
Assessing the Performance of Online Adjunct Instructor Candidates: Process Improvements and Lessons Learned

Steve Stokes

Brigham Young University-Idaho
stokess@byui.edu

Chad McLane

Brigham Young University-Idaho
mclanec@byui.edu

Brian L. Jones

Brigham Young University-Idaho
jonesbri@byui.edu

Abstract

In this study, we outlined some of the obstacles associated with screening and identifying high-quality online adjunct instructor candidates and described the steps that we took to improve our screening and hiring process. We discussed issues associated with collecting and using performance ratings in the context of a remote screening and hiring process. We also demonstrated a novel application of the Many Facet Rasch (MFR) model to ratings of online adjunct instructor candidate performance. The MFR model was useful for validating the results of our processes improvements. Investing in streamlining and focusing our Candidate Assessment Course (CAS) yielded screening results that were more valid and reliable while requiring fewer administrative resources.

Introduction

As the worldwide demand for online education grows, so does the need for institutions to identify and hire high-quality adjunct online instructors (Martin et al., 2019). In 2018, for example, remote adjunct instructors met an estimated 47% of the instructional staff needs in higher education (Snyder et al., 2019). Identifying and hiring remote adjunct instructors, however, represents a significant administrative challenge. Teaching remotely online is a complex task that requires a unique skill set above and beyond technical subject matter expertise. Identifying individuals and hiring individuals who possess this unique skill set can be difficult (Bawane & Spector, 2009; Martin et al., 2019; Richardson, et al., 2015; Thomas & Graham, 2019). In addition to possessing the necessary skills and qualifications, remote online instructors should fit well within the institution's community, culture, scope, and mission. This is especially true for private colleges and universities that strive to foster a unique culture or serve specific audiences (Mandernach, Donnelly, Dailey & Schulte, 2005). In addition, the remote and often asynchronous nature of the screening and hiring process also poses difficulties for institutions and candidates alike. Despite these and other inherent challenges, investing in hiring process improvements for remote online adjunct instructors often yields real and tangible benefits (Sixl-Daniell et al., 2006).

Review of the Literature

We acknowledge that modern distance education in the United States is systemically different from the traditional 19th Century cottage-industry approach to teaching and degree granting. The systems and processes that have supported traditional higher education institutions for generations are not immediately or directly transferable to online distance education environments. For example, Peters (1994) argued that distance education has taken advantage of industrial methods like systematic planning, division of labor, automation, mass production, reliance on front-end work, and quality control. Embracing the hallmarks of industrialization has made it possible for distance education to extend its reach, and improve its efficiency, beyond traditional educational processes (Zawacki-Richter et al., 2020).

However, the same industrial attributes that have contributed to the success of online education also present significant challenges. Recruiting, hiring, and maintaining a high-quality online instructor pool is a prime example. The fast-paced and expanding nature of online education requires online administrators to be simultaneously nimble and stable; flexible enough to expand quickly, but grounded enough to adhere to accreditation requirements and established programs of study. To meet the demands and satisfy the competing interests in online education, it has become increasingly necessary for institutions to develop efficient yet systematic and thorough online instructor recruitment and selection processes (Sixl-Daniell et al., 2006; Patrick & Yick, 2005; Schnitzer & Crosby, 2003).

Indeed, the benefits of a robust effective hiring process must be balanced with the realities of maintaining that process, and the best solutions are often organizationally specific. For example, online instructors can be hired by a centralized office like academic affairs, by individual academic departments themselves, or some combination of both (Magda et al., 2015). Some institutions screen online instructor applicants at the institutional level through a committee of the Chief Academic Officer, the Faculty Affairs manager, the Human Resources Manager, and relevant full-time faculty or Subject Area Coordinator (Sixl-Daniell et al., 2006).

