



New Meridian

Technical Report New Jersey Graduation Proficiency Assessment Spring 2023



Table of Contents

Table of Contents	2
List of Tables	5
List of Figures	7
Chapter 1. Introduction.....	8
1.1. Background	8
1.2. Overview of the Technical Report	8
Chapter 2. Test Design and Development	11
2.1. Overview of the Test.....	11
2.2. ELAGP Claims and Subclaims	11
2.3. MATGP Claims and Subclaims.....	12
2.4. Test Development Activities	13
2.4.1. Item Development Process	13
2.4.1.1. Bank Analysis and Item Development Plan.....	14
2.4.1.2. Text Selection for ELA	14
2.4.1.3. Item Development.....	14
2.4.1.4. Item and Text Review Committees	14
2.4.1.5. Text Review.....	15
2.4.1.6. Content Item Review.....	15
2.4.1.7. Bias and Sensitivity Review	15
2.4.1.8. Editorial Review	15
2.4.2. Data Review	15
2.5. Operational Test Forms.....	16
Chapter 3. Assessment Administration.....	18
3.1 Test Security and Administration Policies	18
3.1.1. Secure vs. Nonsecure Materials	18
3.1.2. Scorable vs. Nonscorable Materials.....	18
3.2. Accessibility Features and Accommodations	19
3.2.1. Participation Guidelines for Assessments	19
3.2.2. Accessibility System.....	20
3.2.3. What are Accessibility Features?.....	20
3.2.4. Accommodations for Students with Disabilities and English Learners	20
3.2.5. Unique Accommodations.....	21

3.2.6. Emergency Accommodations.....	22
3.2.7. Student Refusal Form	22
3.3. Testing Irregularities and Security Breaches	22
3.4. Data Forensics Analyses.....	24
3.4.1. Response Change Analysis	24
3.4.2. Internet and Social Media Monitoring	25
3.4.3. Off-Hours Testing Monitoring.....	25
3.5. Quality Control of Test Administration.....	25
Chapter 4. Item Scoring.....	27
4.1. Machine-scored Items.....	27
4.1.1. Key-based Items	27
4.1.2. Rule-based Items	27
4.2. Human or Hand-scored Items.....	28
4.2.1. Scorer Training	29
4.2.2. Scorer Qualification.....	32
4.2.3. Managing Scoring.....	33
4.2.4. Monitoring Scoring.....	33
4.2.4.1. Second Scoring	33
4.2.4.2. Backreading.....	34
4.2.4.3. Validity	34
4.2.4.4. Calibration Sets	35
4.2.4.5. Inter-rater Agreement.....	35
4.3. Automated Scoring for PCRs	36
4.3.1. Concepts Related to Automated Scoring.....	36
4.3.1.1. Continuous Flow	36
4.3.1.2. Training of IEA Using Operational Data.....	36
4.3.1.3. Smart Routing.....	37
4.3.1.4. Quality Criteria for Evaluating Automated Scoring.....	37
4.3.1.5. Hierarchy of Assigned Scores for Reporting.....	37
4.3.2. Sampling Responses Used for Training IEA	38
4.3.3. Primary Criteria for Evaluating IEA Performance	38
4.3.4. Contingent Primary Criteria for Evaluating IEA Performance	39
4.3.5. Applying Smart Routing	39
4.3.6. Evaluation of Secondary Criteria for Evaluating IEA Performance.....	40

4.3.7. Inter-rater Agreement for Prose Constructed Response.....	42
4.4. Quality Control of Scoring.....	42
Chapter 5. Performance Standards and Standards Validation.....	44
5.1. Performance Standards.....	44
5.2. Standard Setting Processes.....	44
5.2.1. Performance Level Setting (PLS) Process for the PARCC/NMC Affiliate System	44
5.2.1.1. Research Studies.....	45
5.2.1.2. Pre-Policy Meeting.....	45
5.2.1.3. Performance Level Setting Meetings	46
5.2.1.4. Post-Policy Reasonableness Review	46
5.2.2. NJGPA Standards Validation.....	47
Chapter 6. Item Analysis.....	49
6.1. Overview	49
6.2. Data Screening Criteria	49
6.3. Classical Item Analysis.....	49
6.3.1. Item Difficulty	49
6.3.2. Response Option or Score Point Proportions	50
6.3.3. Item-Total Correlations	50
6.3.4. Item Bank Quality Control.....	51
Chapter 7. Item Response Theory Analysis, Calibration and Scaling	54
7.1. Overview	54
7.2. IRT Models.....	54
7.3. Summary Statistics and Distributions from IRT Analyses.....	54
7.4. Scale Scores	55
7.4.1. Establishing the Reporting Scales.....	56
7.4.2. ELA Reading and Writing Claim Scales	56
7.4.3. Subclaims Scales.....	57
7.4.4. Creating Conversion Tables.....	57
7.4.5. Score Distributions.....	63
Chapter 8. Student Demographics and Differential Item Functioning	73
8.1. Rules for Inclusion of Students in Analyses.....	73
8.2. Time to Attempt Assessment Items	73
8.3. Demographics.....	74
8.4. Differential Item Functioning.....	76

8.4.1. DIF Procedures	76
8.4.2. DIF Classification	77
8.4.3. Operational Analysis DIF Comparison Groups	78
8.4.4. Operational DIF Results	79
Chapter 9. Reliability.....	82
9.1. Overview	82
9.2. Reliability and SEM Estimation.....	82
9.2.1. Scale Score Reliability Estimation	83
9.3. Reliability Results.....	84
9.3.1. Raw Score Reliability Results	84
9.3.2. Scale Score Reliability Results	85
9.4. Reliability Results for Demographic Groups of Interest	86
9.5. Reliability Estimates of Subclaim Scores	90
9.6. Reliability of Classification	91
Chapter 10. Validity.....	94
10.1. Overview.....	94
10.2. Evidence Based on Test Content	94
10.3. Evidence Based on Internal Structure	95
10.3.1. Intercorrelations	95
10.3.2. Reliability.....	97
10.3.3. Local Item Dependence	97
10.3.3.1. Evidence from Special Studies	98
References	99
Appendices	101
Appendix 6: Classical Item Analysis Statistics by Form	101
Appendix 7: Item Response Theory Parameters.....	108

List of Tables

Table 2.1 Form Composition for ELAGP	12
Table 2.2 Contribution of Prose Constructed Response Items to ELAGP.....	12
Table 2.3 Form Composition for MATGP.....	13
Table 2.4 Testing Accommodations by Content Area.....	16

Table 6.1 Summary of Post-Administration P-Values for ELAGP Operational Items	51
Table 6.2 Summary of Post-Administration P-Values for MATGP Operational Items	52
Table 6.3 Summary of Post-Administration Item-Total Polyserial Correlations for ELAGP Operational Items.....	52
Table 6.4 Summary of Post-Administration Item-Total Polyserial Correlations for Mathematics Operational Items.53	
Table 7.1 IRT Parameter Estimates Summary for All Items for ELAGP and MATGP	55
Table 7.2 IRT Parameter Distribution by Year for All Items for ELAGP and MATGP	55
Table 7.3 Threshold Scores and Scaling Constants for ELAGP and MATGP	56
Table 7.4 Scaling Constants for Reading and Writing of ELAGP	57
Table 7.5 NJGPA Subclaim Theta Cut Scores	57
Table 7.6 Scale Score Cumulative Frequencies: ELAGP.....	63
Table 7.7 Scale Score Cumulative Frequencies: MATGP.....	65
Table 7.8 Subgroup Performance for ELAGP Scale Scores	69
Table 7.9 Subgroup Performance for MATGP Scale Scores	71
Table 8.1 Time In Seconds for TotalTest and by Unit: ELAGP.....	73
Table 8.2 Time In Seconds for Total Test and by Unit MATGP	74
Table 8.3 ELAGP and MATGP Test Takers.....	74
Table 8.4 DIF Categories for Dichotomous Selected-Response and Constructed-Response Items	78
Table 8.5 DIF Categories for Polytomous Constructed-Response Items	78
Table 8.6 DIF Comparison Groups	79
Table 8.7 Post-Administration Differential Item Functioning for ELAGP	80
Table 8.8 Post-Administration Differential Item Functioning for MATGP	80
Table 9.1 Summary of ELAGP Test Reliability Estimates for Total Group.....	84
Table 9.2 Summary of MATGP Test Reliability Estimates for Total Group.....	85
Table 9.3 ELAGP Scale Score Reliability	85
Table 9.4 MATGP Scale Score Reliability Results.....	85
Table 9.5 Summary of Test Reliability Estimates for Subgroups: ELAGP.....	87
Table 9.6 Summary of Test Reliability Estimates for Subgroups: MATGP	89
Table 9.7 ELAGP Reliability Estimates for Subscores.....	90
Table 9.8 MATGP Reliability Estimates for Subscores	91
Table 9.9 Reliability of Classification: Summary for ELAGP	92
Table 9.10 Reliability of Classification: ELAGP	92
Table 9.11 Reliability of Classification: Summary for MATGP	92
Table 9.12 Reliability of Classification: MATGP	93

Table 10.1 Average Intercorrelations between ELAGP Subclaims	96
Table 10.2 Average Intercorrelations and Reliability between MATGP Subclaims	97

List of Figures

Figure 7.1 ELAGP 2023 Forms TCCs.....	60
Figure 7.2 ELAGP 2023 Forms CSEM Curves.....	60
Figure 7.3 ELAGP 2023 Forms TIF Curves.....	61
Figure 7.4 MATGP 2023 Forms TCCs.....	61
Figure 7.5 MATGP 2023 Forms CSEM Curves	62
Figure 7.6 MATGP 2023 Forms TIF Curves	62
Figure 7.7 2023 ELAGP Score Distribution.....	65
Figure 7.8 2023 MATGP Score Distribution	67
Figure 7.9 2023 Reading Score Distribution.....	68
Figure 7.10 2023 Writing Score Distribution	68

Chapter 1. Introduction

1.1. Background

In May 2021, the New Jersey State Board of Education approved for publication in the New Jersey Register a notice of substantial changes to proposed amendments regarding N.J.A.C. 6A:8, Standards and Assessment. The notice included five new amendments, including the provision that students must take a state graduation proficiency assessment in grade 11. The New Jersey Graduation Proficiency Assessment (NJGPA) was administered for the first time in the spring of 2022 to students of the class of 2023.

On Tuesday, July 5, 2022, Governor Murphy signed P.L.2022, c.60 (ACS for A-3196/S-2349), which requires the State Board of Education to administer the NJGPA as a field test for the class of 2023. Beginning with the class of 2024, a passing score on the NJGPA is the First Pathway for a student to satisfy the New Jersey Graduation Testing requirements.

Purpose of the NJGPA

Referring to the NJGPA, New Jersey State statute § 18A:7C-6.1 indicates that “The test shall measure those basic skills all students must possess to function politically, economically and socially in a democratic society.” There are two test components of the NJGPA, the English Language Arts Graduation Proficiency (ELAGP) component and the Mathematics Graduation Proficiency (MATGP) component. Two core operational forms are administered each year for each component, a computer-based test (CBT) administered online and a paper-based test (PBT) form to support students needing paper-based accommodations. Each core operational test form is constructed with reference to assessment blueprints. Score comparability across the core operational forms is assured by including mostly previously administered items which have been calibrated to item bank measurement scales.

The NJGPA ELA forms are based on the grade 10 ELA standards and evidence statements. Operational ELA forms consist of two units, each designed to be completed in 90 minutes. The ELA forms each contain 20 items worth a total of 74 points. Additionally, the ELA forms each contain a Literary Analysis Task (LAT) and a short passage set in unit 1, and a Research Simulation Task (RST) in unit 2.

MATGP operational forms consist of items from Algebra I and Geometry NJSLA item banks. Operational math forms contain 30 items, worth a total of 55 points, administered in two 90-minute units. The MATGP forms utilize type I, II and III items.

1.2. Overview of the Technical Report

This Technical Report presents the results of analyses for the New Jersey Graduation Proficiency Assessment (NJGPA) operational administration of the spring 2023 forms. In this document, the term “operational items” refers to the scorable items that contribute to students’ raw scores and scale scores. The term “field-test items” refers to the nonscorable items that were newly developed to collect statistics for future test administrations, calibrated by using the field-test spring 2023 administration data. The report begins by providing explanations of the test form construction process, test administration, and scoring of the test items. Subsequent chapters of the report present descriptions of student characteristics

and results of statistical analyses, including classical test theory statistics, item response theory (IRT) scaling and parameters and differential item functioning (DIF).

The technical report contains the following chapters:

Chapter 2 – Test Development

This chapter describes the test design and procedures followed during the development of operational test forms, and characteristics for the 2023 forms.

Chapter 3 – Test Administration

This chapter presents the operational administration schedule, information regarding test security and confidentiality, accessibility features and accommodations, testing irregularities and security breaches and administration quality control procedures.

Chapter 4 – Item Scoring

This chapter explains the key-based and rule-based processes for machine-scored items, as well as the training and monitoring processes for human-scored items.

Chapter 5 – Standard Setting

This chapter describes the NJGPA performance levels and the processes followed to establish the performance thresholds for each subject area.

Chapter 6 – Item Analysis

This chapter describes the classical item-level statistics calculated for the operational test data, the flagging criteria used to identify items that performed differently than expected, and the results of these analyses are presented in this section.

Chapter 7 – IRT Analysis, Calibration and Scoring

This chapter presents the information related to the item response theory (IRT) models and the descriptive statistics of the item parameters. The development of the reporting scales and conversion tables, and scale score distributions are also presented. Note that all tests delivered in 2023 employed a pre-equated model, in which previously estimated item parameters are used to generate scoring tables.

Chapter 8 – Student Demographics and Differential Item Functioning

This chapter describes the rules for inclusion of students in analyses, distributions of students by grade, mode and gender, and distributions of demographic variables of interest. Also, methods for conducting differential item functioning analyses as well as corresponding flagging criteria are described. This is followed by definitions of the comparison groups and subsequent results for the comparison groups.

Chapter 9 – Reliability

This chapter presents the results of scale score reliability and internal consistency reliability analyses and corresponding standard errors of measurement for content area, mode (CBT or PBT) for all students, and subgroups of interest, is provided in this section. This is followed by reliability of classification (i.e., decision accuracy and decision consistency).

Chapter 10 – Validity

This chapter explains the validity evidence based on analyses of the content of the tests and the concurrence of data measuring related properties of the same variables. Correlations between subscores are reported by content area and mode (CBT or PBT) for all students.

Chapter 2. Test Design and Development

2.1. Overview of the Test

The ELAGP blueprint was derived from the Grade 10 NMC summative assessment blueprints with input from New Jersey ELA educators. The MATGP blueprint is similar to the NJSLA Algebra I and Geometry blueprints in item type and quantity and draws exclusively from those two item banks. Core forms are constructed to meet the blueprint and psychometric properties outlined in the test construction specifications.

The NJGPA was administered in either a computer-based test (CBT) or a paper-based test (PBT) format. ELAGP focused on reading comprehension skills and writing effectively when using and analyzing sources. MATGP focused on applying skills and concepts and included multi-step problems that require abstract reasoning and modeling of real-world problems. Each assessment was comprised of multiple units, and one of the mathematics units was split into calculator and non-calculator sections.

2.2. ELAGP Claims and Subclaims

The ELA GP has one master claim: New Jersey High School Graduation Readiness in ELA. Under the master claim, the ELAGP is comprised of two major claims (Reading Complex Text and Writing) as well as five subclaims, Vocabulary Interpretation and Use, Reading Literature, Reading Informational Text, Written Expression, and Knowledge of Language and Conventions. The form composition of the ELAGP is detailed in table 2.1. As described in the ELA Content Guide the ELAGP:

- Includes 20 items for a total of 74 points.
- Consists of two 90-minute units.
- Utilizes one test blueprint embedded with an LAT and short passage set (unit 1), and RST (unit 2).
- Includes items aligned to the grade 10 standards and evidence statements.
- Contains the following item types: Evidence-Based Selected Response (EBSR), Technology-Enhanced Constructed Response (TECR), and Prose Constructed Response (PCR) items.
- Incorporates writing tasks that will be scored using the Research Simulation Task and Literary Analysis Task Scoring Rubric.

ELA assessments contain two PCR tasks. These two writing prompts are each scored for two traits: (1) Reading Comprehension and Written Expression and (2) Knowledge of Language and Conventions. All traits are initially scored as either 0–4 or 0–3. The Written Expression traits are then multiplied by 3 (or weighted) to increase their contribution to the total score, making subclaim scores of 0, 3, 6, 9 and 12 possible. The maximum points possible for ELAGP PCR items are provided in Table 2.2.

Table 2.1 Form Composition for ELAGP

Unit Task	Claims/Subclaims	EBSR/TECR Points	PCR Points
1 Literary Analysis Task	Reading: Literary Text	8	4
1 Literary Analysis Task	Reading: Vocabulary	4	0
1 Literary Analysis Task	Writing: Written Expression	0	12
1 Literary Analysis Task	Writing: Knowledge of Language and Conventions	0	3
1 Short Passage Set	Reading: Informational Text	6	N/A
1 Short Passage Set	Reading: Vocabulary	2	N/A
2 Research Simulation Task	Reading: Informational Text	12	4
2 Research Simulation Task	Reading: Vocabulary	4	0
2 Research Simulation Task	Writing: Written Expression	0	12
2 Research Simulation Task	Writing: Knowledge of Language and Conventions	0	3
Total Points		36 Reading	8 Reading 30 Writing

Table 2.2 Contribution of Prose Constructed Response Items to ELAGP

Score	Literary Analysis Task Points	Research Simulation Task Points
Reading Comprehension	4	4
Written Expression	12	12
Knowledge of Language and Conventions	3	3
Total	19	19

2.3. MATGP Claims and Subclaims

As described in the Mathematics Content Guide for NJGPA, the math component consists exclusively of NJSLA-Algebra I and NJSLA-Geometry items. The form composition of MATGP is described in table 2.3.

Table 2.3 Form Composition for MATGP

Subclaims	Number of Items	Number of Points
Major Content	13–15	16–18
Additional & Supporting Content	9–11	12–14
Expressing Mathematical Reasoning	3	10
Modeling and Applications	3	15
	30	55

2.4. Test Development Activities

The test forms used in NJGPA testing are based on the NMC Affiliate program blueprints for ELA grade 10 and Integrated Math II summative assessments. For the MATGP, only Algebra I and Geometry content were included. Algebra II was excluded. Test development activities were conducted by New Meridian Corporation (NMC) under the guidance of NJDOE content leads, with input from New Jersey educators.

The process of creating forms is undertaken each year. Potential forms which meet the content guidelines and psychometric standards for both test components are identified by the NMC psychometric team and reviewed by NMC content specialists for subject matter appropriateness. Potential forms which meet psychometric and content guidelines are then reviewed by NJDOE staff and by New Jersey educators during test form verification.

Test items used in the ELA and Math components are drawn from banks of items used for NMC Affiliate ELA grade 10, Algebra I and Geometry. There are no field-test items on the MATGP.

2.4.1. Item Development Process

Items used on the NJGPA were created for the PARCC/NMC Affiliate program. The PARCC/NMC Affiliate program item development process is described in this section.

Test and item development activities were conducted by Pearson under the guidance and oversight of the K–12 state leads, the Higher Education Leadership Team, the Technical Advisory Committee, the Operational Working Group (OWG) members from each of the member states, the Text and Content Item Review Committees, and staff members from New Meridian, the project manager.

Developing high quality assessment content with authentic stimuli for computer-based tests (CBT) and paper-based tests (PBT) measuring rigorous standards was a complex process involving the services of many experts including assessment designers, psychometricians, managers, trainers, content providers,

content experts, editors, artists, programmers, technicians, human scorers, advisors, and members of the OWGs.

2.4.1.1. Bank Analysis and Item Development Plan

The summative item bank houses passages and items at each assessed grade level and subject. The bank supports the administration of the assessments, along with item release and practice tests. Items are developed and field tested annually. As the first step in annual item development cycle, the item development teams, in conjunction with members of the OWGs for ELA and mathematics, evaluated the strengths of the bank and considered the needs for future tests to establish an item development plan.

2.4.1.2. Text Selection for ELA

Using the Passage Selection Guidelines, English language arts subject matter experts were trained to search for appropriate passages to support an annual pool of passages for consideration. Guided by the test specifications documents, contracted subject matter experts worked to deliver the number of texts specified in the annual asset development plan. The Passage Selection Guidelines provided a text complexity framework and guidance on selecting a variety of text types and passages that allowed for a range of standards/evidences to be demonstrated to meet the assessment claims. ELA tests are based on authentic texts, including multi-media stimuli. Authentic texts are grade-appropriate texts that are not developed for the purposes of the assessment or to achieve a particular readability metric but reflect the original language of the authors. Staff content experts reviewed the passages for adherence to the Passage Selection Guidelines to meet the annual asset development plan described above in the number and distribution of genres and topics prior to review and consideration by the Text Review Committee. ELA item development was not conducted until after texts were approved by the Text Review Committee.

2.4.1.3. Item Development

Item writers were recruited, trained, and managed to develop the number of items specified in the annual asset development plan. Prior to committee reviews, staff reviewed the items for content accuracy, alignment to the standards, range of difficulty, adherence to universal design principles (which maximize the participation of the widest possible range of students), bias and sensitivity, and copy editing to enable the accurate measurement of the standards.

2.4.1.4. Item and Text Review Committees

Members of the OWGs for ELA and mathematics, state-level experts, local educators, post-secondary faculty, and community members conducted rigorous reviews of every item and passage being developed for the summative assessment system to ensure all test items are of the highest quality, aligned to the standards, and fair for all student populations. All reviewers were nominated by their state education agency. The purpose of the educator reviews was to provide feedback on the quality, accuracy, alignment, and appropriateness of the test passages and items developed annually for the summative assessments. The meetings were conducted either in person or virtually and included large group training on the

expectations and processes of each meeting, followed by breakout meetings of grade/subject working committees where additional training was provided.

2.4.1.5. Text Review

The Text Review Committee meets to review and approve the texts eligible for item development. Participants reviewed and provided feedback to about the grade-level appropriateness, content, and potential bias concerns, and reached consensus about which texts would move forward for development. The Text Review Committee was made up of members of both Content Item Review and Bias and Sensitivity Review Committees.

2.4.1.6. Content Item Review

During Content Item Review, committees reviewed and edited test items for adherence to the foundational documents, basic universal design principles, Accessibility Guidelines, associated item metadata, and the Style Guide. Committees accessed the item content within the Pearson Assessment Banking for Building and Interoperability (ABBI) system that previews how the passages and items will be displayed in an operational online environment. Committees also verified that the appropriate scoring rule had been applied to each item. The Content Item Review Committees were made up of OWG members and educators nominated by participating states.

