard Deviation</i>
Station 1	0.50 – 5.00	3.43	0.97
Station 2	0.50 – 5.00	3.35	0.93
Station 3	0.25 – 5.00	3.21	1.06
Station 4	0.00 – 5.00	2.94	1.13
Station 5	0.00 – 5.00	3.17	0.93
Station 6	0.00 – 5.00	3.25	1.07
Information handling	1.00 – 4.83	3.17	0.65
Relationship building	1.50 – 4.75	3.41	0.64
Emotional understanding	0.00 – 5.00	3.14	0.88
Emotional management	0.33 – 5.00	3.08	0.78
Overall global score	0.67 – 4.75	3.26	0.65
Overall MMI score	0.88 – 4.75	3.22	0.56
GPA (grade point average)	714 – 831	752.94	24.44
MSU (medical scientific understanding)	13.00 – 53.00	31.28	8.43
Application essay	4.00 – 29.60	17.84	4.50

Note. $N = 376$. For station description see Table 2. GPA: 714 points translate into a GPA of 3.6, 1.4 German system; 831 points translate into a GPA of 4.0, 1.0 German system). Station scores were averaged across dimensions, dimension scores were averaged across stations, and the overall MMI score was averaged across stations and dimensions (scale: 0 to 5). The overall global score was excluded from these calculations and calculated separately across stations. The possible range for the MSU was 0 to 60 and for the applicant essay 0 to 40.

Table 4

Variance components and reliability estimates of MMI ratings

Variance component	Total variance (%)	Between-applicant variance (%)	Between-applicant subtotal (%)
Reliable variance			
σ^2 applicants	7.32	10.16	21.22
σ^2 applicants $\times$ dim	1.35	1.87	3.91
σ^2 applicants $\times$ ex	19.01	26.37	55.08
σ^2 applicants $\times$ dim $\times$ ex	6.83	9.47	19.78
Subtotal	34.51	47.87	
Unreliable variance			
σ^2 applicants $\times$ rater:ex	11.55	16.02	30.74
σ^2 applicants $\times$ dim $\times$ rater:ex + residual	18.89	26.21	50.28
σ^2 rater:ex*	5.01	6.95	13.34
σ^2 rater:ex $\times$ dim*	2.12	2.94	5.65
Subtotal	37.57	52.12	
Other components			
σ^2 dim	3.11		
σ^2 ex	1.28		
σ^2 dim $\times$ ex	1.92		
σ^2 sample	18.49		
σ^2 sample $\times$ dim	0.20		
σ^2 sample $\times$ ex	2.64		
σ^2 sample $\times$ dim $\times$ ex	0.27		
Estimated $G(q, 1)$		.48	

Note. $N = 376$. dim = dimension (i.e. social skill); ex = exercise (i.e., station). Estimated $G(q,1) = \text{reliable variance subtotal} / (\text{reliable variance subtotal} + \text{unreliable variance subtotal}) = \text{expected single-rater reliability for any given dimension ratings}$. Raters were nested in exercises (rater:ex). Please also note that all applicants and raters were nested within samples, thus all reliable and unreliable variance components also include the sample component. * These variance components contribute to unreliable variance because applicants are not fully crossed with raters. Given the ill-structured nature of the measurement design, their contributions to between-applicant variance were rescaled using the q -multiplier; see Table 2.

Table 5

Correlations of individual estimates on each of the variance components with other measures

	GPA	MSU	Application essay	Overall judgment
General factor	.08	-.01	.12*	.87***
Station 1	.04	.01	.17***	.21*** (.79***)
Station 2	.04	-.04	.03	.20*** (.77***)
Station 3	-.02	.00	.11*	.25*** (.82***)
Station 4	-.05	-.01	-.12*	.27*** (.73***)
Station 5	.01	-.04	.02	.21*** (.78***)
Station 6	.13*	.06	.01	.32*** (.74***)
Information handling	.12*	.06	.06	.08
Relationship building	-.05	-.07	.13**	.35***
Emotional understanding	.02	.02	-.07	.16**
Emotional management	-.01	.01	.00	.21***
Specific station x dimension interaction	-.10 to .10*	-.07 to .04	-.06 to .14**	-.03 to .23***
GPA (grade point average)	-	.30**	.05	-.06
MSU (medical scientific understand)	-	-	.03	-.01
Application essay	-	-	-	.14**

Note. $N = 376$. * $p < .05$, ** $p < .01$, *** $p < .001$. Overall judgment = aggregated global judgment across stations (results for the aggregation within the corresponding station in parentheses). For station description see Table 2. For the specific station-dimension interaction, we provided the range out of the possible 12 combinations; see online supplement osf.io/d8qmu.