Using high-level administrators to review applicants as outlined by Sixl-Daniell et al. (2006) gives the process a degree of authority

but also makes it more costly and time consuming. Systematic planning and a division of labor can help to address inefficiencies. While shifting applicant evaluations from high-level administrators to other qualified individuals can save time and money, it does not sufficiently address the need for a comprehensive review of the applicant's organizational fit (Schnitzer & Crosby, 2003). This comprehensive review should examine the applicant's presence and teaching practices in online learning environments (Martin et al., 2019; Richardson et al., 2015). Additionally, examining applicants' teaching strengths is best conducted in scenario-based responses, often in one or two-week simulation course environments (Schnitzer & Crosby, 2003).

Rubrics are commonly used to facilitate a systematic review of candidates and rate their performance on simulated course activities (Leyzberg et al., 2017; Patrick & Yick, 2005). Rubrics help maintain consistency among raters when they are sufficiently developed, clear, accurately employed, and evaluators are adequately trained in their use. For example, Patrick and Yick's (2005) Systematic Interview Rubric for Adjunct Applicant Review (SIRAAR) stresses the need for systematically collecting and analyzing data about potential online adjuncts' hireability. Their rubric was based on four categories of essential skills and qualities needed in online instructors, namely, (a) knowledge, (b) philosophical outlook, (c) personal qualities, and (d) skill sets. Patrick and Yick (2005) employed a five-point Likert-type scale to rate the candidates' interview responses in each of the four categories. Ratings from these initial interviews were then used as the basis for further department interviews and subsequent job offers.

While their systematic use of rubrics was noteworthy, Patrick and Yick's (2005) approach had two significant limitations. First, the ratings were based solely on applicant interviews and not on any demonstration of skill or proficiency. Accordingly, administrators were left relying on what applicants said about themselves rather than on any direct observation of skill or ability. Second, their study was limited by a small sample size and the use of three interviewers, which may have caused unwanted variation in the rating outcomes. Even with these limitations, however, Patrick and Yick (2005), along with Barrett (2010) and Sixl-Daniell et al. (2006), effectively demonstrated the importance and effectiveness of following Peters' (1994) recommendation to pursue *industrial* or systems-based approaches to distance education. Establishing tools and processes that help evaluators rate candidates efficiently, consistently, and accurately can improve the institutions' ability to identify highly qualified instructors that fit their mission and instructional model.

Simply having a rubric and conducting a comprehensive review of applicants, however, is not entirely sufficient. Proper application of screening rubrics by raters is vital. Eckes (2009) has demonstrated that even with trained raters, varying interpretations of rating scales and rater severity "threatens the validity of the inferences that may be drawn from the assessment outcomes" (p. 2). Ensuring that rubrics are clear and easily interpretable can assist with improving the reliability of rating outcomes. Additional training and calibration activities can also help address issues of reliability and validity.

Once rating data are collected from online instructor candidates, they can be analyzed in various ways to determine their reliability. For example, Many-facet Rasch (MFR) measurement is part of a family of item response theory (IRT) models (Rasch, 1993) designed to account for error variance associated with various facets of performance or assessment ratings (e.g., rater severity, item difficulty, rating occasion; Linacre, 2020). That is, MFR modeling explicitly weights rating data to account for variation that is introduced by various facets of the rating process. MFR modeling is especially helpful when researchers desire to simultaneously model multiple sources of variability in performance ratings, and can be used to identify or compensate for differences in rater leniency or severity (Myford and Wolfe, 2003).

MFR analysis is commonly reserved for educational settings where ratings are required, like language proficiency or writing assessment (Aryadoust, 2012). However, it can be used any time performance ratings with multiple criteria or multiple raters are conducted. MFR modeling is particularly useful in these contexts because it explicitly accounts for the variance associated with important facets of the rating process (e.g., rater severity, item difficulty, rating occasion) and places each level of each facet on a common scale. In the end, MFR modeling produces a weighted score that accounts for the extraneous variation introduced by each facet (Myford & Wolfe, 2003).