2.4.1.7. Bias and Sensitivity Review

Educators and community members made up the committee that reviewed items and tasks to confirm that there were no bias or sensitivity issues that would interfere with a student's ability to achieve his or her best performance. The committee reviewed items and tasks to evaluate adherence to the Fairness and Sensitivity Guidelines, and to ensure that items and tasks do not unfairly advantage or disadvantage one student or group of students over another. Bias and Sensitivity Committee members made edits and modifications to items and passages to eliminate sources of bias and improve accessibility for all students.

2.4.1.8. Editorial Review

The Editorial Review Committee consisted of editors who reviewed up to 10 percent of the items and tasks. The committee reviewed the items for grammar, punctuation, clarity, and adherence to the Style Guide.

2.4.2. Data Review

Following testing, ELA educator content and bias committee members met to evaluate field-tested items and associated performance data. The focus of data review is to review items with questionable performance data. Items with data which suggest good performance are generally accepted into item banks. Items demonstrating poor performance data are excluded from item banks. Items with questionable performance data are reviewed for appropriateness, level of difficulty, and potential gender, ethnicity, or other bias. The data review committees then recommended acceptance or rejection of questionable field-test items for inclusion in item banks. Items that were approved by the committee are eligible for use on operational summative assessments.

Two data review committees were convened. The ELA educator committee reviews items with potentially problematic statistical parameters calculated from the performance of all students tested. The bias and sensitivity committee reviews items with statistical parameters that may indicate students from different demographic groups scored differentially.

2.5. Operational Test Forms

In conjunction with NJDOE content specialists, NMC creates one online core form and one paper-based core form for each NJGPA subject component... Core forms are constructed to meet the blueprint and psychometric properties outlined in the test construction specifications. Operationally, each component is assessed via one operational online form for the general student population.

Students requiring paper-based accommodations are assessed by the paper accommodated form (A-PBT). Other accommodations are applied to the online core form. Pearson tracks the online accommodated form as an operational form distinct from the operational online form for the general student population. This form is referred to as the online accommodated form (A-CBT). The A-CBT contains the same items as the online operational form. For the MATGP only, a fourth operational form, the Spanish language MATGP, is created by Pearson. The accommodations available in each content area are outlined in Table 2.4. Table 2.4 also details the accommodations applied to online and paper operational forms.

Table 2.4 Testing Accommodations by Content Area

Content Area	Online Accommodated Form (A-CBT)	Paper Accommodated Form (A-PBT)
English Language Arts	- Text-to-Speech - Screen Reader/Non-Screen Reader - American Sign Language - Closed Captioning - Human Reader/Signer - Human Signer for Test Directions - Assistive Technology Device - Refreshable Braille - Electronic Braille Response - Speech-to-Text - Human Scribe - Monitor Test Response - Word Prediction - Extended Time - Word-to-Word Dictionary - Directions Read Aloud in/Clarified in Student’s Native Language	- Paper-Based - Large Print - Braille with Tactile Graphics - Read Aloud - Extended Time - Word-to-Word Dictionary Directions Read Aloud in/Clarified in Student’s Native Language

Mathematics	• Text-to-Speech • Screen Reader/Non-Screen Reader • American Sign Language • Spanish Transadaptation • Spanish Text-to-Speech • Refreshable Braille • Tactile Graphics • Human Signer for Test Directions • Calculation Device and Mathematics Tools • Speech-to-Text • Spanish Speech-to-Text • Human Scribe/Human Signer • Spanish Human Scribe/Human Signer • Assistive Technology Device • Electronic Braille Response • Monitor Test Response • Word Prediction • Extended Time • Word-to-Word Dictionary • Directions Read Aloud in/Clarified in Student's Native Language • Spanish Human Reader	• Paper-Based • Large Print • Braille with Tactile Graphics • Read Aloud • Spanish Transadaptation • Spanish Large Print • Extended Time • Word-to-Word Dictionary • Directions Read Aloud in/Clarified in Student's Native Language
-------------	---	--

Chapter 3. Assessment Administration

3.1 Test Security and Administration Policies

The administration of the NJGPA is a secure testing event. Maintaining the security of test materials before, during, and after the test administration is crucial to obtaining valid and reliable results. School Test Coordinators are responsible for ensuring that all personnel with authorized access to secure materials are trained in and subsequently act in accordance with all security requirements.

School Test Coordinators must implement chain-of-custody requirements for specified materials. School Test Coordinators are responsible for distributing materials to Test Administrators, collecting materials from Test Administrators, returning secure test materials, and securely destroying certain specified materials after testing.

The administration of the summative assessment includes both secure and nonsecure materials, and these materials are further delineated by whether they are “scorable” or “nonscorable,” depending on whether the assessments were administered via paper/pencil (i.e., paper-based assessments) or online (i.e., computer-based assessments).

3.1.1. Secure vs. Nonsecure Materials

Participating states and agencies define secure materials as those that must be closely monitored and tracked to prevent unauthorized access to or prohibited use or distribution of secure content such as test items, reading passages, student work, and so on. For paper-based tests (PBTs), secure materials include both used and unused test booklets and used scratch paper, while for computer-based tests (CBTs), secure materials include student testing tickets, secure administration scripts (e.g., mathematics read-aloud), and used scratch paper. Nonsecure materials are defined as any authorized testing materials that do not include secure content (e.g., test items or student work). These include test administration manuals, unused scratch paper, and mathematics reference sheets that have not been written on, and so on.

3.1.2. Scorable vs. Nonscorable Materials

Paper-based assessments have both scorable and nonscorable materials, while computer-based assessments have only nonscorable materials. Scorable materials for paper-based NJGPA assessments consist only of answer documents. Scorable materials must be returned to the vendor to be scored. All other materials for PBTs, such as blank (i.e., unused) test booklets, test administration manuals, scratch paper, mathematics reference sheets, and so on, are deemed nonscorable. For CBTs, there are no scorable materials as student work is submitted electronically for scoring. Thus, there are limited physical materials to return (e.g., secure administration scripts for certain accommodations).

Students taking the CBT may not have access to secure test materials, including printed student testing tickets, prior to testing. Printed mathematics reference sheets (if applicable) and scratch paper must be new and unmarked.

Students taking the PBT may not have access to scorable or nonscorable secure test content before or after testing. Scorable secure materials that are to be provided by Test Administrators to students include answer documents. Nonscorable secure materials that are distributed by Test Administrators to paper-based testing students include large-print test booklets, braille test booklets, scratch paper (paper used by students to take notes and work through items), and printed mathematics reference sheets.

School Test Coordinators are required to maintain a tracking log to account for collection and destruction of test materials, including mathematics reference sheets and scratch paper written on by students. As part of the test administration policy, schools are required to maintain the Chain-of-Custody Form or tracking log of secure materials for at least three years unless otherwise directed by state policy. Copies of the Chain-of-Custody Form for paper-based testing are included in each local education agency (LEA) or school's test materials shipment.

Test administrators are not to have extended access to test materials before or after administration (except for certain accessibility or accommodations purposes). Test administrators must document the receipt and return of all secure test materials (used and unused) to the school test coordinator immediately after testing.

All test security and administration policies are found in the *Test Coordinator Manual* and the *Test Administrator Manual*.

3.2. Accessibility Features and Accommodations

3.2.1. Participation Guidelines for Assessments

All students, including students with disabilities (SWDs) and English learners (ELs), are required to participate in statewide assessments and have their assessment results be part of the state's accountability systems, with narrow exceptions for ELs in their first year in a U.S. school, and certain SWDs who have been identified by the Individualized Education Program (IEP) team to take their state's alternate assessment. Federal laws governing student participation in statewide assessments include the Individuals with Disabilities Education Act of 2004 (IDEA), Section 504 of the Rehabilitation Act of 1973 (reauthorized in 2008), and the Elementary and Secondary Education Act (ESEA) of 1965, as amended by the Every Student Succeeds Act (ESSA). All students can receive accessibility features on the summative assessments.

Four distinct groups of students may receive accommodations on the summative assessments:

1. SWDs who have an IEP.
2. Students with a Section 504 plan who have a physical or mental impairment that substantially limits one or more major life activities, have a record of such an impairment, or are regarded as having such an impairment but who do not qualify for special education services.
3. Students who are ELs.
4. Students who are ELs with disabilities who have an IEP or 504 plan.

These students are eligible for accommodations intended for both SWDs and ELs. Testing accommodations for SWDs or students who are ELs must be documented according to the guidelines and requirements outlined in the *Accessibility Features and Accommodations Manual*.

3.2.2. Accessibility System

Through a combination of universal design principles and accessibility features, participating states and agencies designed an inclusive assessment system by considering accessibility from initial design and through item development, field testing, and implementation of the assessments for all students, including SWDs, ELs, and ELs with disabilities. Accommodations may still be needed for some SWDs and ELs to assist in demonstrating what they know and can do. However, the accessibility features available to students should minimize the need for accommodations during testing and ensure the inclusive, accessible, and fair testing of the diverse students being assessed.

3.2.3. What are Accessibility Features?

On computer-based assessments, accessibility features are tools or preferences that are either built into the assessment system or provided externally by Test Administrators and may be used by any student taking the summative assessments (i.e., students with and without disabilities, gifted students, ELs, and ELs with disabilities). Since accessibility features are intended for all students, they are not classified as accommodations. Students should have the opportunity to select and practice using the accessibility features prior to testing to determine which are appropriate for use on the assessment. Consideration should be given to the supports a student finds helpful and consistently uses during instruction. Practice tests that include accessibility features are available for teacher and student use throughout the year.

3.2.4. Accommodations for Students with Disabilities and English Learners

It is important to ensure that performance in the classroom and on assessments is influenced minimally, if at all, by a student's disability or linguistic/cultural characteristics that may be unrelated to the content being assessed. For the NJGPA, accommodations are considered adjustments to the testing conditions, test format, or test administration which provide equitable access during assessments for SWDs and students who are ELs. In general, the administration of the assessment should not be the first occasion on which an accommodation is introduced to the student. To the extent possible, accommodations should:

- Provide equitable access during instruction and assessments;
- Mitigate the effects of a student's disability;
- Not reduce learning or performance expectations;
- Not change the construct being assessed; and
- Not compromise the integrity or validity of the assessment.

Accommodations are intended to reduce and/or eliminate the effects of a student's disability and/or English language proficiency level; however, accommodations should never reduce learning expectations by reducing the scope, complexity, or rigor of an assessment. Moreover, accommodations provided to a student on the summative assessments must be generally

consistent with those provided for classroom instruction and classroom assessments. There are some accommodations that may be used for instruction and for formative assessments that are not allowed for a summative assessment because they impact the validity of the assessment results—for example, allowing a student to use a thesaurus or access the internet during an assessment. There may be consequences (e.g., excluding a student's test score) for the use of nonallowable accommodations during assessments. It is important for educators to become familiar with NJDOE policies regarding accommodations used for assessments.

To the extent possible, accommodations should adhere to the following principles:

- Accommodations enable students to participate more fully and fairly in instruction and assessments and demonstrate their knowledge and skills.
- Accommodations should be based upon an individual student's needs rather than on the category of a student's disability, level of English language proficiency alone, level of or access to grade-level instruction, amount of time spent in a general classroom, current program setting, or availability of staff.
- Accommodations should be based on a documented need in the instruction/assessment setting and should not be provided for the purpose of giving the student an enhancement that could be viewed as an unfair advantage.
- Accommodations for SWDs must be described and documented in the student's appropriate plan (i.e., either a 504 plan or an approved IEP) and must be provided if they are listed.
- Accommodations for ELs should be described and documented.
- Students who are ELs with disabilities are eligible to receive accommodations for both SWDs and ELs.
- Accommodations should become part of the student's program of daily instruction as soon as possible after completion and approval of the appropriate plan.
- Accommodations should not be introduced for the first time during the testing of a student.
- Accommodations should be monitored for effectiveness.
- Accommodations used for instruction should also be used, if allowable, on local district assessments and state assessments.
- In the following scenarios, the school must follow each state's policies and procedures for notifying the state assessment office:
 - A student was provided a test accommodation that was not listed in their IEP/504 plan/documentation for an English learner, or
 - A student was not provided a test accommodation that was listed in their IEP/504 plan/documentation for an English learner.

3.2.5. Unique Accommodations

A comprehensive list of accessibility features and accommodations designed to increase access to the summative assessments and that will result in valid, comparable assessment scores was provided in the *Accessibility Features and Accommodations Manual*. However, SWDs or ELs may require additional accommodations that are not already listed. Participating states and agencies individually review requests for unique accommodations in their respective states and provide a determination as to whether the accommodation would result in a valid score for the student, and if so, would approve the request.

3.2.6. Emergency Accommodations

An emergency accommodation may be appropriate for a student who incurs a temporary disabling condition that interferes with test performance shortly before or during the assessment window. A student, whether or not they already have an IEP or 504 plan, may require an accommodation as a result of a recently occurring accident or illness. Cases include a student who has a recently fractured limb (e.g., arm, wrist, or shoulder); a student whose only pair of eyeglasses has broken; or a student returning to school after a serious or prolonged illness or injury. An emergency accommodation should be given only if the accommodation will result in a valid score for the student (i.e., does not change the construct being measured by the test[s]). If the principal (or designee) determines that a student requires an emergency accommodation on the summative assessment, an Emergency Accommodation Form must be completed and maintained in the student's assessment file. The parent must be notified that an emergency accommodation was provided. If appropriate, the Emergency Accommodation Form may also be submitted to the District Assessment Coordinator to be retained in the student's central office file. Requests for emergency accommodations will be approved after it is determined that use of the accommodation would result in a valid score for the student.

3.2.7. Student Refusal Form

If a student refuses an accommodation listed in their IEP, 504 plan, or (if required by the member state) an EL plan, the school should document in writing that the student refused the accommodation, and the accommodation must be offered and remain available to the student during testing. This form must be completed and placed in the student's file and a copy must be sent to the parent on the day of refusal. Principals (or designee) should work with Test Administrators to determine who, if any others, should be informed when a student refuses an accommodation documented in an IEP, 504, or (if required by the member state) EL plan.

3.3. Testing Irregularities and Security Breaches

Any action that compromises test security or score validity is prohibited. These may be classified as testing irregularities or security breaches. Below are examples of activities that compromise test security or score validity (note that these lists are not exhaustive). It is highly recommended that School Test Coordinators discuss other possible testing irregularities and security breaches with Test Administrators during training.

Examples of test security breaches and irregularities include but are not limited to:

Electronic Devices:

- Using a cell phone or other prohibited handheld electronic device (e.g., smartphone, iPod, smart watch, personal scanner) while secure test materials are distributed, while students are testing, after a student turns in their test materials, or during a break. (*Exception:* Test Coordinators, Technology Coordinators, Test Administrators, and Proctors are permitted to use cell phones in the testing environment only in cases of emergencies or when timely administration assistance is needed; LEAs may set additional restrictions on allowable devices as needed.)

Test Supervision:

- Coaching students during testing, including giving students verbal or nonverbal cues, hints, suggestions, or paraphrasing or defining any part of the test.
- Engaging in activities (e.g., grading papers, reading a book, newspaper, or magazine) that prevent proper student supervision at all times while secure test materials are still distributed or while students are testing.
- Leaving students unattended for any period of time while secure test materials are distributed or while students are testing.
- Deviating from testing time procedures.
- Allowing cheating of any kind.
- Providing unauthorized persons with access to secure materials.
- Unlocking a test in PearsonAccess^{next} during nontesting times.
- Failing to provide a student with a documented accommodation or providing a student with an accommodation that is not documented and therefore not appropriate.
- Allowing students to test before or after the state's test administration window.

Test Materials:

- Losing a student test booklet or answer document.
- Losing a student testing ticket.
- Leaving test materials unattended or failing to keep test materials secure at all times.
- Reading or viewing the passages or test items before, during, or after testing (*Exception:* Administration of a human reader/signer accessibility feature for mathematics or accommodation for English language arts, which requires a Test Administrator to access passages or test items).
- Copying or reproducing (e.g., taking a picture of) any part of the passages or test items or any secure test materials or online test forms.
- Revealing or discussing passages or test items with anyone, including students and school staff, through verbal exchange, email, social media, or any other form of communication.
- Removing secure test materials from the school's campus or removing them from locked storage for any purpose other than administering the test.

Testing Environment:

- Allowing unauthorized visitors in the testing environment.
- Failing to follow administration directions exactly as specified in the *Test Administrator Manual*.
- Displaying testing aids in the testing environment (e.g., a bulletin board containing relevant instructional materials) during testing.

All instances of security breaches and testing irregularities must be reported to the School Test Coordinator immediately. The Form to Report a Testing Irregularity or Security Breach must be completed within two school days of the incident.

If any situation occurred that could cause any part of the test administration to be compromised,

schools should refer to the *Test Coordinator Manual* for each state's policy and immediately follow those steps. Instructions for the School Test Coordinator or LEA Test Coordinator to report a testing irregularity or security breach are available in the *Test Coordinator Manual*.

3.4. Data Forensics Analyses

Maintaining the validity of test scores is essential in any high-stakes assessment program, and misconduct represents a serious threat to test score validity. When used appropriately, data forensic analyses can serve as integral components of a wider test security protocol. The results of these data forensic analyses may be instrumental in identifying potential cases of misconduct for further follow-up and investigation.

The following data forensics analyses were conducted on the operational assessments:

- Response Change Analysis.
- Aberrant Response Analysis.
- Plagiarism Analysis.
- Longitudinal Performance Modeling.
- Internet and Social Media Monitoring.
- Off-Hours Testing Monitoring.

An overview of each data forensics analysis method is provided next.

3.4.1. Response Change Analysis

Response change analysis looks at how often student answers are changed, focusing specifically on an excessive number of wrong answers changed to right answers. In traditional paper-based, multiple-choice testing programs, this is sometimes referred to as “erasure analysis.” The rationale for erasure analysis is that a teacher or administrator who is intent on improving classroom performance might be motivated to change student responses after the answer sheets are collected. A clustered number of student answer documents from the same school or classroom with unusually high numbers of answers changed from wrong to right might provide evidence to support follow-up investigation. The response change analysis extended the traditional erasure method to account for issues specific to computer-based testing as well as the variety of item types on the summative assessments, such as partial-credit, multi-part, and multiple-select items.

3.4.2. Aberrant Response Analysis

Aberrant response pattern detection analysis looks at the unusualness of student responses compared with what would be expected. Most simply, this can be thought of as quantifying the extent to which higher-scoring students miss easy questions and lower-scoring students answer difficult questions correctly. While it would be difficult to draw a definitive inference about a single student flagged as having an aberrant response pattern, a cluster of students with aberrant response patterns within a classroom or school might warrant further investigation.

3.4.3. Plagiarism Analysis

Plagiarism analysis compares the responses given for a group of written composition items, looking for high degrees of similarity. For the summative assessments, the primary item type of interest was the prose constructed-response tasks in the English language arts content area. This analysis was conducted for prose constructed-response tasks administered online using some of the same artificial intelligence techniques that are applied in automated essay scoring. Specifically, this method was based on latent semantic analysis (LSA) technology to detect possible plagiarism. Using LSA, the content of each constructed response was compared against the content of every other constructed response, and a measure that indicated the degree of similarity was generated for each pair of response comparison. Because LSA provided a semantic representation of language, rather than a syntactic or word-based representation, it allowed the detection of potential copying behaviors, even when students or administrators substituted synonymous words or phrases.

3.4.4. Internet and Social Media Monitoring

Internet and social media monitoring were conducted by Caveon, LLC. Caveon's team monitored English-language websites and searchable forums that were publicly available for suspected proxy testing solicitations and website postings that contain, or appear to contain, infringements of protected operational test content.

The internet and social media outlets monitored included popular websites (such as Facebook and Twitter), blogs, discussion forums, video archives, document archives, brain dumps, auction sites, media outlets, peer-to-peer servers, and so on. Caveon's process generated regular updates that categorize identified threats by level of actual or potential risk based upon the representations made on the websites, or actual analysis of the proffered content. For example, categorizations typically ranged from "cleared" (lowest risk but bookmarked for continued monitoring) to "severe" (highest risk). Note that this process only considered potential breaches of secure item content, not violations of testing administration policies. Potential breaches were reported directly to the state(s) implicated for further action. Summary reports describing the threats were provided through notification emails.

3.4.5. Off-Hours Testing Monitoring

Off-hours testing monitoring is a process that checks for suspicious activity occurring outside computer-based testing windows. Participating states and agencies have set start and end times for administering computer-based assessments. Authorized users in the state role were allowed to override these start and end times. The off-hours testing monitoring tracked such overrides and logged them in an operational report. States could use this report to follow up with the organizations.

3.5. Quality Control of Test Administration

Pearson provided high-quality materials in a timely and efficient manner to meet the test administration needs. Since the majority of printing work was done in-house, it was possible to fully control the production environment, press schedule, and quality process for print materials.

Additionally, strict security requirements were employed to protect secure materials production. Section 3 provides details on the secure handling of test materials. Materials were produced according to the style guide and to the detailed specifications supplied in the materials list.

Pearson Print Service operates within the sanctions of an ISO 9001:2008 Quality Management System, and practices process improvement through lean principles and employee involvement.

Raw materials (paper and ink) used for scannable forms production were manufactured exclusively for Pearson Print Service using specifications created by Pearson Print Service. Samples of ink and paper were tested by Pearson prior to use in production. Project specialists were the point of contact for incoming production.

Purchase orders and other order information were assessed against manufacturing capabilities and assigned to the optimal production methodology. Expectations, quality requirements, and cost considerations were foremost in these decisions. Prior to release for manufacture, order information was checked against specifications, technical requirements, and other communication that includes expected outcomes. Records of these checks were maintained.

Files for image creation flow through one of two file preparation functions: digital pre-press for digital print methodology, or plateroom for offset print methodology. Both the digital prepress and plateroom functions verify content, file naming, imposition, pagination, numbering stream, registration of technical components, color mapping, workflow, and file integrity. Records of these checks are created and saved.