Study Context

As administrators at a large private religious university in the western United States, we typically hire and onboard over 100 new adjunct online instructors every semester. Our online programs serve thousands of students (nearly 15,000 full-time-equivalents in the fall 2019 semester) who enroll in well over 2,000 online sections each semester. For several years, we utilized an online Candidate Assessment Course (CAS) to evaluate and screen online instructor candidates. During this two-week simulation course, candidates had the opportunity to perform several instructor-related tasks and become acquainted with our institution's goals, processes, and expectations. For example, the candidates were asked to facilitate an online discussion regarding the institution's learning model with their peers. In another activity, they were provided with a sample email from a disgruntled student and asked to compose a response. As it was initially designed, the CAS experience lasted for two weeks. After the course, candidates were given an overall holistic performance rating on a 7-point scale and a simple hiring recommendation (e.g., yes, no, yes with reservations). Successful candidates were approved by department heads and then placed into teaching pools.

This process worked reasonably well for a period of time, but became unmanageable as the demand for instructors and applicants increased. Feedback from instructor managers and candidates indicated that the CAS course was cumbersome and time-consuming while producing inconsistent evaluations of applicants. Most candidates already had full-time employment, making it challenging to devote time and resources to a two-week unpaid evaluation course. Further, the Employment and Scheduling team wanted to be more confident that CAS ratings were valid and reliable indicators of actual online instructor quality and performance.

Study Background

This study started with the university's Employment and Scheduling organization enlisting the help of the Institutional Research & Assessment (IRA) and Online Instruction teams to streamline the screening and hiring process and reduce the workload for hiring

coordinators and the candidates themselves.

The purpose of this project was to update and improve the screening, interviewing, and evaluation processes for online instructor candidates and address the following research questions:

1. Can the MFR model (Linacre, 1990) successfully be used to calibrate and validate instructor candidate performance ratings in a simulated online course environment?
2. Can we create a shorter and more efficient candidate screening process that produces valid and reliable results?
3. What benefits can we realize from simplifying and streamlining our candidate screening process?

Participants

During the Winter 2020 semester, our hiring coordinators evaluated 208 online instructor candidates who participated in an updated and streamlined version of CAS. Ratings were analyzed using the MFR model to identify and account for unwanted variance across raters, rubric items, and CAS sessions. Candidates were not compensated for their time, but satisfactory completion of the CAS course would place them into the university's hiring pool for upcoming semesters. They were informed and gave consent to use their confidential performance ratings for research purposes. Each candidate was enrolled in one of three CAS sessions and evaluated by at least one of ten CAS raters. The ten CAS raters were experienced online instructors with strong performance records and had taught online for an average of over 22 semesters with the university. CAS evaluators were compensated for ten hours of work to evaluate up to 25 candidates each.

Method

Our method consisted of four major steps: (a) construct definition and course improvement, (b) rating system set up, (c) data collection, and (d) analysis.

Construct Definition and Course Improvement

The first goal was to shorten and simplify the CAS process for candidates and raters. To do this, we held an open discussion with administrators to determine the essential abilities and traits that would need to be assessed using performance ratings, and which abilities or qualifications could be objectively evaluated in other ways (e.g., whether or not a candidate holds an appropriate degree, or has teaching experience in their field, etc.). Only candidates who passed an objective evaluation screening were enrolled into CAS.

Candidates were assessed on four constructs meaningful to our private religious institution including, (a) build faith, (b) develop relationships with and among students, (c) inspire a love for learning, and (d) ability to follow CAS assignment instructions. These four abilities were not directly observable in application materials alone. Therefore, our goal was to measure them using performance ratings based on CAS course observations. Defining performance criteria in this way created greater clarity for candidates and hiring coordinators. It also created alignment between the candidate evaluation metrics and the metrics used to evaluate employed online instructors. In other words, it created much more consistency in our evaluation metrics and messaging throughout the entire online instructor lifecycle.

Second, the duration of the CAS course was reduced from two weeks to one week and focused more explicitly on five key activities: (a) self-introduction, (b) foundational principles discussion, (c) sample student email response, (d) teaching demonstration, and (e) providing student feedback. These activities were designed to allow candidates to demonstrate their fit for our program and familiarize themselves with our processes and expectations.

Third, we re-developed the evaluation rubric, course activities, and instructions to align with our identified standards. After testing several different rubric configurations, we settled on a four-point scale: (0) did not participate in or did not complete the activity, (1) candidate should not be recommended for hire, (2) candidate exhibits sufficient capability, and (3) candidate exhibited excellence in this area. The rubric also included brief descriptions of what each level of performance would look like within each activity. A team of experienced and high-performing adjunct online instructors was contracted and trained to rate candidates on each activity using the updated rubric and rating scale.

Rating System Setup

Ideally, performance ratings should be comparable across rating facets, which are the various elements of the rating system that could potentially introduce error variance (e.g., raters, items, occasions, etc.). Whenever that is the goal, the design of the rating system is critical: the collected ratings are only as good as the system used to generate them (Eckes, 2011). A critical aspect of the rating system is the resulting connectedness of data points. In order to determine if performance ratings are comparable across facets (i.e., raters), the resulting elements of each facet must be sufficiently connected to each other (Eckes, 2011). Put simply, to rigorously compare the severity or leniency of raters, all raters must be connected to each other through common rating points (i.e., a common candidate rating). As Eckes described (2001), “[the] lack of connectedness among elements of a particular facet (e.g., among raters) would make it impossible to calibrate all elements of that facet on the same scale; that is, the measures constructed for these elements (e.g., rater severity measures) could not be directly compared” (p. 152). Therefore, we had to ensure the process itself would produce sufficiently connected ratings before collecting any data.

Because fully-crossed data (i.e., where every candidate is rated by every rater) is prohibitively costly and time-consuming, we utilized an “incomplete design” herein every rater was connected to every other rater by a limited number of common rating points (Eckes, 2001, p. 154). In our case, each individual rater was directly connected to every other rater through a maximum of three

common rating observations. This design produced ratings that were sufficiently crossed to make outcomes comparable across raters. It also made the rating process less burdensome or time-consuming for raters and more cost-effective for the organization. Candidates were assigned to raters in advance via an internally developed algorithm executed in R (R Core Team, 2020). These candidate-rater assignments were then entered into the learning management system, as well as the rubric utilized by raters (see Table 1).

Table 1

Example of an incomplete crossed design.

Candidate	Rater1	Rater2	Rater3	Rater4	Rater5
Candidate1	1	1			
Candidate2	1		1		
Candidate3	1			1	
Candidate4	1				1
Candidate5		1	1		
Candidate6		1		1	
Candidate7		1			1
Candidate8			1	1	
Candidate9			1		1
Candidate10				1	1

Note. 1 indicates a rating-candidate assignment. This incomplete design connects every rater to every other rater 1 time in the most efficient way possible.


When rating data are sufficiently crossed, the resulting MFR scores are directly comparable, given what we know about the facets that exist in the model. In other words, the MFR model calibrates candidate scores to be comparable, regardless of which rater reviewed an individual candidate’s performance. It also provides model-data fit statistics that can be used for the purpose of rater training and calibration.

Data Collection

During the Winter 2020 semester, we collected performance rating data on 208 online instructor candidates who participated in the updated CAS course. After candidates had completed each activity, CAS evaluators reviewed the submissions and rated each candidate using the updated rubric (Figure 1). At the top of the rubric, raters were provided with a brief description of each standard to be rated, as shown in Figure 1.

Figure 1

Sample of the updated CAS rubric.



Analysis

Once the data were collected, they were analyzed using a MFR approach in the Facets software (Linacre, 2020). We included three facets in our model: (a) candidates, (b) raters, and (c) rubric items. In addition to estimates of global data-model fit, we also examined the usefulness of the model based on the proportion of outliers, as described by Eckes (2011). We did not experience any convergence issues and always reached normal model termination.