Offset production requires printing that uses a lithographic process. Offline finishing activities are required to create books and package offset output. Digital output may flow through an inkjet digital production line or a sheet-fed toner application process in the Xpress Center. A battery of quality checks was performed in these areas. The checks included color match, correct file selection, content match to proof, litho-code to serial number synchronization, registration of technical components, ink density controlled by densitometry, inspection for print flaws, perforations, punching, pagination, scanning requirements, and any unique features specified for the order. Records of these checks and samples pulled from planned production points were maintained. Offline finishing included cutting, shrink-wrapping, folding, and collating. The collation process has three robust inline detection systems that inspected each book for:

- Caliper validation that detects too few or too many pages. This detector will stop the collator if an incorrect caliper reading is registered.
- An optical reader that will only accept one sheet. Two or zero sheets will result in a collator stoppage.
- The correct bar code for the signature being assembled. An incorrect or upside-down signature will be rejected by the bar code scanner and will result in a collator stoppage.

Pearson's Quality Assurance department personnel inspected print output prior to collation and shipment. Quality Assurance also supported process improvement, work area documentation, audited process adherence, and established training programs for employees.

Chapter 4. Item Scoring

4.1. Machine-scored Items

4.1.1. Key-based Items

Pearson performed a key review prior to the test administration to verify that the scoring (answer) keys were correct for each item. Once the forms were constructed and approved for publication, an independent key review was performed by an experienced third-party vendor. The vendor reviewed each item and confirmed that the key was correct. If discrepancies were identified, a Pearson senior content specialist or content manager reviewed the flagged item(s) and worked with the item developers to resolve the issue.

4.1.2. Rule-based Items

Rule-based scoring refers to item types that use various scoring models. Participating states and agencies use Question and Test Interoperability item type implementation based on scoring model rules. Examples of these item types include choice interaction, which presents a set of choices where one or more choices can be selected; text entry, where the response is entered in a text box; hot spot or text interaction, where an area in a graph or text in a paragraph (for example) can be highlighted; or match interaction, where an association can be made between pairs of choices in a set. These items include the scoring rules and correct responses as part of their item XML (markup language) coding. This subsection of the technical report discusses the creation and initial implementation for the scoring of rule-based items, which continue to be used in all related programs, including the NJGPA. The final two paragraphs of this subsection discuss practices undertaken each year to score rule-based items on the NJGPA.

During the initial stages of item development for the PARCC program, Pearson staff worked closely with participating states and agencies to first delineate the rules for the scoring rubrics and then adjust those rules based on student responses. During item studies in spring 2015, Pearson content staff received input from the staff of participating states and agencies to develop a thorough rule-based scoring process that met their needs.

Pearson worked with the item developers to review initial scoring rules created during item development. Once the rule-based scoring process was approved, and prior to test construction, Pearson content staff worked closely with the item developers to finalize scoring rubrics for items to be scored via the rule-based scoring method. The proposed scoring rubrics were sent for review, and if any additional changes were needed or new rules added, Pearson documented and applied the requested edits.

During initial test construction, Pearson monitored and evaluated the scoring and updated the scoring keys/scoring rules in the item bank. After the tryout items were scored, Pearson prepared a frequency distribution of student responses for each scored item or task using a rule-based approach and compared this to the expected response based on correct answers to ensure that

scoring keys and rules were appropriately applied. The content team analyzed the student response data to determine if scoring was acceptable using the item metadata and the student response file in conjunction with any potential item issues as flagged by psychometrics. These frequency distributions included an indication of right/wrong and other identifying information defined by participating states and agencies, and those items that showed a statistical anomaly, whereby the frequency distribution was outside of the expected range, were sent to content experts to verify that the items were coded with the correct key.

Following the Rule-Based Scoring Educator Committee's review, which occurred prior to year one test construction, Pearson analyzed the feedback from the committees and made recommendations about adjustments to the scoring rubrics based on the results of the reviews. Upon submission of the results, Pearson worked with the staff of participating states and agencies to discuss these findings and determine next steps prior to the completion of scoring. In subsequent years as scoring inquiries arise throughout the process of test construction, forms creation, testing, scoring and psychometric analysis, items with scoring discrepancies are brought before the Priority Alert Task Force for resolution. This committee consists of representatives from each state as well as the content specialists at participating states and agencies and Pearson.

Following the initial development of the rule-based scoring rubrics, Pearson has continued to monitor and evaluate new item development to ensure the scoring rules established are maintained within all item types as approved.

Pearson continues to use several avenues to monitor scoring each year. During testing, Pearson creates early testing files for frequency distribution analysis whereby items for which an incorrect key receives a high distribution of responses are further evaluated for accuracy. After testing, all responses are again evaluated for the distribution of responses and potential scoring abnormalities during psychometric analysis. These processes are the same for both paper and online modes of testing.

4.2. Human or Hand-scored Items

Constructed-response items were scored by human scorers in a process referred to as hand scoring. Online training units were used to train all scorers. The online training units included prompts (items), passages, rubrics, training sets, and qualification sets. Scorers who successfully completed the training and demonstrated they could correctly score student responses based on the guidelines in the online training units were permitted to score student responses using the ePEN2 (Electronic Performance Evaluation Network, second generation) scoring platform. All online and paper responses were scored within the ePEN2 system. Pearson monitored quality throughout scoring.

Pearson staff roles and responsibilities were as follows:

- Scorers applied scores to student responses.
- Scoring supervisors monitored the work of a team of scorers through review of scorer statistics and backreading, which is a review of responses scored by each scorer. When backreading, a supervisor sees the scores applied by scorers, which helps the supervisor provide additional coaching or instruction to the scorer being backread.

- Scoring directors managed the scoring quality of a subset of items and monitored the work of supervisors and scorers for their assigned items. Directors backread responses scored by supervisors and scorers as part of their quality-monitoring duties.
- English language arts (ELA) and mathematics content specialists managed the scoring quality and monitored the work of the scoring directors.
- The project manager documented the procedures, identified risks, and managed day-to-day administrative matters.
- A portfolio manager provided oversight for the entire scoring process.

All Pearson employees involved in the scoring or the supervision of scoring possessed at least a four-year college degree.

4.2.1. Scorer Training

Key steps in the development of scorer training materials were range-finding and rangefinder review meetings where educators and administrators from states met to interpret the scoring rubrics and determine consensus scores for student responses. Range-finding meetings were held prior to scoring field-test items, and rangefinder review meetings were held prior to scoring operational items.

At range-finding meetings, educators and administrators from states reviewed student responses and used scoring rubrics to determine consensus scores. Those responses scored in range-finding were used to create field-test scorer training sets. After items were selected for operational testing, educators and administrators attended rangefinder review meetings to review and approve proposed operational scorer training sets.

When developing scorer training materials, Pearson scoring directors carefully reviewed detailed notes and records from range-finding and rangefinder review committee meetings. Training sets were developed using the responses scored by the committees and additional suitable student response samples (as needed). All scorer training sets were reviewed and approved prior to scorer training.

During training, scorers reviewed training sets of scored student responses with annotations that explained the rationale for the score assigned. The anchor set was the primary reference for scorers as they internalized the rubric during training. Each anchor set consisted of responses that were clear examples of student performance at each score point. The responses selected were representative of typical approaches to the task and arranged to reflect a continuum of performance. All scorers had access to the anchor set when they were training and scoring and were directed to refer to it regularly during scoring.

Practice sets were used in training to help trainees practice applying the scoring guidelines. Scorers reviewed the anchor sets, scored the practice sets, and then were able to compare their assigned scores for the practice sets to the actual assigned scores to help them learn.

Qualification sets were used to confirm that scorers understood how to score student responses accurately. Qualification sets were composed of responses that were clear examples of score

points. Scorers were required to meet specified agreement percentages on qualification sets in order to score student responses.

Pearson has developed two types of training sets to train scorers: prototype and abbreviated sets. Prototype training sets were complete training sets consisting of anchor, practice, and qualification sets (refer to 4.2.2 for information on the qualification process). In ELA, there was one prototype training set per task type (Research Simulation Task and Literary Analysis Task). In mathematics, a prototype training set was built for a grouping of similar items for a total of approximately three to four prototype sets per grade level or course.

The prototype training approach promoted consistency in scoring, as each subsequent abbreviated training set for the ELA task type or mathematics item grouping was based on the prototype. Once a prototype was chosen, full training materials were developed for that item. Then, scorers at each grade level were trained to score a particular item type using the prototype training materials for that type.

Abbreviated training sets were prepared for all items not selected for prototype training sets. The abbreviated training sets included an anchor set and two practice sets so scorers could internalize the scoring standards for these new items, which were similar to prototype items they had previously scored.

Anchor and practice sets for both prototype and abbreviated items included annotations for each response. Annotations are formal written explanations of the score for each student response.

Table 4.1 details the composition of the anchor sets, practice sets, and qualification sets.

Table 4.1 Training Materials Used During Scoring

Training Set Development	
Description	Specification
Anchor Set	
The anchor set is the primary reference for scorers as they internalize the rubric during training. All scorers have access to the anchor set when they are training and scoring and are directed to refer to it regularly.	The anchor set for mathematics prototype items comprises three annotated responses per score point. The anchor set for subsequent abbreviated items for mathematics comprise one to three annotated responses per score point.

The anchor set comprises clear examples of student performance at each score point. The responses selected may be representative of typical approaches to the task or arranged to reflect a continuum of performance.	The anchor sets for ELA prototype items comprise three annotated responses per score point. Anchor sets for prototype items include separate complete anchor sets for each applicable scoring trait (Reading Comprehension and Written Expression, and Knowledge of Language and Conventions for Research Simulation and Literary Analysis Tasks).
Practice Sets	
Practice sets are used to help trainees develop experience in independently applying the scoring guide (the rubric) to student responses. Some of these responses clearly reinforce the scoring guidelines presented in the anchor set. Other responses are selected because they are more difficult to evaluate, fall near the boundary between two score categories, or represent unusual approaches to the task. The practice sets provide guidance and practice for trainees in defining the line between score categories, as well as applying the scoring criteria to a wider range of types of responses.	The practice sets for mathematics prototype and abbreviated items include two to three sets of ten annotated responses.
	ELA practice sets for prototype items include two sets of five annotated responses and two sets of 10 annotated responses. The subsequent ELA practice sets for abbreviated items include two sets of ten annotated responses.
Training Set Development	
Description	Specification
Qualification Sets	
Qualification sets are used to confirm that scorer trainees understand the scoring criteria and are able to assign scores to student responses accurately. The responses in these sets are selected to reinforce the application of the scoring criteria illustrated in the anchor set.	The qualification sets for mathematics prototype items include three sets of 10 responses each (not annotated). The subsequent mathematics abbreviated items for mathematics do not include qualification sets.

Scorer trainees must demonstrate acceptable performance on these sets by meeting a predetermined standard for accuracy in order to qualify to score. Pearson scoring staff defined and documented qualifying standards in conjunction with participating states and agencies prior to scoring.	The qualification sets for ELA prototype items include three sets of 10 responses each (not annotated). The subsequent ELA abbreviated items do not include qualification sets.
---	---

4.2.2. Scorer Qualification

In order to score items, scorers were required to show that they were able to apply scoring methodology accurately through a qualification process. Scorers were asked to apply scores to three qualification sets consisting of 10 responses each. ELA scorers applied a score for each trait on each response in the qualification sets. Literary Analysis and Research Simulation Tasks each had two traits: the Reading Comprehension and Written Expression trait and the Conventions trait. Mathematics scorers applied a score for each part of an item that was a constructed response. The number of constructed-response parts for each mathematics item ranged from one to four. Scorers were required to match the approved score at a percentage agreed to by participating states and agencies in order to qualify.

For ELA qualification, scorers were required to meet the following three conditions:

1. On at least one of the three qualifying sets, at least 70 percent of the ratings on each of the two scoring traits (considered separately) must agree exactly with the approved scores.
2. On at least two of the three qualifying sets, at least 70 percent of the ratings (combined across the three scoring traits) must agree exactly with the approved scores.
3. Combining over the three qualifying sets and across the two scoring traits, at least 96 percent of the ratings must be within one point of the approved scores.

For mathematics qualification, the requirements were based on the item type and score point range. Because mathematics items can have one or more scoring traits, a scorer needed to achieve the requirements as set forth in Table 4.2 separately for each scoring trait (when applicable to the item).

Table 4.2 Mathematics Qualification Requirements

Category	Score Point Range	Perfect Agreement	Within One Point
2	0–1	90%	100%
3	0–2	80%	96%
4	0–3	70%	96%
5	0–4	70%	95%
6	0–5	70%	95%
7	0–6	70%	95%

On at least two of the three qualifying sets, a scorer was required to meet the “perfect agreement” percentage indicated in the table above for each category. “Perfect agreement” was achieved when the scores applied exactly matched the approved scores. Over the three qualifying sets, a scorer was required to meet the “within one point” percentage indicated in the table above for each category. The average is exclusive to each trait, so an item with multiple scoring traits would have multiple trait rating averages within one point of the approved score.

4.2.3. Managing Scoring

Pearson created a hand-scoring specifications document that detailed the hand scoring schedule, customer requirements, range-finding plans, quality management plans, item information, and staffing plans for each scoring administration.

4.2.4. Monitoring Scoring

4.2.4.1. Second Scoring

During scoring, Pearson’s ePEN2 scoring system automatically and randomly distributed a minimum of 10 percent of student responses for second scoring; scorers had no indication whether a response had been scored previously. Humans applied the second score for all mathematics items. Second scoring for ELA was performed either by human scorers or by Pearson’s Intelligent Essay Assessor. If the first and second scores applied were nonadjacent, a third and occasionally a fourth score were assigned to resolve scorer disagreements. When a resolution score (i.e., third score) was nonadjacent to one or both of the first and second scores, the content specialist or scoring director would apply an adjudication score (fourth score). If a response was scored more than once, the rules in Table 4.3 were applied to determine the final score.

Table 4.3 Scoring Hierarchy Rules

Score Type	Rank	Final Score Calculation
Adjudication	1	If an adjudication score is assigned, this is the final score.
Resolution	2	If no adjudication score is assigned, this is the final score.
Backread	3	If no adjudication or resolution score is assigned, the latest backreading score is the final score.
Human First Score	4	If no adjudication, resolution, or backreading score is assigned, this is the final score.
Human Second Score	5	If no adjudication, resolution, backreading, or human first score is assigned, this is the final score.
Intelligent Essay Assessor Score	6	If no human score is assigned, this is the final score.

4.2.4.2. Backreading

Backreading was one of the major responsibilities of Pearson Scoring Supervisors and a primary tool for proactively guarding against scorer drift, where scorers score responses in comparison to one another instead of in comparison to the training responses. Scoring supervisory staff used the ePEN2 backreading tool to review scores assigned to individual student responses by any given scorer in order to confirm that the scores were correctly assigned and to give feedback and remediation to individual scorers. Pearson backread approximately 5 percent of the hand-scored responses. Backreading scores did not override the original score but were used to monitor scorer performance.

4.2.4.3. Validity

Validity responses are prescored responses strategically interspersed in the pool of live responses. These responses were not distinguishable from any other responses so that scorers were not aware they were scoring validity responses rather than live responses. The use of validity responses provided an objective measure that helped ensure that scorers were applying the same standards throughout the project. In addition, validity was at times shared with scorers in a process known as “validity as review.” Validity as review provided scorers automated, immediate feedback: a chance to review responses they mis-scored, with reference to the correct score and a brief explanation of that score. One validity response was sent to scorers for every 25 “live” responses scored.

Validity agreement requirements for scorers are listed in Table 4.4. Scorers had to meet the required validity agreement percentages to continue working on the project. Scorers who did not maintain expected agreement statistics were given a series of interventions culminating in a targeted calibration set: a test of scorer knowledge. Scorers who did not pass targeted calibration were removed from scoring the item, and all the scores they assigned were deleted.

Table 4.4 Scoring Validity Agreement Requirements

Subject	Score Point Range	Perfect Agreement	Within One Point
Mathematics	0–1	90%	96%
Mathematics	0–2	80%	96%
Mathematics	0–3	70%	96%
Mathematics	0–4	65%	95%
Mathematics	0–5	65%	95%
Mathematics	0–6	65%	95%
ELA	Multi-trait	65%	96%

*A numerical score compared to a blank or condition code score will have a disagreement greater than 1 point.

4.2.4.4. Calibration Sets

Calibration sets are special sets created during scoring to help train scorers on particular areas of concern or focus. Scoring directors used calibration sets to reinforce range-finding standards, introduce scoring decisions, or address scoring issues and trends. Calibration was used either to correct a scoring issue or trend, or to continue scorer training by introducing a scoring decision. Calibration was administered regularly throughout scoring.

4.2.4.5. Inter-rater Agreement

Inter-rater agreement is the agreement between the first and second scores assigned to student responses and is the measure of how often scorers agree with each other. Pearson scoring staff used inter-rater agreement statistics as one factor in determining the needs for continuing training and intervention on both individual and group levels. Inter-rater agreement expectations are shown in Table 4.5.

Table 4.5 Inter-Rater Agreement Expectations and Results

Subject	Score Point Range	Perfect Agreement Expectation	Perfect Agreement Result	Within One Point Expectation*	Within One Point Result
Mathematics	0–1	90%	98%	96%	100%
Mathematics	0–2	80%	97%	96%	100%
Mathematics	0–3	70%	96%	96%	99%
Mathematics	0–4	65%	94%	95%	99%
Mathematics	0–5	65%	91%	95%	98%
Mathematics	0–6	65%	95%	95%	98%
ELA	Multi-trait	65%	83%	96%	100%

*A numerical score compared to a blank or condition code score will have a disagreement greater than 1 point.

Pearson's ePEN2 scoring system included comprehensive inter-rater agreement reports that

allowed supervisory personnel to monitor both individual and group performance. Based on reviews of these reports, scoring experts targeted individuals for increased backreading and feedback, and if necessary, retraining.

The perfect agreement rate for mathematics responses scored by two scorers ranged from 65 to 98 percent and the within-one-point rate ranged from 95 to 100 percent. For all ELA responses scored by two scorers, the perfect agreement rate ranged from 83 percent to 100 percent and the within-one-point rate ranged from 96 percent to 100 percent.

The results for the ELAGP are provided in Section 4.3.7, “Inter-rater Agreement for Prose Constructed Response.”

4.3. Automated Scoring for PCRs

Automated scoring performed by Pearson’s Intelligent Essay Assessor (IEA) was the default option for scoring the summative assessment’s online prose constructed response (PCR) tasks. Under the default option, it was assumed that operational scores for approximately 90 percent of the online PCR responses would be assigned by IEA for the spring administration. The operational scores for the remaining online responses were assigned by human scorers. Human scoring was applied to responses that were scored while IEA was being trained as well as to additional responses routed to human scoring when there was uncertainty about the automated scores.

For 10 percent of responses, a second “reliability” score was assigned. The purpose of the reliability score was to provide data for evaluating the consistency of scoring, which is done by evaluating scoring agreement. When IEA provided the first score of record, the second reliability score was a human score.

4.3.1. Concepts Related to Automated Scoring

The sections below describe concepts related to automated scoring.

4.3.1.1. Continuous Flow

Continuous flow scoring results in an integrated connection between human scoring and automated scoring. It refers to a system of scoring where either an automated score, a human score, or both can be assigned based on a predetermined asynchronous operational flow.

4.3.1.2. Training of IEA Using Operational Data

Continuous flow scoring facilitates the training of IEA using human scores assigned to operational online data collected early in the administration. Once IEA obtains sufficient data to train, it can be “turned on” and becomes the primary source of scoring (although human scoring continues for the 10 percent reliability sample and other responses that may be routed accordingly).

4.3.1.3. Smart Routing

Smart routing refers to the practice of using automated scoring results to detect responses that are likely to be challenging to score and applying automated routing rules to obtain one or more additional human scores. Smart routing can be applied prompt by prompt to the extent needed to meet scoring quality criteria for automated scoring.

4.3.1.4. Quality Criteria for Evaluating Automated Scoring

The state leads approved specific quality criteria for evaluating automated scoring. The primary evaluation criteria for IEA was based on responses to validity papers with “known” scores assigned by experts. For each prompt scored, a set of validity papers is used to monitor the human-scoring process over time. Validity papers are seeded into human scoring throughout the administration. The expectation is that IEA can score validity papers at least as accurately as humans can.

Additional measures of inter-rater agreement for evaluating automated scoring were proposed based on the research literature (Williamson et al., 2012). These measures were previously utilized in Pearson’s automated scoring research and include Pearson correlation, kappa, quadratic-weighted kappa, exact agreement, and standardized mean difference. These measures are computed between pairs of human scores, as well as between IEA and humans, to evaluate how performance was the same or different. Criteria for evaluating the training of IEA given these measures include the following:

- Pearson correlation between IEA-human should be within 0.1 of human-human.
- Kappa between IEA-human should be within 0.1 of human-human.
- Quadratic-weighted kappa between IEA-human should be within 0.1 of human-human.
- Exact agreement between IEA-human should be within 5.25 percent of human-human.
- Standardized mean difference between IEA-human should be less than 0.15.

The specific criteria for evaluating IEA included both primary and secondary criteria and are noted below:

- Primary Criteria — Based on responses to validity papers: With smart routing applied as needed, IEA agreement is as good as or better than human agreement for each trait score.
- Contingent Primary Criteria — Based on the training responses if validity responses are not available: With smart routing applied as needed, IEA-human exact agreement is within 5.25 percent of human-human exact agreement for each trait score.
- Secondary Criteria — Based on the training responses: With smart routing applied as needed, IEA- human differences on statistical measures for each trait score are within the Williamson et al. (2012) tolerances for subgroups with at least 50 responses.

4.3.1.5. Hierarchy of Assigned Scores for Reporting

When multiple scores are assigned for a given response, the following hierarchy determines which score was reported operationally:

- The IEA score is reported if it is the only score assigned.

- If an IEA score and a human score are assigned, the human score is reported.
- If a first human score and a second human score are assigned, the first human score is reported.
- If a backread score and human and/or IEA scores are assigned, the backread score is reported if there is no resolution or adjudication score assigned.
- If a resolution score is assigned and an adjudicated score is not assigned, the resolution score is reported (note that if nonadjacent scores are encountered, responses are automatically routed to resolution).
- If an adjudicated score is assigned, it is reported (note that if a resolution score is nonadjacent to the other scores assigned, responses are automatically routed to adjudication).

4.3.2. Sampling Responses Used for Training IEA

For prompts trained using 2022 operational data, the early performance of human scoring was closely monitored to verify that an appropriate set of data would be available for training IEA. In particular, several characteristics of the human scoring data were monitored, including:

- Exact agreement between human scorers (the goal was for this to be at least 65 percent for each trait).
- Exact agreement between human scores conditioned on score point (the goal was for this to be at least 50 percent for each trait).
- The number of responses at each score point (the goal was to have at least 40 responses at the highest score points in the training samples used by IEA).
- The number of responses with two human scores assigned (note that IEA “ordered” additional scoring of responses during the sampling period as needed).

Although the desired characteristics of the training data were easily achieved for some prompts, they were more challenging to achieve for others. For some prompts, a subset of scores were reset and clarifying directions were provided to scorers to improve human-human agreement. For other prompts, special sampling approaches were used to increase the numbers of responses that received top scores. In addition, a healthy percentage of responses were backread during the sampling period and these scores as well as double human scores were all part of the data used to train IEA.

4.3.3. Primary Criteria for Evaluating IEA Performance

The primary criteria for evaluating IEA performance are based on evaluating validity papers and is stated as follows: With smart routing applied as needed, IEA agreement is as good as or better than human agreement for each trait score.

To operationalize the primary criteria for a given prompt, the following general steps are undertaken:

1. Determine agreement of the human scores with the validity papers for each trait.
2. Calculate agreement of the IEA scores with the validity papers for each trait.
3. Compare the IEA validity agreement with the human agreement.
4. If the IEA validity agreement is greater than or equal to the human agreement for each trait, IEA can be deployed operationally.

In addition to looking at overall validity agreement, conditional agreement was also examined. In general, it was desirable for IEA to exceed 65 percent agreement at every score point as well as be close to or exceed the human validity agreement at each score point.

4.3.4. Contingent Primary Criteria for Evaluating IEA Performance

For many of the prompts trained in 2022, it was not possible to utilize human-scored validity responses in evaluating IEA performance. In these cases, IEA was evaluated based on IEA-human exact agreement for each trait score and compared to agreement based on responses that were double-scored by humans. A portion of the data was held out for evaluating IEA-human exact agreement according to the following steps:

1. Determine exact agreement of the two human scores with each other for each trait.
2. Calculate agreement of the IEA scores with the human scores for each trait.
3. Compare the IEA-human agreement with the human-human agreement.
4. If the IEA-human agreement is within 5.25 percent of the human-human agreement, IEA can be deployed operationally.

In addition to the overall comparison, the following performance thresholds were targeted in the test data set: (1) at least 65 percent overall IEA-human agreement; and (2) 50 percent IEA-human agreement by score point (i.e., conditioned on the human score). These targets went beyond the contingent primary criteria approved by the state leads.

4.3.5. Applying Smart Routing

With smart routing, the quality of automated scoring can be increased by routing responses that are more likely to disagree with a human score to receive an additional human score.

When human scorers read a paper, they typically apply integer scores based on a scoring rubric. When there is strong agreement between two independent human readers, the readers might both assign a score of 3 such that the average score over both raters is also a 3 (i.e., $(3+3)/2 = 3$). IEA simulates this behavior, but because its scores come from an artificial intelligence algorithm, it generates continuous (i.e., decimalized) scores. In this case, the IEA score might be a 2.9 or 3.1. When human readers disagree on the score for a paper, say one reader gives the paper a score of 3 and another reader gives the paper a score of 4, the average of the two scores would be 3.5 (i.e., $3+4=7/2=3.5$). For this paper, IEA would likely provide a score between 3 and 4, say 3.4 or 3.6. Because this continuous score needs to be rounded to an integer score for reporting, it might be reported as a 3 or a 4, depending on the rounding rules. Smart routing involves routing those responses with “in between” IEA scores to additional human scoring because the nature of the responses suggests there may be less confidence in the IEA score. Since these “in between” IEA scores are based on modeling human scores, it follows that human scores may be less certain as well, and thus such responses tend to be the ones that it makes sense to have double-scored and possibly to resolve if the IEA and human scores are nonadjacent.

Smart routing was utilized as needed to help IEA achieve targeted quality metrics (e.g., validity agreement or agreement with human scorers). Smart routing involved the application of the following four steps:

1. The continuous IEA score for each of the two trait scores was rounded to the nearest score interval of 0.2, starting from zero. For example, IEA scores between 0 and 0.1 were rounded to an interval score of 0, scores between 0.1 and 0.3 were rounded to an interval score of 0.2, scores between 0.3 and 0.5 were rounded to an interval score of 0.4, and so on.
2. Within each of these intervals, the percentage of exact agreement between IEA integer scores and the human scores was calculated for each trait.
3. For each prompt, agreement rates were evaluated by rounding interval. Those intervals for which the agreement rates were below a designated threshold for either trait were identified.
4. Once IEA scoring was implemented, responses within intervals for which IEA-human agreement was below the designated threshold were routed for additional human scoring.

In training IEA, the scoring models without smart routing were evaluated first by applying either the primary validity criteria or the contingent criteria as described in Section 4.3.4. For those prompts that did not meet these criteria, increasing smart routing thresholds were applied in an iterative fashion to filter scores and evaluate the remaining scores against the criteria. That is, in any one iteration a particular smart routing threshold was applied such that only scores falling in intervals for which exact agreement exceeded the threshold were included in evaluating the criteria. If the primary or contingent criteria were not met with this level of smart routing, an increased smart routing threshold was applied iteratively until the primary or contingent criteria were met, or the maximum threshold reached. If the criteria were still not met after a maximum threshold was applied, different models were investigated and/or additional human scoring data utilized until an IEA scoring model was found that met the criteria.

4.3.6. Evaluation of Secondary Criteria for Evaluating IEA Performance

The secondary criteria for evaluating IEA performance involved comparing agreement indices for IEA-human scoring for various demographic subgroups. Because of the importance of protecting personally identifiable information, student demographic data is stored and managed separately from the performance scoring data. For this reason, it was not possible to evaluate subgroup performance in real time as IEA was being trained.

For those prompts trained on early operational data, attempts were made to prioritize the data being returned from the field to include data from states or districts where more diverse populations of students were anticipated. In addition, requests for additional human scores were made to increase the likelihood that there would be sufficient numbers of responses with two human scores for most of the demographic subgroups of interest.

Once IEA was trained and deployed, scoring sets used in training were matched to demographic information so that agreement between IEA and human scorers could be evaluated across subgroups. The analysis was conducted for the eleven comparison groups outlined in Table 4.6.

Table 4.6 Comparison Groups

Group Type	Comparison Groups
Sex	Female Male
Ethnicity/Race	Hispanic/Latino* American Indian/Alaska Native Asian Black/African American Native Hawaiian or Other Pacific Islander Two or More Races White
Special Instructional Needs	English Language Learners (ELL) Students with Disabilities (SWD)

*Hispanic/Latino is an ethnic categorization based on cultural affiliation. Students may identify as Hispanic and Latino and as one of more of the other racial categories listed.

IEA-human agreement indices were calculated for all cases with an IEA score and at least one human score. Human-human agreement was calculated for all cases with two human scores.

To evaluate the training of IEA for subgroups, the following criteria approved by the state leads for subgroups with at least 50 IEA-human scores and at least 50 human-human scores were applied:

- Pearson correlation between IEA-human should be within 0.1 of human-human.
- Kappa between IEA-human should be within 0.1 of human-human.
- Quadratic-weighted kappa between IEA-human should be within 0.1 of human-human.
- Exact agreement between IEA-human should be within 5.25 percent of human-human.
- Standardized mean difference between IEA-human should be less than ± 0.15 (this criterion was applied to subgroups with at least 50 IEA-human scores).

Although it was not expected that these criteria would be met for all subgroups for all prompts, if results of the evaluation between IEA and human scoring for subgroups for any prompt indicated that IEA performance persistently failed on the criteria listed above, consideration would be given to resetting the responses scored by IEA and reverting to human scoring until such time that an alternate IEA model could be established with improved subgroup performance.

In addition to the secondary criteria approved by the State Leads, the performance of IEA was compared to the following targets on the various measures for subgroups with at least 50 responses:

- Pearson correlation between IEA-human should be 0.70 or above.
- Kappa between IEA-human should be 0.40 or above.
- Quadratic-weighted kappa between IEA-human should be 0.70 or above.
- Exact agreement between IEA-human should be 65 percent or above.

These targets were not intended to be directly applied in decisions about whether to deploy IEA

operationally or not. Such targets may or may not be met by human scoring for any particular prompt and/or subgroup, and if they are not met by human scoring, they are unlikely to be met by IEA scoring. Nevertheless, comparisons to these targets provided additional information about IEA performance (and human scoring) in an absolute sense.

4.3.7. Inter-rater Agreement for Prose Constructed Response

This section presents the inter-rater agreement for operational results for the online PCR tasks by trait. PCR items are scored on two traits: (1) Reading Comprehension and Written Expression and (2) Knowledge of Language and Conventions for Research Simulation for Literary Analysis Tasks.

For 10 percent of responses, a second “reliability” score was assigned. The purpose of the reliability score is to provide data for evaluating the consistency of scoring, which is done by evaluating scoring agreement. Inter-rater agreement is the agreement between the first and second scores assigned to student responses and is the measure of how often scorers agree with each other. Pearson scoring staff used inter-rater agreement indices as one factor in determining the needs for continuing training and intervention on both individual and group levels. Inter-rater agreement expectations are provided in Table 4.5 in Section 4.2.4. For ELA PCR traits, the expectation for agreement is an inter-rater agreement of 65 percent or higher between two scorers. When IEA provided the first score of record, the second reliability score was a human score. For a subset of responses, the first and second score were both human scores.

Table 4.7 presents the average agreement across the PCRs for the ELAGP by trait. The number of prompts included in the analyses is listed. The agreement indices (exact agreement, kappa, quadratic-weighted kappa, and Pearson correlation) were calculated separately by PCR for each trait (Reading Comprehension and Written Expression or Knowledge of Language and Conventions). The agreement indices were averaged across the PCRs. Table 4.7 presents the average count and the average for the agreement indices.

The exact agreement for the PCR traits is above the criteria of a 65 percent agreement rate for all PCRs. The strength of agreement between raters is moderate to substantial agreement as defined by Landis and Koch (1977) for all PCRs. The quadratic-weighted kappa (QW Kappa) distinguishes between differences in ratings that are close to each other versus larger differences. The weighted kappa is substantial to almost perfect agreement. The Pearson correlations (*r*) ranged from 0.94 to 0.95.

Table 4.7 Prose Constructed Response Average Agreement Indices by Test

Test	Number of PCRs	Count	Written Expression				Conventions			
			Exact	Kappa	QW Kappa	<i>r</i>	Exact	Kappa	QW Kappa	<i>r</i>
ELAGP	4	1155	89.5	0.85	0.95	0.95	88.83	0.84	0.94	0.94

4.4. Quality Control of Scoring

Quality control of answer document processing and scoring involves all aspects of the scoring

procedures, including key-based and rule-based machine scoring and hand scoring for constructed-response items and performance tasks.

For the 2015 operational administration, Pearson's validation team prepared test plans used throughout the scoring process. Test plan preparation was organized around detailed specifications.

Based on lessons learned from previous administrations, the following quality steps were implemented:

- Raw score validation (e.g., score key validation; evidence statement; field-test non-score; double-grid combinations; possible correct combination, if applicable; out-of-range/negative test cases).
- Matching (e.g., validation of high-confidence criteria, low-confidence criteria, cross document, external or forced matching by customer; prior to and after data updates; extract file of matched and unmatched documents).
- Demographic update tests (e.g., verification of data extract against corresponding layout; valid values for updatable fields; invalid values for updatable/non-updatable fields; negative test for nonexistent record or empty file).

The following components were added to the quality control process specifically for the program. These additional steps were introduced to address issues with item-level scoring that were identified in the 2014 field-test administration:

- XML Validation: A combination of automated validation against 100 percent of item XMLs and human inspection of XML from selected difficult item types or composite items.
- Administration/End-to-End Data Validation: An automated generation of response data from approved test maps that have known conditions against the operational scoring systems and data generation systems to verify scoring accuracy.
- Psychometric Validation: Verification of data integrity using criteria typically used in psychometric processes (e.g., statistical key checks) and categorization of identified issues to help inform investigation by other groups.
- Content Validation: An examination, by subject matter experts, of all items using a combination of automated tools to generate response and scoring data.

In addition to the steps described above, the following quality control process for answer keys and scoring that was implemented for the first operational administration was as follows:

- Pearson's psychometrics team conducted empirical analyses based on preliminary data files and flagged items based on statistical criteria.
- Pearson's content team reviewed the flagged items and provided feedback on the accuracy of content, answer keys, and scoring.
- Items potentially requiring changes were added to the product validation log for further investigation by other Pearson teams.
- Staff was notified of items for which keys or scoring changes were recommended.
- Participating states and agencies approved or rejected scoring changes.
- All approved scoring changes were implemented and validated prior to the generation of the data files used for psychometric processing.

Chapter 5. Performance Standards and Standards Validation

The score scales and performance standards for ELAGP and MATGP are based on score scales created and maintained for the PARCC/NMC Affiliate program. The ELAGP score scale derives from the 10th grade NMC summative assessment scale and the MATGP from the NMC Affiliate Integrated Math II scale. The performance standard levels of 725 scale score relates to the performance levels established for the PARCC/NMC Affiliate program.

5.1. Performance Standards

Student performance on both components of the NJGPA is compared with performance standards to determine if each student meets criteria to be considered Graduation Ready. Those students not meeting the criteria are Not Yet Graduation Ready. The cut scores that delineate Graduation Ready from Not Yet Graduation Ready were validated by New Jersey subject matter experts (SMEs) as described in the next section of this chapter. The suggested cut scores were approved by the New Jersey State Board of Education in May 2023.

5.2. Standard Setting Processes

This chapter will discuss two distinct standard setting processes that apply to the NJGPA. The first subsection will explain the standard setting processes of the PARCC/NMC Affiliate program, including the initial calibration of performance levels in high school ELA and mathematics assessments, which contributed to NJGPA development. The final subsection will outline a standards validation meeting held in February 2022 that validated recommended scale score level-725 cut scores.

The extensive standard setting activities described in the following section regarding the PARCC/NMC Affiliate program allowed for a streamlined, standard validation process for NJGPA cut scores. PARCC/NMC Affiliate assessments report five performance levels. PARCC/NMC Affiliate cut scores corresponding to level three are defined as approaching academic performance expectations. Level three is anchored to a scale score of 725. The 725 scale score was chosen as the target scale score for the NJGPA standards validation meetings.

5.2.1. Performance Level Setting (PLS) Process for the PARCC/NMC Affiliate System

NJGPA score scales were derived from score scales previously created for the PARCC/NMC Affiliate ELA grade 10, Algebra I and Geometry summative assessments. One of the main objectives of the PARCC/NMC Affiliate system was to provide information to students, parents, educators, and administrators as to whether students are on track in their learning for success after high school, defined as college- and career-readiness. To set performance levels associated with this objective, participating states and agencies used the evidence-based standard setting (EBSS) method (Beimers et al., 2012) for the PLS process. The EBSS method is a systematic method for combining various considerations into the

process for setting performance levels, including policy considerations, content standards, educator judgment about what students should know and be able to demonstrate, and research to support policy goals related to college- and career-readiness. A defined multistep process was used to allow a diverse set of stakeholders to consider the interaction of these elements in recommending performance-level threshold scores for each assessment.

The seven steps of the EBSS process that were followed to establish performance standards for the summative assessments are:

- Step 1: Define outcomes of interest and policy goals.
- Step 2: Develop research, data collection, and analysis plans.
- Step 3: Synthesize the research results.
- Step 4: Conduct pre-policy meeting.
- Step 5: Conduct PLS meetings with panels.
- Step 6: Conduct reasonableness review with post-policy panel.
- Step 7: Continue to gather evidence in support of standards.

A summary of key components within these steps is provided below. Additional details about each step in the PLS process are provided in the *Performance Level Setting Technical Report*.

5.2.1.1. Research Studies

Participating states and agencies conducted two research studies in support of their policy goals—the benchmarking study and the postsecondary educators’ judgment (PEJ) study. The benchmarking study included a review of the literature relative to college- and career-readiness as well as consideration of the percentage of students obtaining a level equivalent to college- and career-readiness on a set of external assessments (e.g., ACT, SAT, NAEP). The PEJ study involved a group of nearly 200 college faculty reviewing items on the Algebra II and ELA grade 11 assessments and making judgments about the level of performance needed on each item to be academically ready for an entry-level college-credit bearing course in mathematics or ELA. Additional details about the benchmarking study can be found in the *Performance Level Setting Technical Report* as well as in the *PARCC Benchmarking Study Report*. Additional details about the PEJ study can be found in the *Performance Level Setting Technical Report* as well as in the *Postsecondary Educators’ Judgment Study Final Report*.

5.2.1.2. Pre-Policy Meeting

Prior to the PLS meetings, a pre-policy meeting was convened to determine reasonable ranges that would be shown to panelists during the high school PLS meetings. Pre-policy meeting participants included representatives from both K–12 and higher education who served in roles such as commissioner/superintendent, deputy/assistant commissioner, state board member, director of assessment, director of academic affairs, senior policy associate, and so on. The reasonable ranges recommended by the pre-policy meeting defined the minimum and maximum percentage of students that would be expected to be classified as college- and career-ready. The pre-policy meeting participants reviewed the test purpose, how the performance standards will be used, and the results of the research studies to provide the recommendations for the reasonable ranges without viewing any student performance data.

5.2.1.3. Performance Level Setting Meetings

The task of the PLS committee was to recommend four threshold scores that would define the five performance levels for each assessment. Participating states and agencies solicited nominations from all Affiliate states that had administered the assessments in 2014–2015 for panelists to serve on the PLS committees. Nominations were solicited both from state departments of public education (K–12) and higher education (primarily for participation on the high school panels). When selecting panelists, an emphasis was placed on those educators who had content knowledge as well as experience with a variety of student groups and attempted to balance the panels in terms of state representation.

Participating states and agencies used an extended modified Angoff (Yes/No) method to collect educator judgments on the items. This method asked panelists to review each item on a reference form of the assessment and to make the following judgment: How many points would a borderline student at each performance level likely earn if they answered the question?

This extension to the Yes/No standard setting method (Plake et al., 2005) allowed for incorporation of the multipoint items by asking educators to evaluate (Yes or No) whether a borderline student would earn the maximum number of points on an item, a lesser number of points on an item, or no points on the item. In the case of a single point or multiple-choice item, this task simplifies to the standard Yes/No method.

After receiving training on the PLS procedure, panelists participated in three rounds of judgments for each assessment. Within each round, panelists were asked to consider the items in the test form, starting with the performance-based assessment (PBA) component and then the end-of-year (EOY) component. Each panelist made a judgment for the Level 2 performance level, followed by judgments for the Level 3 performance level, the Level 4 performance level, and the Level 5 performance level, in this order. The panelists entered their item judgments for each round by completing an online item judgment survey. Educator judgments were summed across items to create an estimated total score on the reference form for each performance level threshold. Feedback data relative to panelist agreement, student performance on the items, and student performance on the test as a whole were provided in between each of the three rounds of judgment. Panelists were shown the pre-policy reasonable ranges prior to making their Round 1 judgments and again as feedback data following each round of judgment.

A dry run of the PLS meeting process was held for grade 11 ELA and Algebra II in order to evaluate the implementation of the PLS method with the innovative characteristics of the summative assessments. The results of the dry run PLS meeting were used to implement improvements in the process for the operational PLS meetings. Additional information about the methods and results of the dry run PLS meeting is available in the full report in the *Performance Level Setting Dry-Run Meeting Report*.

Additional information about the methods and results of the PLS meetings is available in the *Performance Level Setting Technical Report*.

5.2.1.4. Post-Policy Reasonableness Review

Performance standards for all summative assessments were recommended by PLS committees and reviewed by the Governing Board and (for the Algebra II, Integrated Mathematics III, and ELA grade 11 assessments) the Advisory Committee on College Readiness as part of a post-policy reasonableness review. This group reviewed both the median threshold score recommendations from each committee and

the variability in the threshold scores as represented by the standard error of judgment (SEJ) of the committee. Adjustments to the median threshold scores that were within 2 SEJ were considered to be consistent with the PLS panels' recommendation.

In addition to voting to adopt the performance standards based on the committees' recommendations, this group also voted to conduct a shift in the performance levels to better meet the intended inferences about student performance. Holding the college- and career-ready (or on-track) expectations (i.e., the current level 4) constant, performance levels above this expectation were combined and performance levels below this expectation were expanded to create the final system of performance levels with three below and two above the college- and career-ready (or on-track) expectation. The shift in performance levels was accomplished using a scale anchoring process that involved two primary steps. In the first step, the top two performance levels, above college- and career-ready (or on-track), were combined into a single performance level and an additional performance level below college- and career-ready (or on-track) was created by empirically determining the midpoint between the existing two levels. In the second step, the performance level descriptors (PLDs) were updated using items that discriminated student performance well at this level to create a PLD aligned with the new empirically determined performance level. At this same time, PLDs for all performance levels were reviewed for consistency and continuity. Members of the original PLS committees were recruited to participate in this process. Additional information about this process can be found in the *Performance Level Setting Technical Report*.

5.2.2. NJGPA Standards Validation

The New Jersey Department of Education (NJDOE) conducted virtual standards validation meetings on February 1–2, 2022. Educators from throughout the state of New Jersey participated in these two-day meetings. The staff of NMC facilitated the meetings. The main goal of the meetings was to have SMEs recommend psychometric cut scores which delineate two performance levels: "Graduation Ready" and "Not Yet Graduation Ready" for the two components of the NJGPA, the ELAGP and the MATGP. Two additional aims of the meetings designed to prepare and assist SMEs in the identification of cut scores were:

- To allow workshop participants (i.e., subject matter experts [SMEs]) to gain an understanding of the assessment content and performance level descriptors (PLDs).
- To have SMEs learn a standard-setting methodology based on the Bookmark standard-setting procedure.

NMC facilitated Bookmark standard-setting procedures to validate the psychometric cut scores for the ELAGP and MATGP set at the 725 scale score level. The Bookmark standard-setting procedure (Lewis, Mitzel & Green, 1996) was developed in 1996 and fully documented in *Setting Performance Standards: A Guide to Establishing and Evaluating Performance Standards on Tests* (Cizek & Bunch, 2007). The procedure is an iterative set of activities intended to produce suggested cut scores based on the judgments of SMEs. The Bookmark standard-setting procedure was developed to provide a less complex method for judges to determine a cut score (Mitzel et al., 2001) and is one of the most used procedures for establishing cut scores for large-scale K–12 assessments (Baldwin et al., 2019).

For the NJGPA standards validation meeting, the standard bookmark procedure was modified. The SMEs were provided with the cut score levels, and discussion centered on how well the scores defined the

knowledge, skills and abilities that would distinguish students who are ready for graduation from those who are not yet ready. The SMEs did place their bookmarks in the locations they felt best described the graduation-ready level of performance. The calculated cut scores were then compared to the established cut score levels.