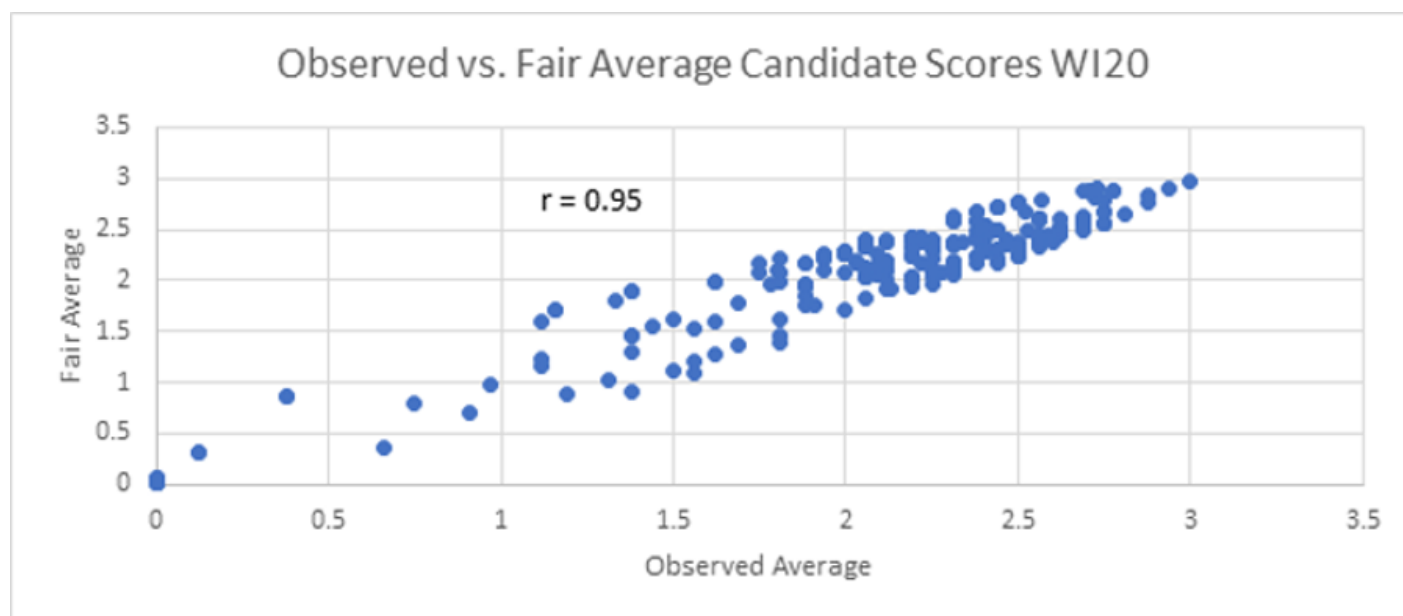
Findings and Conclusions

Our first research question was *can the many-facet Rasch (MFR) model successfully be used to calibrate and validate instructor candidate performance ratings in a simulated online course environment*. We were able to confidently differentiate between high and lower-quality online instructor candidates using an updated CAS and rating system and analyzing the data using the MFR model. This differentiation controlled for the severity or leniency of the rater and the relative difficulty of specific rubric items. However, as is the case with all statistical models, the first step was to check our assumptions and determine whether the data sufficiently fit the model. We found that the chi-squared value for global data-model fit indicated a significant *lack* of global data-model fit ($X^2=8,408.7$, $df = 4,822$, $p < 0.0001$). While this may seem concerning, a statistical lack of model-data fit should “come as no surprise; it is rather what can be predicted for nearly any set of empirical observations” (Eckes, 2011, p. 69). That is, real-world empirical data rarely conform ideally to theoretical expectations. Thankfully, there are alternative approaches to assessing global model fit that may help establish the practical usefulness of the model for the purpose of measurement. Eckes (2011) and Linacre (2014), for example, have suggested examining the model residuals (i.e., the standardized distance between the actual observed values and the model-predicted values). If relatively few of the standardized residuals are statistically large, that is a good indication of the practical utility of the model. Eckes (2011) noted, “satisfactory model fit is indicated when about 5% or less of (absolute) standardized residuals are ≥ 2 and about 1% or less of (absolute) standardized residuals are ≥ 3 ” (p. 69). In our study, out of the 5,102 data points used for estimation, 100 (< 2%) had (absolute) standardized residuals ≥ 2 and only 22 (< 1%) had (absolute) standardized residuals ≥ 3 . The small difference between the model-predicted and observed values gave us confidence that our model would be productive for measuring candidate performance and validating our rating approach.

Another encouraging observation was that the weighted “fair average” scores controlling for variation across facets and the raw “observed average” scores for each candidate were very highly correlated ($r = 0.95$; Figure 2). This meant that the statistically weighted scores and the raw observed average scores were very strongly related. Thus, weighting the scores to control for variation across facets did not have a large practical impact on outcomes. Ultimately, we determined that statistical weighting to control for variation across facets was probably unnecessary to achieve reliable results.

Figure 2

Weighted “fair average” scores vs. raw “observed average” scores.



The MFR model output seemed to be most useful for validation during the intermediate planning and development stages. We felt the need to determine the extent to which extraneous variation across facets might impact the rating outcomes and improve rater calibration and training if necessary. The MFR model helped us answer that question and reassured us that our updated course, rubrics, and processes were working properly.

In sum, we found that with the shortened CAS course, improved rubric, and focus on rater calibration, our raters could produce raw scores that were sufficiently valid and reliable. We concluded that statistical weighting across facets was unnecessary to do continually. We periodically monitor MFR results in order to promote ongoing improvement and to check our assumptions. Still, over time we have found that our updated processes produce raw scores that are sufficiently valid and reliable to meet our needs.

Our second research question was *can we create a shorter and more efficient candidate screening process that produces valid and reliable results*. It did not take long for us to recognize that the goals of streamlining the course and improving its content validity would go hand in hand. Shortening the course and removing extraneous content forced us to focus intently on the most important aspects of instructor performance outlined in our online instructor standards. As we refined the aspects of the course that directly targeted the standards and removed content that was perhaps only tangentially related, the length of the course was reduced. We were left with a leaner, shorter, and more focused candidate assessment course. All involved administrators agreed that our efforts to streamline the course improved its content validity and focused it more intently on the topics and constructs most valuable to our institution.


From a statistical perspective, the findings already discussed (e.g., very few large standardized residuals and a high correlation between observed and model-predicted outcomes) indicated that the data were acceptably reliable. Combined with the improved content validity of the course itself, the strong relationship between observed scores and model-predicted scores gave us confidence that CAS score outcomes were valid and reliable indicators of candidate ability level.

Our third research question was *what benefits can we realize from simplifying and streamlining our candidate screening process*. One of the primary advantages of using the MFR model and Facets (Linacre, 2020) software is that it provided us with a detailed analysis of each model facet. In our case, the facets raters, candidates, and rubric items. MFR model output provided us with a great deal of insight into the severity of the raters and the performance of candidates and the relative difficulty of individual rubric items, and the relative performance of candidates. For example, we learned that the most difficult task for candidates was Build Faith in the Providing Feedback activity, while the easiest task was to “follow directions” generally throughout the course. Develop Relationships in the Foundational Principles Discussion activity was also a relatively easy task. These findings confirmed many of our assumptions and helped us make targeted refinements and adjustments to the course content as well as the rubric itself.