The key activities for the NJGPA standards validation included:

- Group discussion of PLDs and students at the threshold.
- Taking the NJGPA.
- SME examination of the test items (no item statistics provided).
- Group discussion of each item.
- Individual SME bookmark placement (Round 1 rating).
- Group discussion of Round 1 results.
- Individual SME bookmark placement (Final rating).

An Ordered Item Booklet (OIB) is integral to the Bookmark standard-setting procedure. Each NJGPA OIB contained items and metadata from the item bank. Each OIB consisted of a number of items which appeared on the Spring 2022 forms and some other items chosen to represent the range of topics and difficulty defining the overall item bank available for forms construction. NMC psychometricians created OIBs by ordering items from easiest to most difficult according to an item response theory (IRT) measure. The item statistics of record from previous administrations were used to order the items in the OIB.

During standard setting, SMEs proceeded through the OIB one page at a time and asked themselves whether they believe the “just barely meeting expectations,” or “threshold,” student would have at least a 0.67 probability of answering the item correctly (or obtaining the given score point). If the answer was “Yes,” then the SME proceeded to the next page. The page on which the SME answered “No” became the bookmark between the adjacent performance levels Not Yet Graduation Ready and Graduation Ready. A key element to the procedure was the reconciliation of the PLDs and a threshold student’s demonstrated knowledge, skills, and abilities (KSAs).

Following each of the bookmark placement rounds, NMC collected the SMEs’ bookmark placements (provided as a page number) and computed the median, median absolute difference, and minimum and maximum placement by group (i.e., breakout room) and by committee. NMC psychometricians then presented these results to the SMEs. Following the first round of bookmark placement, SMEs had the opportunity to discuss the subset of items within the individually bookmarked pages, considering the PLDs and the KSAs required by a threshold student. Following the first group discussions, SMEs had the opportunity to update their bookmark location based on the shared discussions. At the culmination of the final round, the committee’s median value was used to calculate the associated psychometric cut.

The cut scores suggested by the SMEs were calculated as IRT theta scores, which allows their application across various forms of the assessment. The cut score theta for the ELAGP is -0.31. The cut score theta for the MATGP is -0.049. The cut score for both subjects corresponds to a scale score of 725.

At the conclusion of the workshop, the SMEs were asked to complete a process evaluation survey. The survey collected information about the SMEs’ experiences in recommending a psychometric cut score for the NJGPA.

Chapter 6. Item Analysis

6.1. Overview

This chapter describes the results of the classical item analyses conducted for operational items on the spring operational forms. For the NJGPA, all forms were pre-equated, meaning that the scoring was based on item parameters estimated using data from prior assessment administrations. By contrast, item analysis results in this section are derived from the item statistics from post-administration analysis.

An operational item may appear on multiple test forms. The tables below list unique item counts for the assessments, and the reported item statistics may be based on student responses across multiple occurrences of an item. PCR items in the ELAGP are scored on two traits: (1) Written Expression and (2) Knowledge of Language and Conventions. ELAGP item counts include both scored PCR traits.

Spoiled or “do not score” items, if any, are excluded from the total test score in item analysis. These items are removed from scoring because of item performance, technical scoring issues, content concerns, or multiple/no correct answers. There were no spoiled items in the 2023 NJGPA testing. Thus, there were no exclusions.

For each item, analyses were conducted at the item level for each form. These analyses included difficulty (p-value and pseudo p-value) and discrimination (item-total correlation).

6.2. Data Screening Criteria

Analyses were performed on an Incomplete Data Matrix (IDM) that was generated from the results file. These analyses were done by form. Student records were removed prior to running the analyses if the records met any of the criteria indicated in the test calibration specs.

Items may not be scored due to a student omitting the item or to the student not yet reaching an item within the test. “Omitted” (i.e., skipped) items are items for which a student did not provide a response when items coming before and after have student responses. Item response scores for “omits” are re-coded as ‘0’ in the classical test theory (CTT) analyses and IRT IDM files whereas “Not reached” and “Not administered” items are considered missing and therefore do not contribute to the statistics.

6.3. Classical Item Analysis

6.3.1. Item Difficulty

The p-value represents the mean item score as a proportion of the maximum obtainable score points, indicating the item difficulty. Values range between 0 and 1. Higher values indicate easier items while lower values indicate more difficult items. For dichotomous items (item scored as either correct = 1 or incorrect = 0) the formula is

Equation 6.1

$$p\text{-value} = \bar{x}_i = \frac{1}{n} \cdot \sum_1^n x_i$$

where x_i are the individual student item scores on item i and n is the total number of students for whom the item was administered.

For polytomous items, the pseudo p-value is calculated by further dividing by the maximum obtainable points possible for the item:

Equation 6.2

$$\text{pseudo } p\text{-value} = \frac{\bar{x}}{T}$$

where $\bar{x}$ is the mean item score and T is the maximum obtainable points possible for the item.

Frequently, the p-value is reported as a percentage by multiplying by 100. For instance, a p-value of 0.67 means that 67 percent of the students answered a dichotomous item correctly.

6.3.2. Response Option or Score Point Proportions

A dichotomous item's alternate response options (i.e., distractors) are plausible but incorrect options that are included to test common misconceptions or miscalculations. Ideally, all response options should garner a proportion of student selections. These are calculated by the simple proportion formula:

Equation 6.3

$$\text{proportion} = \frac{N_o}{N_T}$$

where N_o is the number of students selecting the specific option and N_T is the total number of students for whom the question was administered, including those who did not record a response (i.e., omitted the item).

In the case of polytomous items, the numerator becomes the number of students obtaining the specific score point (N_{SP}):

Equation 6.4

$$\text{proportion} = \frac{N_{SP}}{N_T}$$

6.3.3. Item-Total Correlations

The Item-Total Correlation is the relationship between students' performance on the item and students' performance. Possible values range between -1 and $+1$. The correlation will be positive when the mean test score of the students answering the item correctly is greater than the mean test score of the students answering the item incorrectly. Negative values may indicate that an item has multiple correct answers or an incorrect answer key.

The point-biserial correlation (Crocker & Algina, 1986) is one possible Item-Total Correlation for dichotomously scored items. However, the correlation will be spuriously high because the item of interest

is also included in the total test score (i.e., correlating with itself; Henrysson, 1963). Therefore, a correction is made by using the means with the item deleted (i.e., the total operational test score not including the item of interest) from the calculation:

Equation 6.5

$$r_{pbis} = \frac{(\overline{M}'_+ - \overline{M}')}{S'} \sqrt{p/(1 - p)}$$

where $\overline{M}'_+$ is the mean score with the item deleted for students who answered the item correctly, $\overline{M}'$ is the mean score with the item deleted for all students, S' is the standard deviation with the item deleted for all students, and p is the item p-value (difficulty).

The Pearson correlation (polyserial) with the item of interest deleted is typically calculated for polytomous items by this equation:

Equation 6.6

$$r = \frac{\sum (x_i - \bar{x})(y'_i - \bar{y}')}{\sqrt{\sum (x_i - \bar{x})^2 \sum (y'_i - \bar{y}')^2}}$$

where x_i is the student score point on the item, $\bar{x}$ is the mean score for the item, y'_i is the total score with the item deleted for the student, and $\bar{y}'$ is the mean total score with the item deleted for all students (Lemke & Wiersma, 1976).

6.3.4. Item Bank Quality Control

Item analysis results are summarized by forms. In the below tables, OL stands for online form, which is the general student population testing form. AC1 is accommodated form 1 (e.g., 2023 TTS forms), and AC2 is accommodated form 2, the paper-based form (e.g., all other accommodations rolled forward from 2022). AC1 shares the same items with OL, and AC2 has different items. The Spanish form for Math shares the same items as AC2 but is disambiguated here for analysis purposes.

Table 6.1 Summary of Post-Administration P-values for ELAGP Operational Items.

Grade	Form	N of Unique Items	Mean p-Value	SD p-Value	Min p-Value	Max p-Value	Median p-Value
ELAGP	OL	20	0.53	0.12	0.34	0.75	0.53
	AC1	20	0.31	0.10	0.11	0.51	0.31
	AC2	20	0.37	0.13	0.19	0.72	0.36

Note. SD = standard deviation.

Table 6.1 presents post-administration p-value information for the ELAGP operational items. For the OL form, which includes 20 unique items, the mean p-value is 0.53, with a standard deviation of 0.12. The p-values range from a minimum of 0.34 to a maximum of 0.75, and the median p-value is 0.53.

For the AC1 form with 20 unique items, the mean p-value is 0.31, and the standard deviation is 0.10. The p-values for AC1 range from a minimum of 0.11 to a maximum of 0.51, with a median p-value of 0.31.

For the AC2 form, also comprising 20 unique items, the mean p-value is 0.37, with a standard deviation of 0.13. The p-values for AC2 range from a minimum of 0.19 to a maximum of 0.72, and the median p-value is 0.36.

Table 6.2 Summary of Post-Administration P-values for MATGP Operational Items.

Grade	Form	N of Unique Items	Mean p-value	SD p-value	Min p-value	Max p-value	Median p-value
MATGP	OL	30	0.44	0.20	0.05	0.75	0.48
	AC1	30	0.33	0.18	0.03	0.64	0.35
	AC2	30	0.32	0.18	0.03	0.70	0.33
	Spanish	30	0.20	0.16	0.01	0.53	0.16

Note. SD = standard deviation.

Table 6.2 presents a summary of post-administration p-value information for the mathematics operational items. For the OL form, there are 30 unique items with a mean p-value of 0.44, a standard deviation of 0.20, a minimum p-value of 0.05, a maximum p-value of 0.75, and a median p-value of 0.48.

Similarly, for forms AC1 and AC2, each comprising 30 unique items, the mean p-values are 0.33 and 0.32, with standard deviations of 0.18 for both. The minimum and maximum p-values vary, ranging from 0.03 to 0.64 for AC1 and 0.03 to 0.70 for AC2. The median p-values for AC1 and AC2 are 0.35 and 0.33, respectively.

In the case of the Spanish form, which also includes 30 unique items, the mean p-value is 0.20, with a standard deviation of 0.16. The minimum and maximum p-values for Spanish are 0.01 and 0.53, with a median p-value of 0.16.

Table 6.3 Summary of Post-Administration Item-Total Polyserial Correlations for ELAGP Operational Items.

Grade	Form	N of Unique Items	Mean ITC	SD ITC	Min ITC	Max ITC	Median ITC
ELAGP	OL	20	0.57	0.14	0.43	0.93	0.53
	AC1	20	0.49	0.16	0.30	0.87	0.45
	AC2	20	0.58	0.16	0.30	0.95	0.55

Note. SD = standard deviation.

Table 6.3 presents post-administration item-total correlations for the ELAGP operational items, broken down by different forms. For the OL form, there are 20 unique items with a mean polyserial correlation of 0.57, a standard deviation of 0.14, a minimum correlation of 0.43, and a maximum correlation of 0.93. The median polyserial correlation for OL is 0.53.

Similarly, for the AC1 form, there are also 20 unique items, but the mean polyserial correlation is slightly lower at 0.49, with a standard deviation of 0.16. The minimum and maximum polyserial correlations for AC1 are 0.30 and 0.87, respectively, and the median is 0.45.

For the AC2 form, again with 20 unique items, the mean polyserial correlation is 0.58, and the standard deviation is 0.16. The minimum and maximum polyserial correlations for AC2 are 0.30 and 0.95, and the median is 0.55.

Table 6.4 Summary of Post-Administration Item-Total Polyserial Correlations for Mathematics Operational Items.

Grade	Form	N of Unique Items	Mean ITC	SD ITC	Min ITC	Max ITC	Median ITC
MATGP	OL	30	0.68	0.13	0.37	0.87	0.70
	AC1	30	0.63	0.14	0.37	0.84	0.64
	AC2	30	0.65	0.13	0.41	0.82	0.70
	Spanish	30	0.39	0.12	0.19	0.71	0.40

Note. SD = standard deviation.

Table 6.4 presents post-administration item-total correlations for the mathematics MATGP operational items. In the OL form, there are 30 unique items with a mean polyserial correlation of 0.68, a standard deviation of 0.13, a minimum polyserial correlation of 0.37, a maximum polyserial correlation of 0.87, and a median polyserial correlation of 0.70.

Similarly, for the AC1 and AC2 forms, there are 30 unique items each, with mean polyserial correlations of 0.63 and 0.65, standard deviations of 0.14 and 0.13, minimum polyserial correlations of 0.37 and 0.41, maximum polyserial correlations of 0.84 and 0.82, and median polyserial correlations of 0.64 and 0.70, respectively.

In the Spanish form there are 30 unique items with a mean polyserial correlation of 0.39, a standard deviation of 0.12, a minimum polyserial correlation of 0.19, a maximum polyserial correlation of 0.71, and a median polyserial correlation of 0.40.

Chapter 7. Item Response Theory Analysis, Calibration and Scaling

7.1. Overview

The ELA and mathematics core forms are linked to their respective base reporting scales via pre-equating. This section of the technical report describes the item response theory (IRT) models used for pre-equating, item calibration and calculation of students' scale scores for ELAGP and MATGP. Descriptive statistics of the distributions of item parameter estimates for each assessment component are also included in this chapter.

7.2. IRT Models

The items on mathematics operational core forms were calibrated using the IRT two-parameter logistic (2PL) model (Birnbaum, 1968) for dichotomously scored items and the Generalized Partial Credit Model (GPCM) (Muraki, 1997) for polytomously scored items. ELA items were calibrated using the GPCM. To be concise, we express the two models using a single formula since the 2PL model can be considered algebraically nested within the GPCM, which is denoted as

Equation 7.1

$$P_{im}(\theta_j) = \frac{\exp[\sum_{k=0}^m D a_i(\theta_j - b_i + d_{ik})]}{\sum_{v=0}^{M_i-1} \exp[\sum_{k=0}^v D a_i(\theta_j - b_i + d_{ik})]}$$

where $a_i(\theta_j - b_i + d_{ik}) \equiv 0$; $P_{im}(\theta_j)$ is the probability of the j^{th} student with θ_j getting score m on item i ; D is the IRT scale constant (1.7); a_i is the discrimination parameter of item i ; b_i is the item difficulty parameter of item i ; d_{ik} is the k^{th} step deviation parameter for item i ; M_i is the number of score categories of item i with possible item scores as consecutive integers from zero to $M_i - 1$; and v indexes the response categories and is iterated from zero to $M_i - 1$. For items with just two response categories, with one category being scored as a correct response and the other category scored as an incorrect response, the d_{ik} parameter becomes zero since no step parameters are needed, making the 2PL model a special case of the GPCM.

7.3. Summary Statistics and Distributions from IRT Analyses

Tables 7.1 and 7.2 present summary statistics for the IRT (b - and a -) parameter estimates, the standard errors (SEs) of the parameter estimates, and the IRT model fit values (chi-square and adjusted fit) for ELA and mathematics components. The summary statistics for IRT parameter estimates include all the items administered in the spring administration.

Table 7.1 shows the b - and a -parameter estimates for the ELAGP and MATGP components. In IRT, the a -parameter refers to the item's ability to discriminate the performance of test takers. The b -parameter represents the item's level of difficulty, with larger, positive values reflecting a harder item. The items summarized in Table 7.1 include all unique items on online general forms, paper forms and online

accommodated forms. ELAGP PCR item traits are included in the item counts as IRT parameters are calculated for each trait.

IRT calibration of 2023 data occurred in the summer and fall after spring 2023 testing. During IRT calibration, items on the NJGPA spring 2023 online regular forms were calibrated using the data from the NJGPA spring 2023 online testing and scaled to the NMC bank scale. Student responses to online regular forms are the only data used in IRT calibration.

Items on spring 2023 paper- and online-accommodated forms have pre-equated IRT parameters calculated during item calibration in 2022.

Table 7.1 IRT Parameter Estimates Summary for All Items for ELAGP and MATGP

Test Code	No. of Items	No. of Points	a Estimates Summary				b Estimates Summary			
			Mean	SD	Min.	Max.	Mean	SD	Min.	Max.
ELAGP	44	148	0.50	0.27	0.16	1.18	0.55	0.67	-1.08	1.83
MATGP	56	106	0.64	0.26	0.16	1.34	0.84	1.19	-1.50	3.74

Note. SD = standard deviation.

Table 7.2 shows the source year for the IRT item parameters for the ELAGP and MATGP assessments.

Table 7.2 IRT Parameter Distribution by Year for All Items for ELAGP and MATGP

Test Code	No. of Items	2022	2023
ELAGP	44	22	22
MATGP	56	26	30

7.4. Scale Scores

Student NJGPA results in ELA and mathematics are reported using scales that designate student performance into one of two performance levels that delineate the knowledge, skills, and practices students are able to demonstrate. Students meeting the graduation performance standards described in Chapter 5 of this report are described as Graduation Ready. Students not meeting these standards are Not Yet Graduation Ready.

The ELA and mathematics components are designed to measure and report results in categories called master claims and subclaims. Major claims (or simply “claims”) are at a higher level than subclaims, with content representing multiple subclaims contributing to each claim outcome. A summative scale score is reported for the mathematics assessment. A summative scale score and claim scores for Reading and Writing are reported for the ELA assessment.

7.4.1. Establishing the Reporting Scales

There are 201 defined summative scale score points for both the ELA and mathematics components, with the scale ranging from 650 (lowest obtainable scale score) to 850 (highest obtainable scale score). Scale scores are calculated from the IRT theta score calculated for each student from their test responses. Scale score calculations are based on a linear transformation between two sets of score values: IRT theta scores and scale scores. Not all possible scale scores may be realized in a scoring table. The relationship between raw and scale scores is expressed through scaling constants that define the mathematical relationship between the two score scales. The threshold scores and scaling constants for the overall ELAGP and MATGP scales are reported in table 7.3.

To facilitate the linear transformation between the IRT-derived theta scale and the reporting scale score range, anchor points were specified. Anchors were selected at the 725 and 750 scale score levels. The 725 scale score level was chosen as an anchor due to the cut score being set at this level. Psychometric analysis was conducted to identify other possible anchor points. The 700 and 750 levels were considered. These levels are among the defined performance levels in the grade 10 ELA and Integrated Math II NMC assessments, but would only serve as 2nd anchor points for the ELAGP and MATGP scales. Pearson and NMC psychometricians calculated the 750 scale score anchor points for ELAGP and MATGP scales at approximately the same theta score differences from the 725 threshold scores, as would define the 750-scale score level in similar NMC Affiliate program test scales.

Table 7.3 Threshold Scores and Scaling Constants for ELAGP and MATGP

Assessment	Theta	Scale Score	A	B
ELAGP	-0.310	725	43.2227	738.399
MATGP	-0.049	725	29.80448	726.4604

7.4.2. ELA Reading and Writing Claim Scales

As with the summative scale scores, scale scores for Reading and Writing were defined for each test as a linear transformation of the IRT theta scale. There are 81 defined scale score points possible for Reading, ranging from 10 to 90. There are 51 defined scale score points possible for Writing, ranging from 10 to 60. Not all possible scale scores may be realized in a scoring table. The same IRT theta scale was used for Reading and Writing as was used for the ELA summative scores. The scaling constants for the Reading and Writing scales are reported in Table 7.4.

Table 7.4 Scaling Constants for Reading and Writing ELAGP Subscores

Assessment	Reading		Writing	
	AR	BR	AW	BW
ELAGP	17.28907	45.35961	8.644537	32.67981

7.4.3. Subclaims Scales

Table 7.5 reports the threshold scores for ELA and mathematics subclaim scales. These values are set at 1.5 times greater than and less than the standard error (+/- 1.5 SE) calculated for each threshold score. For ELA 1.5 SE was calculated to be .470 theta, and mathematics 1.5 SE = .4526 theta.

Table 7.5 NJGPA Subclaim Theta Cut Scores

Assessment	Theta
ELAGP	-0.7804
	0.1605
MATGP	-0.5164
	0.4148

7.4.4. Creating Conversion Tables

A conversion table relates the number of points earned by a student for the ELA summative score, the mathematics summative score, the Reading claim score, or the Writing claim score to the corresponding scale score for the test form administered to that student. An IRT inverse test characteristic curve (TCC) approach is used to develop the relationship between point scores and IRT ability parameters or theta scores. In carrying out the calculations, estimates of item parameters and thetas are substituted for parameters in the formulas in each of the following steps.

Step 1: Calculate the expected item score (i.e., estimated item true score) for every theta in the selected range (between -15 and +15, in 0.0001 increments) based on the GPCM for both dichotomous and polytomous items:

Equation 7.2

$$s_i(\theta_j) = \sum_{m=0}^{M_i-1} m P_{im}(\theta_j)$$

Equation 7.3

$$P_{im}(\theta_j) = \frac{\exp[\sum_{k=0}^m D a_i(\theta_j - b_i + d_{ik})]}{\sum_{v=0}^{M_i-1} \exp[\sum_{k=0}^v D a_i(\theta_j - b_i + d_{iv})]}$$

where $a_i(\theta_j - b_i + d_{i0}) \equiv 0$; $s_i(\theta_j)$ is the expected item score for item i on theta, θ_j ; $P_{im}(\theta_j)$ is the probability of a student, j , with θ_j getting score m on item i ; m_i is the number of score categories of item i , with possible item scores as consecutive integers from 0 to $M_i - 1$; D is the IRT scale constant (1.7); a_i is the item slope parameter; b_i is the item location parameter reflecting overall item difficulty; d_{ik} is a location parameter incrementing the overall item difficulty to reflect the difficulty of earning score category k ; v is the number of score categories. Since the 2PL model can be considered a special case of the GPCM, the latter can be used to calibrate dichotomously scored items, resulting in 2PL model item parameters.

Step 2: Calculate the expected (weighted) test score for every theta in the selected range:

Equation 7.4

$$T_j = \sum_{i=1}^I w_i s_i(\theta_j)$$

where T_j is the expected (weighted) test score on theta, θ_j ; w_i is the item weight for item i I is the total number of items in the test form.

Step 3: Calculate the estimated conditional standard error of measurement (CSEM) for each theta in the selected range:

Equation 7.5

$$CSEM_j = \sqrt{\frac{1}{\sum_{i=1}^I L_i(\theta_j)}}$$

Equation 7.6

$$L_i(\theta_j) = (Da_i)^2[s_{i2}(\theta_j) - s_i^2(\theta_j)]$$

Equation 7.7

$$s_{i2}(\theta_j) = \sum_{m=0}^{M_i-1} m^2 P_{im}(\theta_j)$$

where $L_i(\theta_j)$ is the estimated item information function for item i on theta, θ_j .