Finally, Facets (Linacre, 2020) software provided an output table titled “Unexpected Responses.” This table summarized the instances where a rating from a particular evaluator for a specific candidate and item were statistically extreme or abnormal (Figure 3). For example, when Evaluator 3 rated Candidate 203120WI.1 on the Develop Relationships item of the Foundational Principles Discussion activity, the MFR model predicted a rating of 2.5. However, the score actually given was 0, which was statistically unexpected. These insights were especially useful in discussions about rater calibration and training, as well as model-data fit. In addition, they provided us with specific examples of places where we could be more focused and deliberate in our evaluator training and calibration efforts. Indeed, we learned that we still have some room for improvement in rater consistency, calibration, and training.

Figure 3

Facets software unexpected responses output.



Simplifying and streamlining the CAS course, in conjunction with MFR analysis, allowed us to be more confident and efficient in our hiring decisions. We found that we were able to collect more focused data on candidate performance using fewer administrative resources while reducing the evaluators' workload and the time commitment from the candidates themselves.

Discussion

In this study, we applied the MFR model to online instructor candidate performance ratings during a simulated online course to gain insight into our streamlined rating and screening processes. We found that MFR modeling provided useful insight into rubric development, rater calibration and training, and final score validation.

However, it should be noted that regardless of the statistical model used to analyze performance ratings, the scores that come out are only as good as the data put in. A thorough examination of the goals, needs, and objectives of the hiring process is necessary to establish a strong case for content validity. A well-designed rating system is also crucial to ensure that outcomes are valid and reliable from a statistical perspective. In our case, shortening the CAS course, focusing the content on the mission of our online programs, and aligning the rubric with the activity instructions proved invaluable. The MFR model served to validate and refine our understanding of candidate performance. Still, its use represented only a portion of the total effort that went into refining and improving our screening and hiring processes. Additional future analyses will be conducted to determine if actual online instructor performance was impacted by our improved screening and hiring processes.

Finally, a growing body of research suggests that diversity, equity, and inclusion (DEI) initiatives can greatly enhance discussions of institutional community, culture, and scope (Paredes-Collins, 2013; Griffin, 2019). Paredes-Collins (2013) pointed out, for example, the value of “compositional diversity” and stated that “a positive racial climate is essential to promote spiritual growth for students from diverse backgrounds. As a result, diversity is a compelling interest for Christian institutions” (p. 122). As online programs and initiatives spread to serve a broader and more diverse study body, research on role and impact of DEI in online adjust teaching pools

becomes critical. Hiring processes may also benefit from an evaluation of their online adjunct teaching pools in terms of DEI, as diversity and inclusion are relevant aspects of institutional community, culture, and scope.

References

- Aryadoust, V. (2012). Evaluating the psychometric quality of an ESL placement test of writing: A many-facets Rasch study. *Linguistics Journal*, 6(1), 8-33.
- Barrett, B. (2010). Virtual teaching and strategies: Transitioning from teaching traditional classes to online classes. *Contemporary Issues in Education Research*, 3(12), 17-20. DOI: <https://doi.org/10.19030/cier.v3i12.919>
- Bawane, J., & Spector, J. M. (2009). Prioritization of online instructor roles: Implications for competency based teacher education programs. *Distance Education*, 30(3), 383-397. <https://doi.org/10.1080/01587910903236536>
- Eckes, T. (2009). Many-facet Rasch measurement. *Reference supplement to the manual for relating language examinations to the Common European Framework of Reference for Languages: Learning, teaching, assessment*, 1-52.
- Eckes, T. (2011). Introduction to many-facet Rasch measurement. *Discourse Studies*, 8(1), 4149.
- Griffin, K. A. (2019). Institutional barriers, strategies, and benefits to increasing the representation of women and men of color in the professoriate: looking beyond the pipeline. *Higher Education: Handbook of Theory and Research: Volume 35*, 1-73.
- Linacre, J. (1990). *Many-faceted Rasch measurement*. MESA Press.
- Linacre, J. (2014) A user's guide to FACETS: Rasch-model computer programs. Chicago: <https://www.winsteps.com/index.htm>. Retrieved from <http://www.winsteps.com/facets.html>.
- Linacre, J. (2020) Facets computer program for many-facet Rasch measurement, version 3.83.2. Beaverton, Oregon: Winsteps.com.
- Leyzberg, D., Lumbroso, J., & Moretti, C. (2017, June). Nailing the TA interview: Using a rubric to hire teaching assistants. In *Proceedings of the 2017 ACM Conference on Innovation and Technology in Computer Science Education* (pp. 128-133). DOI: <https://doi.org/10.1145/3059009.3059057>
- Magda, A. J., Poulin, R., & Clinefelter, D. L. (2015). Recruiting, orienting, & supporting online adjunct faculty: A survey of practices. *WICHE Cooperative for Educational Technologies (WCET)*.
- Mandernach, B., Donnelly, E., Dailey, A., & Schulte, M. (2005). A faculty evaluation model for online instructors: Mentoring and evaluation in the online classroom. *Online Journal of Distance Learning Administration*, 8(3).
- Martin, F., Budhrani, K., Kumar, S., & Ritzhaupt, A. (2019). Award-winning faculty online teaching practices: Roles and competencies. *Online Learning*, 23(1), 184-205.
- Myford, C., & Wolfe, E. (2003). Detecting and measuring rater effects using many-facet Rasch measurement: Part I. *Journal of Applied Measurement*, 4(4), 386-422.
- Paredes-Collins, K. (2013) Cultivating Diversity and Spirituality: A Compelling Interest for Institutional Priority, *Christian Higher Education*, 12:1-2, 122-137. DOI: [10.1080/15363759.2013.739436](https://doi.org/10.1080/15363759.2013.739436)
- Patrick, P., & Yick, A. (2005). Standardizing the interview process and developing a faculty interview rubric: An effective method to recruit and retain online instructors. *The Internet and Higher Education*, 8(3), 199-212. DOI: <https://doi.org/10.1016/j.jheduc.2005.06.002>
- Peters, O. (1994). Distance education and industrial production: A comparative interpretation in outline (1973). *Otto Peters on distance education: The industrialization of teaching and learning*, 107-127.
- R Core Team (2020). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL: <https://www.R-project.org/>.
- Rasch, G. (1993). *Probabilistic models for some intelligence and attainment tests*. MESA Press.
- Richardson, J. C., Koehler, A. A., Besser, E. D., Caskurlu, S., Lim, J., & Mueller, C. M. (2015). Conceptualizing and investigating instructor presence in online learning environments. *The International Review of Research in Open and Distributed Learning*, 16(3). DOI: <https://doi.org/10.19173/irrodl.v16i3.2123>
- Schnitzer, M., & Crosby, L. S. (2003). Recruitment and development of online adjunct instructors. *Online Journal of Distance Learning Administration*, 6(2), 1-7.
- Sixl-Daniell, K., Williams, J. B., & Wong, A. (2006). A quality assurance framework for recruiting, training (and retaining) virtual

adjunct faculty. *Online Journal of Distance Learning Administration*, 9(1), 1-12.

Snyder, T., de Brey, C., & Dillow, S. (2019). *Digest of education statistics 2018*, NCES 2020-009. National Center for Education Statistics.

Thomas, J. E., & Graham, C. R. (2019). Online teaching competencies in observational rubrics: What are institutions evaluating?. *Distance Education*, 40(1), 114-132. DOI: <https://doi.org/10.1080/01587919.2018.1553564>

Zawacki-Richter, O., Conrad, D., Bozkurt, A., Aydin, C., Bedenlier, S., Jung, I. & Kerres, M. (2020). Elements of open education: an invitation to future research. *International Review of Research in Open and Distributed Learning*, 21(3), 319-334. DOI: <https://doi.org/10.19173/irrodl.v21i3.4659>CopiedAn error has occurred

Online Journal of Distance Learning Administration, XXIV, Number 4, Winter 2021
University of West Georgia, Distance Education Center
[Back to the Online Journal of Distance Learning Administration Contents](#)