Step 4: Match every raw score with a theta. θ_j is the theta for a raw score r_h , if $T_j - r_h$ is minimum across all T_j .

Step 5: Calculate the reported scale score. Using the A and B scaling constants in Appendix 7, convert each theta value to a scale score and each theta CSEM to a scale score CSEM:

Equation 7.8

$$ScaleScore = A \times \theta + B$$

Equation 7.9

$$CSEM = CSEM_{\theta} \times A$$

The scale scores are rounded to the nearest whole number, and CSEMs are rounded to the tenths place. Furthermore, the scale scores are truncated with the lowest obtainable scale score (LOSS) of 650 and highest obtainable scale score (HOSS) of 850.

Figures 7.1–7.3 contain the TCC, estimated CSEM and estimated test information function (TIF) curves for all operational forms of the ELAGP. Within each figure, the curve is reported on the theta scale; and the vertical dotted line indicates the performance level threshold score on the theta scale.

Figures 7.4–7.6 contain TCC, CSEM, and TIF curves for the general population online core MATGP form.

Figure 7.1 ELAGP 2023 Forms TCCs




Figure 7.2 ELAGP 2023 Forms CSEM Curves




Figure 7.3 ELAGP 2023 Forms TIF Curves




Figure 7.4 MATGP 2023 Forms TCCs




Figure 7.5 MATGP 2023 Forms CSEM Curves





Figure 7.6 MATGP 2023 Forms TIF Curves



7.4.5. Score Distributions

The overall performance of students taking the 2023 NJGPA core forms is reported in the cumulative score distributions given in tables 7.6 and 7.7. Score distribution information is also illustrated in Figures 7.7 and 7.8. The vertical axis of each graph represents the proportion of students earning the scale score point indicated along the horizontal axis. For the summative score distribution, the y-axis ranges from 0 to .08 and the x-axis from 650 to 850.

Table 7.6 Scale Score Cumulative Frequencies: ELAGP

Score Band	Count	Percent	Cumulative Count	Cumulative Percent
650-654	2,667	2.83	2,667	2.83
655-659	—	0.00	2,667	2.83
660-664	967	1.03	3,634	3.86
665-669	998	1.06	4,632	4.92
670-674	—	0.00	4,632	4.92
675-679	961	1.02	5,593	5.94
680-684	969	1.03	6,562	6.97
685-689	948	1.00	7,510	7.97
690-694	971	1.03	8,481	9.00
695-699	971	1.03	9,452	10.03
700-704	1,020	1.08	10,472	11.11
705-709	2,003	2.12	12,475	13.23
710-714	982	1.04	13,457	14.27
715-719	2,122	2.25	15,579	16.52
720-724	1,123	1.19	16,702	17.71
725-729	2,388	2.53	19,090	20.24
730-734	2,584	2.74	21,674	22.98
735-739	2,913	3.09	24,587	26.07
740-744	1,507	1.60	26,094	27.67
745-749	3,221	3.41	29,315	31.08
750-754	3,459	3.67	32,774	34.75
755-759	3,771	4.00	36,545	38.75
760-764	4,005	4.25	40,550	43.00
765-769	4,354	4.62	44,904	47.62

Score Band	Count	Percent	Cumulative Count	Cumulative Percent
770-774	4,485	4.75	49,389	52.37
775-779	2,295	2.43	51,684	54.80
780-784	4,457	4.72	56,141	59.52
785-789	4,456	4.72	60,597	64.24
790-794	4,374	4.64	64,971	68.88
795-799	4,332	4.59	69,303	73.47
800-804	2,036	2.16	71,339	75.63
805-809	3,906	4.14	75,245	79.77
810-814	3,695	3.92	78,940	83.69
815-819	3,260	3.46	82,200	87.15
820-824	1,506	1.60	83,706	88.75
825-829	2,608	2.76	86,314	91.51
830-834	1,156	1.23	87,470	92.74
835-839	1,958	2.08	89,428	94.82
840-844	813	0.86	90,241	95.68
845-849	773	0.82	91,014	96.50
850	3,324	3.52	94,338	100.00

Figure 7.7 2023 ELAGP Score Distribution





Table 7.7 Scale Score Cumulative Frequencies: MATGP

Score Band	Count	Percent	Cumulative Count	Cumulative Percent
650-654	268	0.30	268	0.30
655-659	720	0.82	988	1.12
660-664	—	0.00	988	1.12
665-669	—	0.00	988	1.12
670-674	1,517	1.72	2,505	2.84
675-679	—	0.00	2,505	2.84
680-684	2,650	3.01	5,155	5.85
685-689	3,655	4.16	8,810	10.01
690-694	—	0.00	8,810	10.01
695-699	4,448	5.06	13,258	15.07
700-704	4,623	5.26	17,881	20.33
705-709	4,479	5.09	22,360	25.42
710-714	4,209	4.79	26,569	30.21

Score Band	Count	Percent	Cumulative Count	Cumulative Percent
715-719	3,794	4.31	30,363	34.52
720-724	6,779	7.71	37,142	42.23
725-729	3,028	3.44	40,170	45.67
730-734	5,656	6.43	45,826	52.10
735-739	2,556	2.91	48,382	55.01
740-744	4,584	5.21	52,966	60.22
745-749	4,151	4.72	57,117	64.94
750-754	3,746	4.26	60,863	69.20
755-759	3,346	3.80	64,209	73.00
760-764	2,992	3.40	67,201	76.40
765-769	2,820	3.21	70,021	79.61
770-774	2,497	2.84	72,518	82.45
775-779	2,394	2.72	74,912	85.17
780-784	2,251	2.56	77,163	87.73
785-789	2,049	2.33	79,212	90.06
790-794	1,949	2.22	81,161	92.28
795-799	1,867	2.12	83,028	94.40
800-804	1,628	1.85	84,656	96.25
805-809	696	0.79	85,352	97.04
810-814	646	0.73	85,998	97.77
815-819	921	1.05	86,919	98.82
820-824	336	0.38	87,255	99.20
825-829	260	0.30	87,515	99.50
830-834	—	0.00	87,515	99.50
835-839	167	0.19	87,682	99.69
840-844	113	0.13	87,795	99.82
845-849	—	0.00	87,795	99.82
850	156	0.18	87,951	100.00

Figure 7.8 2023 MATGP Score Distribution



Score distributions are also presented for the Reading and Writing sub scores, in figures 7.9 and 7.10, respectively. For the Reading distribution, the y-axis ranges from 0 to .09 and the x-axis from 10 to 90. For the Writing distribution, the y-axis ranges from 0 to .08 and the x-axis from 10 to 60.

Figure 7.9 2023 Reading Score Distribution





Figure 7.10 2023 Writing Score Distribution



The performance of demographic groups of students is summarized in Tables 7.8 and 7.9 for ELA and mathematics, respectively. Each table reports the number of students in ethnic groups, and other

demographic variables including gender, economic and English Learner status, and students with disabilities. For each demographic grouping the mean scale score and standard deviation is presented along with the minimum and maximum scale scores for reference. Table 7.8 also summarizes group performance on the Reading and Writing scores.

Table 7.8 Subgroup Performance for ELAGP Scale Scores

Group Type	Subgroup	N	Mean	SD	Min.	Max.
Full Summative Score		94,338	767.69	48.11	650	850
Gender	Female	46,698	774.71	45.96	650	850
	Male	47,497	760.72	49.17	650	850
Ethnicity	American Indian/Alaska Native	139	756.16	53.03	650	850
	Asian	10,329	800.19	38.57	650	850
	Black or African American	12,851	747.48	45.5	650	850
	Hispanic/Latino	28,349	749.53	50.16	650	850
	Native Hawaiian or Pacific Islander	312	777.54	48.33	650	850
	Two or more races	2,140	774.37	44.24	650	850
	White	39,610	778.26	41.68	650	850
Economic Status*	Not Economically Disadvantaged	67,023	775.82	45.46	650	850
	Economically Disadvantaged	27,315	747.73	48.61	650	850
English Learner Status	Non-English Learner	89,506	771.98	44.82	650	850
	English Learner	4,832	688.18	36.25	650	850
Disabilities	Students without Disabilities	78,149	773.14	46.04	650	850
	Student with Disability (SWD)	16,189	741.36	49.2	650	850
Reading Summative Score		94,338	56.23	19.84	10	90
Gender	Female	46,698	57.99	19.13	10	90
	Male	47,497	54.48	20.37	10	90

Group Type	Subgroup	N	Mean	SD	Min.	Max.
Ethnicity	American Indian/Alaska Native	139	52.11	20.77	10	90
	Asian	10,329	69.96	16.95	10	90
	Black or African American	12,851	48.31	18.41	10	90
	Hispanic/Latino	28,349	48.22	19.92	10	90
	Native Hawaiian or Pacific Islander	312	59.16	19.62	10	90
	Two or more races	2,140	59.83	18.56	10	90
	White	39,610	60.7	17.45	10	90
Economic Status*	Not Economically Disadvantaged	67,023	59.76	18.94	10	90
	Economically Disadvantaged	27,315	47.6	19.34	10	90
English Learner Status	Non-English Learner	89,506	57.94	18.66	10	90
	English Learner	4,832	24.72	13.74	10	90
Disabilities	Students without Disabilities	78,149	58.24	19.19	10	90
	Student with Disability (SWD)	16,189	46.53	20.08	10	90
Writing Summative Score		94,338	38.28	12.21	10	60
Gender	Female	46,698	40.58	11.32	10	60
	Male	47,497	36	12.63	10	60
Ethnicity	American Indian/Alaska Native	139	35.14	13.91	10	60
	Asian	10,329	45.25	9.06	10	60
	Black or African American	12,851	33.53	12.44	10	60
	Hispanic/Latino	28,349	34.61	13.27	10	60
	Native Hawaiian or Pacific Islander	312	41.05	12.06	10	60
	Two or more races	2,140	39.21	11.23	10	60
	White	39,610	40.54	10.62	10	60

Group Type	Subgroup	N	Mean	SD	Min.	Max.
Economic Status	Not Economically Disadvantaged	67,023	39.97	11.42	10	60
	Economically Disadvantaged	27,315	34.12	13.08	10	60
English Learner Status	Non-English Learner	89,506	39.28	11.41	10	60
	English Learner	4,832	19.57	11.44	10	60
Disabilities	Students without Disabilities	78,149	39.66	11.5	10	60
	Student with Disability (SWD)	16,189	31.58	13.29	10	60

Note. SD = standard deviation. Economic status was based on participation in National School Lunch Program (NSLP): receipt of free or reduced-price lunch (FRL).

Table 7.9 Subgroup Performance for MATGP Scale Scores

Group Type	Subgroup	N	Mean	SD	Min.	Max.
Full Summative Score		87,951	736.71	36.48	650	850
Gender	Female	43,628	736.34	35.19	650	850
	Male	44,187	737.07	37.73	650	850
Ethnicity	American Indian/Alaska Native	127	732.46	35.87	659	807
	Asian	9,759	772.77	35.87	650	850
	Black or African American	11,590	715.87	28.43	650	829
	Hispanic/Latino	26,289	720.23	29.9	650	850
	Native Hawaiian or Pacific Islander	299	738.75	35.36	659	843
	Two or more races	2,010	743.44	36.34	650	850
	White	37,322	744.82	32.89	650	850
Economic Status	Not Economically Disadvantaged	63,118	743.69	36.57	650	850
	Economically Disadvantaged	24,833	718.99	29.6	650	850
English Learner Status	Non-English Learner	83,068	738.85	36.03	650	850

Group Type	Subgroup	N	Mean	SD	Min.	Max.
	English Learner	4,883	700.37	22.09	650	815
Disabilities	Students without Disabilities	73,323	740.3	36.12	650	850
	Student with Disability (SWD)	14,628	718.74	32.77	650	850
Language Form	Spanish	1,820	697.11	19.25	650	798

Note. SD = standard deviation. Economic status was based on participation in National School Lunch Program (NSLP): receipt of free or reduced-price lunch (FRL).

Chapter 8. Student Demographics and Differential Item Functioning

More than 90,000 high school juniors took the two components of the NJGPA—the ELAGP and the MATGP—in the spring of 2023. Chapter 8 reports important demographic descriptions of the testing populations who completed both components, as well as investigation of Differential Item Functioning (DIF). DIF analyses are utilized to detect systematic differences in performance on the test items between demographic groups.

8.1. Rules for Inclusion of Students in Analyses

Criteria for inclusion of students were implemented prior to all operational analyses. These rules were established for the PARCC/NMC Affiliate assessment program and help to determine which, if any, student records should be removed from analyses. This data screening process resulted in higher quality, albeit slightly smaller, data sets.

Student response data were included in analyses when:

- 1. Valid form numbers were observed for each unit for online assessments or for the full form for paper assessments,
- 2. Student records were not flagged as “void” (i.e., do not score), and
- 3. The student attempted at least 25 percent of the items in each unit or form.

Additionally, in cases where students had more than one valid record, the record with the higher raw score was chosen. Records for students with administration issues or anomalies were excluded from analyses.

8.2. Time to Attempt Assessment Items

It is important to understand how long students may take to answer each item on the assessment. Table 8.1 reports key timing information for the ELAGP items. For each operational form, Table 8.1 reports the minimum, maximum, median, and mean time to answer the items on the form and the items in each unit which makes up the test form. Also reported are the standard deviation of the mean, the time for 80% of respondents to answer (P80).

Table 8.1 Time in seconds for all test items and items by unit: ELAGP

Form_ID	N	Min	Median	P80	Max	Mean	SD	Unit 1 Mean	Unit 2 Mean	Unit 3 Mean
AC1	3,406	0	185	461	11,584	313.83	422.00	293.98	334.49	329.34
AC2	175	0	30	253	5,943	190.56	407.38	194.28	196.39	150.43
OL	94,338	0	210	458	13,404	353.63	463.29	310.70	397.51	388.68

Note. SD=Standard Deviation.

Similarly, Table 8.2 reports key timing information for the MATGP items. For each operational form, Table 8.2 reports the minimum, median, maximum, and mean time to answer the item on the form and the items in each unit. Also reported are the standard deviation of the mean, the time for 80% of respondents to answer (P80).

Table 8.2 Time in seconds for all test items and items by unit: MATGP

Form_ID	N	Min	Median	P80	Max	Mean	SD	Unit 1 Mean	Unit 2 Mean
AC1	10,392	0	122	304	7302	212.59	269.57	197.32	232.57
AC2	181	0	32	182	3887	125.38	243.52	123.66	127.64
OL	86,131	0	125	312	7370	219.85	272.31	199.16	246.91
OS1	1,820	0	125	306	5616	211.09	266.18	206.06	217.66

Note. SD=Standard Deviation.

8.3. Demographics

Table 8.3 presents the number and percentage of students who took the ELAGP and MATGP in each mode, either computer-based test (CBT) or paper-based test (PBT). Table 8.3 also provides this information for students who took the mathematics component in Spanish.

Markedly more students tested online than on paper for both content areas. For ELA, the percentage of online students by grade level was greater than 99 percent. For all mathematics students, the percentage of students testing online was greater than 99 percent. The percentage of students taking Spanish-language mathematics online forms was greater than or equal to 99 percent.

Table 8.3 ELAGP and MATGP Test Takers

Test	% of All Students	N Students	N CBT	% CBT	N PBT	% PBT
ELAGP	100.0	97,919	97,810	99.9	109	0.1
MATGP	100.0	98,524	98,400	99.9	124	0.1
MATGP Spanish	100.0	2,110	2,107	99.9	n/r	n/r

Note. n/r = not reported due to n<20. CBT=Computer-based test. PBT=Paper-based test.

Table 8.4 summarizes demographic information for students with valid ELA scores, and Table 8.5 presents demographics for students with valid mathematics scores. Percentages are not reported in instances where fewer than 20 students were tested.

Table 8.4 Demographic Information: ELAGP Test Takers

Demographic	N	Percent
Economically Disadvantaged	28,637	29.2
Student with Disability (SWD)	19,749	20.2
English Learner	4,903	5.0
Male	49,795	50.9
Female	47,977	49.0
American Indian/Alaska Native	144	0.1
Asian	10,421	10.6
Black or African American	13,455	13.7
Hispanic/Latino	29,536	30.2
White/Caucasian	41,218	42.1
Native Hawaiian or Pacific Islander	319	0.3
Two or more races	2,197	2.2
Unknown Ethnicity	629	0.6

Table 8.5 Demographic Information: MATGP Test Takers

Demographic	N	Percent
Economically Disadvantaged	29,024	29.5
Student with Disability (SWD)	19,704	20.0
English Learner	5,612	5.7
Male	50,099	50.8
Female	48,277	49.0
American Indian/Alaska Native	146	0.1
Asian	10,457	10.6
Black or African American	13,431	13.6
Hispanic/Latino	30,075	30.5
White/Caucasian	41,268	41.9
Native Hawaiian or Pacific Islander	319	0.3
Two or more races	2,204	2.2
Unknown Ethnicity	624	0.6

8.4. Differential Item Functioning

Differential Item Functioning (DIF) was calculated for items used on the online forms for ELAGP and MATGP. If an item performs differentially across identifiable subgroups (e.g., gender or ethnicity) when students are matched on ability, the item may be measuring something other than the intended construct. It is important, however, to recognize that item performance differences flagged for DIF might be related to actual differences in relevant knowledge or skills (item impact) or statistical Type I error (i.e., false positive error). As a result, DIF statistics are used to identify *potential* item bias. Subsequent reviews by content experts and bias/sensitivity committees are required to determine the source and meaning of performance differences.

8.4.1. DIF Procedures

The Mantel-Haenszel (MH) DIF statistic was calculated for selected-response items and for dichotomously scored constructed-response items. In this method, students are classified into relevant subgroups of interest (e.g., gender or ethnicity). Using the raw score total as the criteria, students in a certain total score category in the focal group (e.g., females) are compared with students in the same total score category in the reference group (e.g., males). For each item, students in the focal group are also compared to students in the reference group who performed equally well on the whole test. The common odds ratio is estimated across all categories of matched student ability using the formula illustrated in Equation 8.1 (Dorans & Holland, 1993), and the resulting estimate is interpreted as the relative likelihood of success on a particular item for members of two groups when matched on ability.

Equation 8.1

$$\hat{\alpha}_{MH} = \frac{\sum_{s=1}^S \frac{R_{rs}W_{fs}}{N_{ts}}}{\sum_{s=1}^S \frac{R_{fs}W_{rs'}}{N_{ts}}}$$

in which:

S = the number of score categories,

R_{rs} = the number of students in the reference group who answer the item correctly,

W_{fs} = the number of students in the focal group who answer the item incorrectly,

R_{fs} = the number of students in the focal group who answer the item correctly,

W_{rs} = the number of students in the reference group who answer the item incorrectly, and

N_{ts} = the total number of students.

To facilitate the interpretation of MH results, the common odds ratio is frequently transformed to the delta scale using the following formula (Holland & Thayer, 1988):

Equation 8.2

$$MH\ D - DIF = -2.35 \ln(\hat{\alpha}_{MH})$$

Positive values indicate DIF in favor of the focal group (i.e., positive DIF items are differentially easier for the focal group), whereas negative values indicate DIF in favor of the reference group (i.e., negative DIF items are differentially easier for the reference group).

For polytomously scored constructed-response items, the MH D-DIF statistic is not calculated; instead, the standardization DIF (Dorans & Schmitt, 1991; Zwick et al., 1997; Dorans, 2013), in conjunction with the Mantel chi-square statistic (Mantel, 1963; Mantel & Haenszel, 1959), is used to identify items with DIF.

The standardization DIF compares the item means of the two groups after adjusting for differences in the distribution of students across the values of the matching variable (i.e., total test score) and is calculated using the following formula:

Equation 8.3

$$STD - EISDIF = \frac{\sum_{s=1}^S N_{fs} \times E_f(Y|X=s)}{\sum_{s=1}^S N_{fs}} - \frac{\sum_{s=1}^S N_{rs} \times E_r(Y|X=s)}{\sum_{s=1}^S N_{rs}},$$

in which:

X = the total score,

Y = the item score,

S = the number of score categories,

N_{rs} = the number of students in the reference group in score category s ,

N_{fs} = the number of students in the focal group in score category s ,

E_r = the expected item score for the reference group, and

E_f = the expected item score for the focal group.

A positive STD-EISDIF value means that, conditional on the total test score, the focal group has a higher mean item score than the reference group. In contrast, a negative STD-EISDIF value means that, conditional on the total test score, the focal group has a lower mean item score than the reference group.

8.4.2. DIF Classification

Based on the DIF statistics and significance tests, items are classified into three categories and assigned values of A, B, or C (Zieky, 1993). Category A items contain negligible DIF, Category B items exhibit slight to moderate DIF, and Category C items possess moderate to large DIF values. Items classified with C level DIF are items that require more thorough examination and review. Positive values indicate that, conditional on the total score, the focal group has a higher mean item score than the reference group. In contrast, negative DIF values indicate that, conditional on the total test score, the focal group has a lower mean item score than the reference group. The flagging criteria for dichotomously scored items are presented in Table 8.4, and the flagging criteria for polytomously scored constructed-response items are provided in Table 8.5.

Table 8.4 DIF Categories for Dichotomous Selected-Response and Constructed-Response Items

DIF Category	Criteria
A (negligible)	Absolute value of the MH D-DIF is not significantly different from zero, or is less than one.
B (slight to moderate)	1. Absolute value of the MH D-DIF is significantly different from zero but not from one, and is at least one; or 2. Absolute value of the MH D-DIF is significantly different from one, but is less than 1.5. Positive values are classified as “B+” and negative values as “B-”.
C (moderate to large)	Absolute value of the MH D-DIF is significantly different from one, and is at least 1.5. Positive values are classified as “C+” and negative values as “C-”.

Table 8.5 DIF Categories for Polytomous Constructed-Response Items

DIF Category	Criteria
A (negligible)	Mantel Chi-square p-value > 0.05 or $ STD-EISDIF/SD \leq 0.17$
B (slight to moderate)	Mantel Chi-square p-value < 0.05 and $ STD-EISDIF/SD > 0.17$
C (moderate to large)	Mantel Chi-square p-value < 0.05 and $ STD-EISDIF/SD > 0.25$

Note: STD-EISDIF = standardized DIF; SD = total group standard deviation of item score.

8.4.3. Operational Analysis DIF Comparison Groups

DIF analyses were conducted on each test form for designated comparison groups defined on the basis of demographic variables including gender, race/ethnicity, economic disadvantage, and special instructional needs such as students with disabilities (SWD) or English learners (EL). Student demographic information was provided by the state and stored in PearsonAccess^{next} by means of a student data upload. The demographic data was verified by the state prior to score reporting. These comparison groups are specified in Table 8.6.

Table 8.6 DIF Comparison Groups

Grouping Variable	Focal Group	Reference Group
Gender	Female	Male
Ethnicity	American Indian/Alaska Native (AmerIndian)	White
	Asian	White
	Black or African American	White
	Hispanic/Latino	White
	Native Hawaiian or Pacific Islander	White
	Multiple Race Selected	White
Economic Status*	Economically Disadvantaged (EcnDis)	Not Economically Disadvantaged (NoEcnDis)
Special Instructional Needs	English Learner (ELY)	Non-English Learner (ELN)
	Students with Disabilities (SWDY)	Students without Disabilities (SWDN)

Note: *Economic status was based on participation in National School Lunch Program (receipt of free or reduced-price lunch).

DIF analyses were conducted when the following sample size requirements were met:

- The smaller group, reference or focal, had at least 100 students, and
- The combined group, reference and focal, had at least 400 students.

8.4.4. Operational DIF Results

Results for the ELAGP are reported in Table 8.7, and results for the MATGP are reported in Table 8.8. For the ELAGP, DIF is calculated for PCR items at the item level rather than for the two writing traits which contribute to PCR scoring. There were a relatively large number of ELN versus ELY C-DIF items identified in the ELAGP form.

Table 8.7 Post-Administration Differential Item Functioning for ELAGP

DIF Comparison	N Items	C- DIF		B- DIF		A DIF		B+ DIF		C+ DIF	
		N	%	N	%	N	%	N	%	N	%
Female vs. Male	20	1	5	1	5	18	90				
White vs. Black/African American	20					20	100				
White vs. Hispanic/Latino	20					20	100				
White vs. Asian	20					20	100				
White vs. American Indian/ Alaska Native	20					20	100				
White vs. Native Hawaiian or Pacific Islander	20					20	100				
White vs. Two or more races	20					20	100				
Not Economically Disadvantaged vs. Economically Disadvantaged	20					20	100				
Non-English Learner vs. English Learner	20	5	25	3	15	12	60				
Student without Disability vs. Student with Disability	20					20	100				

Table 8.8 Post-Administration Differential Item Functioning for MATGP

DIF Comparison	N Items	C- DIF		B- DIF		A DIF		B+ DIF		C+ DIF	
		N	%	N	%	N	%	N	%	N	%
Female vs. Male	30			1	3	29	97				
White vs. Black or African American	30					30	100				
White vs. Hispanic/Latino	30					30	100				
White vs. Asian	30					26	87	3	10	1	3
White vs. American Indian/ Alaska Native	30			1	3	29	97				

DIF Comparison	N Items	C- DIF		B- DIF		A DIF		B+ DIF		C+ DIF	
		N	%	N	%	N	%	N	%	N	%
White vs. Native Hawaiian or Pacific Islander	30			1	3	29	97				
White vs. Two or more races	30					30	100				
Not Economically Disadvantaged vs. Economically Disadvantaged	30					30	100				
Non-English Learner vs. English Learner	30	1	3	2	7	26	87	1	3		
Student without Disability vs. Student with Disability	30					30	100				

Chapter 9. Reliability

9.1. Overview

Reliability focuses on the extent to which score differences reflect true differences in the knowledge, skills, and abilities being assessed rather than chance fluctuations. Thus, reliability measures the consistency of the scores across conditions that can be assumed to differ at random; for example, which form of the assessment the student is administered or which raters are assigned to score constructed-response prompts. In statistical terms, the variance in the distributions of scores, essentially the differences among students, is partly due to real differences in the knowledge, skills, and abilities being assessed (true variance) and partly due to random errors in the measurement process (error variance). Reliability is an estimate of the proportion of the total variance that is true variance.

There are several different ways to estimate reliability. The type of raw score reliability estimate reported here is an internal-consistency measure, which is derived from analysis of the consistency of the performance of students across items within an assessment. It is used because it serves as a good estimate of alternate forms reliability, but it does not consider form-to-form variation due to lack of test form parallelism, nor is it responsive to day-to-day variation due to, for example, the student's state of health or the administration environment.

Reliability coefficients range from 0 to 1. The higher the reliability coefficient for a set of scores, the more likely students would be to obtain very similar scores upon repeated administrations if the students do not change in their level of the knowledge or skills measured by the assessment. Acceptable ranges of reliability tend to exceed 0.6, with values over 0.8 considered good to excellent (Cortina, 1993; Schmitt, 1996). Estimates lower than 0.5 may indicate a lack of internal consistency.

The CTT-based standard error of measurement (SEM) quantifies the amount of error in the scores. SEM is the extent by which students' scores tend to differ from the scores they would receive if the test were perfectly reliable. As the SEM increases, the variability of students' observed scores is likely to increase across repeated administrations. Observed scores with large SEMs pose a challenge to the valid interpretation of a single score. Reliability and SEM estimates were calculated at the full assessment level.

9.2. Reliability and SEM Estimation

Coefficient alpha (Cronbach, 1951) is a reliability measure for dichotomously or polytomously scored items (Brennan, 2001). The coefficient is calculated by substituting the variance of both items and total raw scores as follows:

Equation 9.1

$$\alpha_x = \frac{n}{n-1} \left(1 - \frac{\sum_{i=1}^n \sigma_i^2}{\sigma_x^2} \right)$$

in which n is the number of items, σ_i^2 is the variance of scores on each item, and σ_x^2 is the variance of the total raw score. When other administration constraints are held constant, the more items the assessment includes, the greater the reliability coefficient. Conversely, when sample sizes become smaller and more homogeneous, lower reliability estimates are obtained.

The formula for the classical SEM given as:

Equation 9.2

$$SEM = \sigma_x \sqrt{1 - \alpha_x}$$

where σ_x is the standard deviation of the raw score and α_x is the estimated coefficient alpha computed above.

9.2.1. Scale Score Reliability Estimation

Like the stratified alpha coefficients, scale score reliability coefficients range from 0 to 1. The higher the reliability coefficient for a set of scores, the more likely individuals would be to obtain similar scores upon repeated testing occasions, if the students do not change in their level of knowledge or skills measured by the test. Because the scale scores are computed from a total score and do not have an item-level component, a stratified alpha coefficient cannot be computed for scale scores. Instead, Kolen et al.'s (1996) method for scale score reliability was used.

The general formula for a reliability coefficient,

Equation 9.3

$$\rho = 1 - \frac{\sigma^2(E)}{\sigma^2(X)}$$

involves the error variance, $\sigma^2(E)$ and the total score variance, $\sigma^2(X)$. Using Kolen et al.'s (1996) method, conditional raw score distributions are estimated using Lord and Wingersky's (1984) recursion formula. The conditional raw score distributions are transformed into conditional scale score distributions. Denote X as the raw sum score ranging from 0 to X , and s as a resulting scale score after transformation. The conditional distribution of scale scores is written as $P(X = x|\theta)$. The mean and variance, $\sigma^2[s(X)]$, of this distribution can be computed using these scores and their associated probabilities.

The average error variance of the scale scores is computed as

Equation 9.4

$$\sigma^2(Error_{scale}) = \int_{\theta} \sigma^2(s(X)|\theta) g(\theta) d\theta$$

where $g(\theta)$ is the ability distribution. The square root of the error variance is the conditional standard error of measurement of the scale scores.

Just as the reliability of raw scores is one minus the ratio of error variance to total variance, the reliability of scale scores is one minus the ratio of the average variance of measurement error for scale scores to the total variance of scale scores.

Equation 9.5

$$\rho_{scale} = 1 - \frac{\sigma^2(Error_{scale})}{\sigma^2[s(X)]}$$

The Windows program POLYCSEM (Kolen, 2004) was used to estimate scale score error variance and reliability.

9.3. Reliability Results

9.3.1. Raw Score Reliability Results

Tables 9.1 and 9.2 summarize test reliability estimates for the total testing group for ELA and mathematics, respectively. Please note that during operational test form construction (described in section 2.6 of this technical report), multiple parallel operational forms of the accommodated core form may be constructed to facilitate delivery and scoring. The tables below report statistics for all operational forms created for the ELAGP and MATGP. Estimates were calculated only for groups of 100 or more students who were administered a specific test form.

Table 9.1 Summary of ELAGP Test Reliability Estimates for Total Group

Grade Level	Form	Max Score	N Students	SEM	Alpha
ELAGP	OL	74	94,338	5.31	0.89
	AC1	74	3,406	4.49	0.86
	AC2	74	175	5.11	0.89

Table 9.1 provides a summary of the test reliability estimates for the ELAGP assessment, specifically focusing on different forms. The forms include OL (Online English), AC1 (Accommodated Form 1), and AC2 (Accommodated Form 2).

For the OL form, the maximum score is 74, with 94,338 students being assessed. The SEM for this form is 5.31, and the reliability coefficient (Alpha) is 0.89.

Moving on to AC1, the maximum score remains at 74, but the number of students assessed is 3,406. The SEM for this form is 4.49, and the reliability coefficient (Alpha) is 0.86. For the AC2 form, again with a maximum score of 74, there are 175 students assessed. The SEM for this form is 5.11, and the reliability coefficient (Alpha) is 0.89.

Table 9.2 Summary of MATGP Test Reliability Estimates for Total Group

Grade Level	Form	Max Score	N Students	SEM	Alpha
MATGP	OL	55	86,131	3.42	0.91
	AC1	55	10,392	2.95	0.91
	AC2	55	181	3.14	0.89
	Spanish	55	1,820	2.27	0.64

Table 9.2 provides a summary of the test reliability estimates for the ELAGP, broken down by different forms. The OL form, with a maximum score of 55, was taken by 86,131 students, and the SEM was 3.42 with an alpha reliability coefficient of 0.91.

The AC1 also had a maximum score of 55, with 10,392 students. The SEM for this form was 2.95, and the alpha reliability coefficient was 0.91. The AC2 also had a maximum score of 55, with 181 students. The SEM was 3.14 with an alpha reliability coefficient of 0.89.

The Spanish form, with a maximum score of 55, was taken by 1,820 students. The SEM for this form was 2.27, and the alpha reliability coefficient was 0.64.

9.3.2. Scale Score Reliability Results

Tables 9.3 and 9.4 report the scale score reliability and SEM estimates for ELAGP and MATGP, respectively, for spring 2022.

Table 9.3 ELAGP Scale Score Reliability

Grade Level	Form	Max Score	N Students	SEM	Alpha
ELAGP	OL	74	94,338	13.65	.96
	AC1	74	3,406	13.03	.93
	AC2	74	175	13.43	.95

Table 9.4 MATGP Scale Score Reliability Results

Grade Level	Form	Max Score	N Students	SEM	Alpha
MATGP	OL	55	86,131	12.29	.92
	AC1	55	10,392	11.34	.88
	AC2	55	181	13.02	.88
	Spanish	55	1,820	12.79	.60

9.4. Reliability Results for Demographic Groups of Interest

Raw score reliability statistics were also calculated for student demographic groups. The results for these gender, ethnic and student needs groups, along with similar statistics calculated for the students taking various accommodated forms, are reported in table 9.3 for the ELAGP and 9.4 for the MATGP.

Table 9.5 presents test reliability estimates for various subgroups in the context of the ELAGP assessment. All forms had a maximum score of 74. The OL Form was administered to a total of 94,338 students. The SEM is reported as 5.31, and the Cronbach's Alpha reliability coefficient is 0.89 for the overall test. When examining gender-based subgroups, results for the 46,698 female students and 47,497 male students demonstrated similar reliability, with SEM values of 5.27 and 5.22, and Alpha coefficients of 0.88 and 0.90, respectively. The reliability estimates are also presented based on ethnicity, socioeconomic status, presence of disabilities, English learner status and accommodation forms.

In terms of ethnicity, variations in reliability are observed across subgroups, such as American Indian/Alaska Native, Asian, Black or African American, Hispanic/Latino, Native Hawaiian or Pacific Islander, Two or More Races, and White. Additionally, the table provides reliability estimates for students based on their economic disadvantage status, disability status, English learner status, and the type of accommodation they received. Notably, students taking accommodated forms also have reliability estimates reported. This includes students using American Sign Language, closed captioning, a human reader, a non-screen reader, a screen reader, paper, or text-to-speech. The comprehensive nature of these reliability estimates allows for a detailed understanding of the assessment's performance across diverse student subgroups and testing conditions.

Table 9.5 Summary of Test Reliability Estimates for Subgroups: ELAGP

	Max Score	N Students	SEM	Alpha
ELAGPTotal Component	74	94,338	5.31	0.89
Gender	Max Score	N Students	SEM	Alpha
Female	74	46,698	5.27	0.88
Male	74	47,497	5.22	0.90
Ethnicity	Max Score	N Students	SEM	Alpha
American Indian/Alaska Native	74	139	5.55	0.90
Asian	74	10,329	5.01	0.86
Black or African American	74	12,851	5.11	0.88
Hispanic/Latino	74	28,349	5.29	0.89
Native Hawaiian or Pacific Islander	74	312	5.27	0.89
Two or More Races	74	2,140	5.20	0.88
White	74	39,610	5.19	0.87
Students with instructional needs	Max Score	N Students	SEM	Alpha
Not Economically Disadvantaged	74	67,023	5.26	0.88
Economically Disadvantaged	74	27,315	5.25	0.89
Students without Disabilities	74	78,149	5.26	0.88
Students with Disabilities (SWD)	74	16,189	5.09	0.89
Non-English Learner	74	89,506	5.27	0.88
English Learner	74	4,832	3.85	0.83
Students Taking Accommodated Forms	Max Score	N Students	SEM	Alpha
American Sign Language	74	19	n/r	n/r
Closed Captioning	74	54	n/r	n/r
Human Reader	74	29	n/r	n/r
Non-Screen Reader	74	15	n/r	n/r
Screen Reader	74	3	n/r	n/r
Paper	74	109	4.96	0.89
Text-to-Speech	74	3,352	4.45	0.86

Table 9.6 presents test reliability estimates for subgroups in the MATGP. The maximum score for all subgroups is 55.

In terms of gender, both females' (42,773 students) and males' (43,222 students) results exhibit similar reliability measures, with SEM values of 3.41 and 3.42, respectively. The alpha coefficients for both groups are 0.91 and 0.92, respectively.

There are variations in the reliability of scores for various ethnicities. Asian students (9,758) have a higher SEM of 3.83, while Black or African American students (11,585) have a lower SEM of 2.77. The alpha coefficients range from 0.88 for Black or African American students to 0.92 for American Indian/Alaska Native students.

Considering economic status, economically disadvantaged students (23,729) demonstrate a lower SEM of 2.91 compared to their non-economically disadvantaged counterparts (62,402), whose SEM is 3.54. The alpha coefficients for these groups are 0.88 and 0.91, respectively.

Reliability estimates for the assessment results of students with disabilities (SWD), are slightly lower, with an SEM of 2.94 and an alpha coefficient of 0.90. English learners' results (3,063) exhibit a lower SEM of 2.48 and a notable alpha coefficient of 0.81.

The results of students using various methods of student test accommodation demonstrate different reliability measures, with Text-to-Speech (10,105 students) having an SEM of 2.96 and a high alpha coefficient of 0.91. Spanish-speaking students (1,823) present a notably low SEM of 2.28 but a higher alpha coefficient of 0.64. Reliability estimates are not reported for accommodation methods which were used by fewer than 100 students.

Table 9.6 Summary of Test Reliability Estimates for Subgroups: MATGP

	Max Score	N Students	SEM	Alpha
MATGP Total Component	55	86,131	3.42	0.91
Gender	Max Score	N Students	SEM	Alpha
Female	55	42,773	3.41	0.91
Male	55	43,222	3.42	0.92
Ethnicity	Max Score	N Students	SEM	Alpha
American Indian/Alaska Native	55	127	3.31	0.92
Asian	55	9,758	3.83	0.90
Black or African American	55	11,585	2.77	0.88
Hispanic/Latino	55	24,490	2.97	0.89
Native Hawaiian or Pacific Islander	55	297	3.47	0.90
Two or More Races	55	2,009	3.53	0.91
White	55	37,310	3.49	0.90
Students with instructional needs	Max Score	N Students	SEM	Alpha
Not Economically Disadvantaged	55	62,402	3.54	0.91
Economically Disadvantaged	55	23,729	2.91	0.88
Students without Disabilities	55	71,520	3.49	0.91
Students with Disabilities (SWD)	55	14,611	2.94	0.90
Non-English Learner	55	83,068	3.44	0.91
English Learner	55	3,063	2.48	0.81
Students Taking Accommodated Forms	Max Score	N Students	SEM	Alpha
American Sign Language	55	16	n/r	n/r
Human Reader	55	30	n/r	n/r
Non-Screen Reader	55	8	n/r	n/r
Screen Reader	55	3	n/r	n/r
Paper	55	121	3.11	0.90
Text-to-Speech	55	10,105	2.96	0.91
Spanish	55	1,823	2.28	0.64

9.5. Reliability Estimates of Subclaim Scores

Table 9.7 presents the average reliability estimates for various ELAGP subclaim scores across different forms, providing insights into the consistency and dependability of these scores. The table includes subclaim scores for Reading (RD), Reading: Literature (RL), Reading: Information (RI), Reading: Vocabulary (RV), Writing (WR), Written Expression (WE), and Knowledge of Language and Conventions (WKL).

For Form OL, the reliability estimates (Alpha) for the subclaim scores range from 0.60 to 0.87. Specifically, Reading (RD) has an Alpha of 0.87, indicating high reliability, while Reading: Information (RI) has an Alpha of 0.74, suggesting good reliability. Other subclaim scores, such as Reading: Vocabulary (RV) and Written Expression (WE), also demonstrate reliability, with Alphas of 0.83 and 0.84, respectively.

Similarly, Form AC1 and AC2 exhibit varying reliability estimates for the subclaim scores. Reading (RD) consistently shows relatively high reliability across forms. Overall, the table provides a comprehensive overview of the average reliability estimates for MATGP subclaim scores across different forms, offering valuable information for evaluating the consistency and dependability of the assessments in various academic domains.

Table 9.7 ELAGP Reliability Estimates for Subscores

	RD		RL		RI		RV		WR		WE		WKL	
Form	Score Points	Alpha	Score Points	Alpha	Score Points	Alpha	Score Points	Alpha	Score Points	Alpha	Score Points	Alpha	Score Points	Alpha
OL	44	0.87	12	0.69	22	0.74	10	0.60	30	0.83	24	0.84	6	0.83
AC1	44	0.81	12	0.64	22	0.61	10	0.52	30	0.81	24	0.78	6	0.78
AC2	44	0.87	12	0.71	20	0.78	12	0.61	30	0.85	24	0.89	6	0.90

Table 9.8 presents MATGP reliability estimates for different subscores across various forms. The subscores include Major Content (MC), Additional & Supporting Content (ASC), Mathematics Reasoning (MR), and Modeling Practice (MP). Each form is associated with corresponding score points and alpha reliability coefficients for the mentioned subscores.

For form OL, MC has an average score of 14 points with an alpha reliability coefficient of 0.80. ASC has 16 points with an alpha of 0.78, Mathematics Reasoning (MR) has 10 points with an alpha of 0.72, and MP has 15 points with an alpha of 0.76.

Similarly, forms AC1 and AC2 exhibit alpha coefficients for each subscore similar to those of the OL form. Notably, the Spanish form shows lower reliability estimates across all subscores compared to the other forms, with alpha coefficients ranging from 0.25 to 0.41.

In summary, the table details the average MATGP reliability estimates for different subscores across distinct forms, highlighting variations in scores and alpha coefficients, with Spanish displaying comparatively lower reliability.

Table 9.8 MATGP Reliability Estimates for Subscores

	MC		ASC		MR		MP	
Form	Score Points	Alpha	Score Points	Alpha	Score Points	Alpha	Score Points	Alpha
OL	14	0.80	16	0.78	10	0.72	15	0.76
AC1	14	0.77	16	0.75	10	0.71	15	0.73
AC2	16	0.80	14	0.69	10	0.77	15	0.68
Spanish	14	0.41	16	0.34	10	0.25	15	0.39

9.6. Reliability of Classification

Classification accuracy is defined as the extent to which the actual classifications of test takers (on the basis of their test scores) agree with those that would be made on the basis of their true scores—if their true scores could somehow be known. The term consistency refers to the agreement between classifications based on two nonoverlapping, equally difficult forms of the test (parallel forms) (Livingston & Lewis, 1995).

We used Livingston and Lewis's (1995) approach, which is intended to handle situations where items are not equally weighted and/or some or all the items are polytomously scored. This method is formulated as

Equation 9.6

$$\tilde{n} = \frac{(\mu_x - X_{min})(X_{max} - \mu_x) - r\sigma_x^2}{\sigma_x^2(1-r)},$$

where X_{min} is the lowest score for X , X_{max} is the highest score, μ_x is the mean, σ_x^2 is the variance, and r is the reliability. This method models the distribution of the true scores and of scores on a parallel form by using a four-parameter beta distribution.

As seen in the above formula, classification accuracy and consistency indices rely on the interaction between several different factors related to test design and standard-setting decisions. These factors include the number of score cuts, test reliability, measurement accuracy at the cut score, distance between adjacent cuts, location of the cut scores on the ability scale, and percentage of students around a cut score (Ercikan & Julian, 2002; Lee, Hanson, & Brennan, 2002). Because these statistics are influenced by the interplay between a variety of factors, only a very limited number of studies to date have investigated the ideal or expected levels of decision consistency and accuracy needed for educational assessments.

Classification accuracy indices quantify the percentage of students who are accurately placed below and/or above a given cut score. For example, a classification accuracy index of 0.886 for cut score 1/2 means that were students to be classified twice, once according to their observed score and once according to their true score, 88% of those students would be classified in the same category both times.

Similarly, classification consistency indices give the percentage of students classified consistently below and/or above a given cut score. For example, a classification index of 0.840 for cut score 1/2 means that were two parallel forms to be administered to students, 84% of those students would be classified in the same way for both forms.

Table 9.9 reports the decision accuracy and consistency for classifying each student into one of the two performance levels for ELA. The estimated proportion of those accurately classified is 0.92, and the proportion who would be consistently classified on two different test forms is 0.88.

Table 9.9 Reliability of Classification: Summary for ELAGP

	Decision Accuracy: Proportion Accurately Classified	Decision Consistency: Proportion Consistently Classified
Level	Exact Level	Exact Level
ELAGP	0.92	0.88

Table 9.10 provides more detailed information about the accuracy and consistency of the classification of students into performance levels for ELA. Each cell shows the estimated proportion of students who would be classified into a particular combination of performance levels. The sum of the two shaded values on the diagonal is approximately equal to the level of decision accuracy or consistency presented in Table 9.9. Note that the sums based on values in Table 9.10 may not match the values in Table 9.9 due to truncation and rounding.

Table 9.10 Reliability of Classification: ELAGP

	Full Summative Scale Score	Not Yet Graduation Ready	Graduation Ready	Category Total
Decision Accuracy	650–749	0.56	0.040	0.60
	750–850	0.04	0.36	0.40
Decision Consistency	650–749	0.54	0.057	0.60
	750–850	0.062	0.34	0.40

Table 9.11 shows reliability for classifying each student into one of the two performance levels for MATGP. The estimated proportion of those accurately classified is 0.92, and the proportion who would be consistently classified on two different test forms is 0.89.

Table 9.11 Reliability of Classification: Summary for MATGP

	Decision Accuracy: Proportion Accurately Classified	Decision Consistency: Proportion Consistently Classified
Level	Exact Level	Exact Level
MATGP	0.92	0.89

Table 9.12 provides more detailed information about the accuracy and consistency of the classification of students into performance levels for MATGP. Each cell shows the estimated proportion of students who

would be classified into a particular combination of performance levels. The sum of the two shaded values on the diagonal is approximately equal to the level of decision accuracy or consistency presented in Table 9.9. Note that the sums based on values in Table 9.9 may not match the values in Table 9.9 due to truncation and rounding.

Table 9.12 Reliability of Classification: MATGP

	Full Summative Scale Score	Not Yet Graduation Ready	Graduation Ready	Category Total
Decision Accuracy	650–749	0.47	0.04	0.51
	750–850	0.037	0.46	0.49
Decision Consistency	650–749	0.45	0.054	0.50
	750–850	0.054	0.44	0.50

Chapter 10. Validity

10.1. Overview

The term validity addresses the question of how well a test measures what it is designed to measure. A good description of assessment validity is given in the *Standards for Educational and Psychological Testing*, issued jointly by the American Educational Research Association (AERA), American Psychological Association (APA), and National Council on Measurement in Education (NCME) (2014). The Standards state the following:

Validity refers to the degree to which evidence and theory support the interpretations of test scores for proposed uses of tests. Validity is, therefore, the most fundamental consideration in developing tests and evaluating tests. The process of validation involves accumulating relevant evidence to provide a sound scientific basis for the proposed score interpretations (p. 11).

The purpose of test validation is not to validate the test itself but to validate interpretations of the test scores for particular uses. Test validation is not a quantifiable property but an ongoing process, beginning at initial conceptualization and continuing throughout the lifetime of an assessment. Every aspect of an assessment provides evidence in support of its validity (or evidence of lack of validity), including design, content specifications, item development, and psychometric characteristics. The spring 2022–2023 NJGPA administration provided an opportunity to gather evidence of validity based on both test content and on the internal structure of the tests.

NMC applies the principles of universal design, as articulated in materials developed by the National Center for Educational Outcomes (NCEO) at the University of Minnesota (Thompson et al., 2002).

10.2. Evidence Based on Test Content

This section describes evidence of validity from the creation of items used on NJGPA forms which come from PARCC/NMC Affiliate ELA Grade 10, Algebra I and Geometry banks.

Evidence based on content of achievement tests is supported by the strength and clarity of the relationship between test items and content standards. The degree to which the test measures what it claims to measure is known as construct validity. The NJGPA adheres to the principles of evidence-centered design, in which the standards to be measured are identified and the performance a student needs to achieve to meet those standards is delineated in the evidence statements. Test items are reviewed for adherence to universal design principles, which maximize the participation of the widest possible range of students.

During the creation of ELA grade 10, Algebra I and Geometry items, Pearson built spreadsheets at the evidence statement level that incorporate the probability statements from the test blueprints and attrition rates at committee review and data review. The basis of our entire item development is driven by the use of these item development target spreadsheets. Before item development begins, these target spreadsheets are used to develop an internal item writing plan to correlate with the expectations of the test design. These were reviewed and approved by state leads and NMC.

In addition to the evidence statements, content was aligned through the articulation of performance in the performance level descriptors. At the policy level, the performance level descriptors include policy claims

about the educational achievement of students who attain a particular performance level. They also include a broad description of the grade-level knowledge, skills, and practices students performing at a particular achievement level are able to demonstrate. Those policy-level descriptors were the foundation for the subject- and grade-specific performance level descriptors, which, along with the evidence frameworks, guided the development of the items and tasks.

The college- and career-ready determinations (CCRD) in English language arts (ELA) and mathematics describe the academic knowledge, skills and practices students must demonstrate to show readiness for success in entry-level, credit-bearing college courses and relevant technical courses. To validate the determinations, a postsecondary educator judgment study and a benchmark study were conducted of the SAT, ACT, National Assessment of Educational Progress (NAEP), Trends in International Mathematics and Science Study (TIMSS), Programme of International Student Assessment (PISA), and Progress in International Reading Literacy Study (PIRLS) tests (McClarty et al., 2015).

Gathering construct validity evidence for the assessments was embedded in the process by which the assessment content was developed. At each step in the assessment development process, participating state educators, assessment experts, and bias and sensitivity experts reviewed texts, items, and tasks for accuracy, appropriateness, and to eliminate bias.

Please see Chapter 2 for an overview of the operational forms and the form development process. For more information on the origins of items, forms and test blueprints, please consult chapter 10 of the most recent New Meridian [Affiliate Technical Report Alternate Blueprint](#). The items and tasks were field tested prior to their use on an assessment. Items demonstrating questionable statistical performance, either within the testing population as a whole or within demographic groups, are reviewed at data review meetings.

10.3. Evidence Based on Internal Structure

Analyses of the internal structure of a test typically involve studies of the relationships among test items and/or test components (i.e., subclaims) in the interest of establishing the degree to which the items or components appear to reflect the construct on which a test score interpretation is based (AERA et al., 2014, p. 16). The term construct is used here to refer to the characteristics that a test is intended to measure; in the case of the NJGPA, the characteristics of interest are the knowledge and skills defined by the test blueprint for the ELA and mathematics components.

NJGPA scoring provides the overall ELA and mathematics test scores, the Reading claim score, and the Writing claim score as well as ELA subclaim and mathematics subclaim scores. The goal of reporting at this level is to provide criterion-referenced data to assess the strengths and weaknesses of a student's achievement in specific topics for each test component.

10.3.1. Intercorrelations

The ELAGP comprise two claim scores, Reading (RD) and Writing (WR), and five subclaim scores: Reading Literature (RL), Reading Information (RI), Reading Vocabulary (RV), Written Expression (WE), and Knowledge of Language and Conventions (WKL). The RD claim score is a composite of RL, RI, and RV. The writing claim score, a composite of WE and WKL, comprises only PCR items, and the same PCR items are in each subclaim. The ELA operational test analyses were performed by evaluating the separate trait scores of WE and WKL, and for some PCR items also RL or RI; therefore, the trait scores were used for the intercorrelations.

The MATGP has four subclaim scores—Major Content (MC), Mathematical Reasoning (MR), Modeling Practice (MP), and Additional and Supporting Content (ASC).

High total group internal consistencies as well as similar reliabilities across subgroups provide additional evidence of validity. High reliability of test scores implies that the test items within a domain are measuring a single construct, which is a necessary condition for validity when the intention is to measure a single construct. Refer to Chapter 9 for reliability estimates for the overall population, subgroups of interest, as well as for claims and subclaims for ELA and subclaims for mathematics.

Another way to assess the internal structure of a test is through the evaluation of correlations among scores. These analyses were conducted between the ELA Reading and Writing claim scores and the ELA subclaims (RL, RI, RV, WE, and WKL) and between the mathematics subclaims. If these components within a content area are strongly related to each other, this is evidence of unidimensionality.

A series of tables are provided to summarize the results for the spring 2023 administration. Table 10.1 presents the correlations observed between the ELA Reading and Writing subclaim scores for ELA. The intercorrelations for mathematics are provided in Table 10.2.

Table 10.1 Average Intercorrelations between ELAGP Subclaims

Subclaims	N							
	Students	RD	RL	RI	RV	WR	WE	WKL
RD	97,919	1						
RL	97,919	0.87	1					
RI	97,919	0.94	0.72	1				
RV	97,919	0.83	0.63	0.67	1			
WR	97,919	0.79	0.77	0.74	0.56	1		
WE	97,919	0.78	0.76	0.74	0.56	1.00	1	
WKL	97,919	0.78	0.76	0.73	0.57	0.97	0.95	1

Note. RD=Reading, RL=Reading: Literature, RI=Reading: Information, RV=Reading: Vocabulary, WR=Writing, WE=Written Expression, WKL=Knowledge of Language and Conventions

ELA intercorrelations range from fair to good. The intercorrelations of reading subclaim scores range from .67 to .94 and likely indicate measurement of a single reading construct, with Reading Vocabulary possibly being a related construct. The intercorrelations of writing sub scores range from .95 to 1.0, indicating there is likely only one writing dimension being measured. The intercorrelations between reading and writing sub scores, with the exception of Reading Vocabulary, likely indicate a unidimensional measurement factor.

Table 10.2 Average Intercorrelations between MATGP Subclaims

Subclaims	N Students	MC	ASC	MR	MP
MC	98,524	1			
ASC	98,524	0.72	1		
MR	98,524	0.77	0.79	1	
MP	98,524	0.81	0.72	0.77	1

MC=Major Content, ASC=Additional & Supporting Content, MR=Mathematics Reasoning, MP=Modeling Practice.

The mathematics intercorrelations are moderate. The main observable pattern in the mathematics intercorrelations is that the Major Content subclaim generally has slightly higher correlations with the Additional and Supporting Content, Mathematical Reasoning, and Modeling Practice subclaims; the intercorrelations among the Additional and Supporting Content, Mathematical Reasoning, and Modeling Practice subclaims are usually slightly lower. The mathematics intercorrelations are sufficiently high to suggest that the mathematics tests are likely to be unidimensional with some minor secondary dimensions.

10.3.2. Reliability

The reliability analyses presented in Chapter 9 of this technical report provide information about the internal consistency of the NJGPA components. Internal consistency is typically measured via correlations amongst the items on an assessment and provides an indication of how much the items measure the same general construct. The reliability estimates, computed using coefficient alpha (Cronbach, 1951), are presented in Tables 9.1 and 9.2. The reliabilities for the ELAGP forms range from .86 to .89, and from .64 to .91 for the MATGP. Table 9.5 summarizes test reliability for groups of interest for the ELAGP, and Table 9.6 summarizes test reliability for groups of interest for MATGP. Along with the subclaim intercorrelations, the reliability estimates indicate that the items within each assessment are measuring the same construct and provide further evidence of unidimensionality.

10.3.3. Local Item Dependence

Local item independence was evaluated for ELA and mathematics. Local independence is one of the primary assumptions of item response theory (IRT) that states the probability of success on one item is not influenced by performance on other items when controlling for ability level. This implies that ability or theta accounts for the associations among the observed items. Local item dependence (LID) when present essentially overstates the amount of information predicted by the IRT model. It can exert other undesirable psychometric effects and represents a threat to validity since other factors besides the construct of interest are present. Classical statistics are also affected when LID is present because estimates of test reliability like IRT information can be inflated (Zenisky et al., 2003).

The LID issue affects the choice of item scoring in IRT calibrations. Specifically, if evidence suggests these items indeed have local dependence, then it might be preferable to sum the item scores into clusters or

testlets as a method of minimizing LID. However, if these items do not appear to have strong local item dependence, then retaining the scores as individual item scores in an IRT calibration is preferred since more information concerning item properties is retained. During the initial operational administration of the summative assessments in spring 2015, a study that included two methods of investigating the presence of LID was conducted. A description of the study methods and findings is available in Chapter 10 of the most recent New Meridian affiliate technical report.

10.3.3.1. Evidence from Special Studies

During the creation and evolution of the testing programs which inform ELA grade 10, Algebra I, and Geometry, several research studies provided additional validity evidence for the goals of assessing more rigorous academic expectations, helping to prepare students for college and careers and providing information back to teachers and parents about their students' progress toward college and career readiness. Details of these studies can be found in the most recent version of the New Meridian affiliate technical report.

References

- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education (2014). *Standards for educational and psychological testing*. Washington, DC: American Educational Research Association.
- Baldwin, P., Margolis, M.J., Clauser, B.E., Mee, J., Winward, M. (2019). The choice of response probability in bookmark standard setting: An experimental study. *Educational Measurement: Issues and Practice*, 39(1), 37–44.
- Beimers, J. N., Way, W. D., McClarty, K. L., & Miles, J. A. (2012). Evidence based standard setting: Establishing cut scores by integrating research evidence with expert content judgments. *Bulletin*, Issue 21. Pearson Education, Inc.
- Birnbaum, A. (1968). Some latent trait models and their use in inferring an examinee's ability. In F.M. Lord & M.R. Novick, *Statistical Theories of Mental Test Scores* (pp 397–472). Reading, MA: Addison Wesley Publishing.
- Brennan, R. L. (2001). An essay on the history and future of reliability from the perspective of replications. *Journal of Educational Measurement*, 38(4), 295–317.
- Cizek, G. & Bunch, M. (2007). *Standard setting: A guide to establishing performance standards on tests*. Thousand Oaks, CA: Sage Publications.
- Cortina, J. M. (1993). What is coefficient alpha? An examination of theory and applications. *Journal of Applied Psychology*, 78(1), 98–104.
- Crocker, L., & Algina, J. (1986). *Introduction to classical and modern test theory*. New York: Holt, Rinehart and Winston.
- Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, 16(3), 297–334.
- Dorans, N. J. (2013). *ETS contributions to the quantitative assessment of item, test and score fairness (ETS R&D Science and Policy Contributions Series, ETS SPC-13-04)*. Princeton, NJ: Educational Testing Service.
- Dorans, N. J. & Holland, P. W. (1993). DIF detection and description: Mantel-Haenszel and standardization. In P. W. Holland & H. Wainer (Eds.), *Differential item functioning* (pp. 35–66). Hillsdale, NJ: Erlbaum.
- Dorans, N. J., & Schmitt, A. P. (1991). *Constructed response and differential item functioning: A pragmatic approach. (ETS Research Report No. 91-47)*. Princeton, NJ: Educational Testing Service
- Ercikan, K. & Julian, M. (2002). Classification accuracy of assigning student performance to proficiency levels: Guidelines for assessment design. *Applied Measurement in Education*, 15(3), 269–294.
- Henrysson, S. (1963). Correction of Item-Total Correlations in item analysis. *Psychometrika*, 28(2), 211–218.
- Holland, P. W. & Thayer, D. T. (1988). Differential item performances and the Mantel-Haenszel procedure. In H. Wainer & H. Braun (Eds.), *Test validity* (pp. 129–145). Hillsdale, NJ: Erlbaum.
- Kolen, M. J. (2004). POLYSEM windows console version [Computer software]. Iowa City, IA: The Center for Advanced Studies in Measurement and Assessment (CASMA), The University of Iowa.
- Kolen, M. J., Zeng, L., & Hanson, B. A. (1996). Conditional standard errors of measurement for scale scores using IRT. *Journal of Educational Measurement*, 33(2), 129–140.
- Lee, W., Hanson, B. A., & Brennan, R. L. (2002). Estimating consistency and accuracy indices for multiple classifications. *Applied Psychological Measurement*, 26, 412–432.
- Lemke, E., & Wiersma, W. (1976). *Principles of psychological measurement*. Chicago, Ill: McNally.
- Lewis, D., Mitzel, H., & Green, D. (1996, June). Standard setting: A bookmark approach. In D. Green (Chair), *IRT-based standard setting procedures utilizing behavioral anchoring*. Symposium conducted at the Council of Chief State School Officers National Conference on Large-Scale Assessment, Phoenix, AZ.

- Livingston, S. A., & Lewis, C. (1995). Estimating the consistency and accuracy of classifications based on test scores. *Journal of Educational Measurement*, 32, 179–197.
- Lord, F. M., & Wingersky, M. S. (1984). Comparison of IRT true score and equipercentile observed-score “equatings.” *Applied Psychological Measurement*, 8, 453–461.
- Mantel, N. (1963). Chi-square tests with one degree of freedom: Extensions of the Mantel-Haenszel procedure. *Journal of the American Statistical Association*, 58(303), 690–700.
- Mantel, N. & Haenszel, W. (1959). Statistical aspects of the analysis of data from retrospective studies of disease. *Journal of the National Cancer Institute*, 22, 719–748.
- McClarty, K. L., Korbin, J. L., Moyer, E., Griffin, S., Huth, K., Carey, S., & Medberry, S. (2015). *PARCC benchmarking study*. Pearson Educational Measurement, Iowa City, IA: Pearson.
- Mitzel, H., Lewis, D., Patz, R., & Green, D. (2001). The Bookmark procedure: Psychological perspectives. In G. Cizek (Ed.), *Setting performance standards: Concepts, methods and perspectives* (pp. 249–281). Mahwah, NJ: Erlbaum.
- Muraki, E. (1992). A generalized partial credit model: Application of an EM algorithm. *Applied Psychological Measurement* 16(2), 159–176.
- Plake, B. S., Ferdous, A. A., Impara, J. C., & Buckendahl, C. W. (2005). Setting multiple performance standards using the Yes/No method: An alternative item mapping method. *Meeting of the National Council on Measurement in Education*. Montreal. Canada.
- Schmitt, N. (1996). Uses and abuses of coefficient alpha. *Psychological Assessment*, 8(4), 350–353.
- Thompson, S. J., Johnstone, C. J., & Thurlow, M. L. (2002). *Universal design applied to large scale assessments*. Synthesis Report.
- Zenisky, A. L., Hambleton, R. K., & Sireci, S.C. (2003). *Effects of local dependence on the validity of IRT item test, and ability statistics*. (Technical Report). American College Admissions Test.
- Zieky, M. (1993). Practical questions in the use of DIF statistics in test development. In P. Holland & H. Wainer (Eds.), *Differential item functioning* (pp. 337–348). Hillsdale, NJ: Erlbaum
- Zwick, R., Thayer, D. T., & Mazzeo, J. (1997). *Describing and Categorizing DIF in Polytomous Items* (ETS Research Report RR-97-05). Princeton, NJ: Educational Testing Service.

Appendices

Appendices are numbered in reference to the chapter in the body of the text for which they expand or amplify the information reported.

Appendix 6: Classical Item Analysis Statistics by Form

A.6.1 ELAGP Online Form Statistics

Item	P-value	Polyserial
Item 1	0.68	0.63
Item 2	0.38	0.45
Item 3	0.75	0.63
Item 4	0.50	0.51
Item 5	0.41	0.49
Item 6	0.55	0.52
Item 7	0.43	0.47
Item 8	0.53	0.90
Item 9	0.46	0.43
Item 10	0.53	0.55
Item 11	0.75	0.59
Item 12	0.54	0.51
Item 13	0.47	0.69
Item 14	0.69	0.43
Item 15	0.56	0.62
Item 16	0.34	0.93
Item 17	0.61	0.53
Item 18	0.56	0.49
Item 19	0.56	0.58
Item 20	0.40	0.46

A.6.2 ELAGP Accommodated Form 1 Statistics

Item	P-value	Polyserial
Item 1	0.42	0.58
Item 2	0.24	0.35
Item 3	0.51	0.56
Item 4	0.32	0.39
Item 5	0.24	0.38
Item 6	0.33	0.51
Item 7	0.26	0.33
Item 8	0.25	0.87
Item 9	0.30	0.37
Item 10	0.28	0.48
Item 11	0.49	0.59
Item 12	0.31	0.39
Item 13	0.17	0.54
Item 14	0.46	0.49
Item 15	0.29	0.62
Item 16	0.11	0.86
Item 17	0.38	0.43
Item 18	0.33	0.34
Item 19	0.31	0.40
Item 20	0.20	0.30

A.6.1 ELAGP Accommodated Form 2 Statistics

Item	P-value	Polyserial
Item 1	0.45	0.71
Item 2	0.21	0.54
Item 3	0.50	0.47
Item 4	0.35	0.51
Item 5	0.37	0.35
Item 6	0.41	0.64

Item	P-value	Polyserial
Item 7	0.25	0.30
Item 8	0.34	0.50
Item 9	0.19	0.89
Item 10	0.38	0.50
Item 11	0.39	0.38
Item 12	0.32	0.55
Item 13	0.46	0.64
Item 14	0.32	0.50
Item 15	0.48	0.55
Item 16	0.23	0.95
Item 17	0.33	0.67
Item 18	0.24	0.56
Item 19	0.72	0.72
Item 20	0.46	0.72

A.6.2 MATGP Online Form Statistics

Item	P-value	Polyserial
Item 1	0.16	0.74
Item 2	0.53	0.61
Item 3	0.18	0.83
Item 4	0.57	0.79
Item 5	0.54	0.72
Item 6	0.59	0.59
Item 7	0.5	0.6
Item 8	0.6	0.78
Item 9	0.36	0.64
Item 10	0.42	0.54
Item 11	0.64	0.69

Item	P-value	Polyserial
Item 12	0.69	0.48
Item 13	0.21	0.87
Item 14	0.24	0.8
Item 15	0.26	0.5
Item 16	0.05	0.55
Item 17	0.34	0.74
Item 18	0.16	0.79
Item 19	0.56	0.81
Item 20	0.15	0.66
Item 21	0.73	0.75
Item 22	0.45	0.81
Item 23	0.46	0.45
Item 24	0.75	0.74
Item 25	0.33	0.59
Item 26	0.49	0.84
Item 27	0.28	0.65
Item 28	0.57	0.77
Item 29	0.69	0.37
Item 30	0.75	0.58

A.5. 3 MATGP Accommodated Form 1 Statistics

Item	P-value	Polyserial
Item 1	0.09	0.60
Item 2	0.42	0.60
Item 3	0.09	0.77
Item 4	0.41	0.78
Item 5	0.40	0.64
Item 6	0.47	0.55

Item	P-value	Polyserial
Item 7	0.40	0.48
Item 8	0.48	0.74
Item 9	0.23	0.55
Item 10	0.33	0.40
Item 11	0.52	0.74
Item 12	0.60	0.48
Item 13	0.11	0.80
Item 14	0.14	0.77
Item 15	0.20	0.44
Item 16	0.03	0.44
Item 17	0.19	0.72
Item 18	0.08	0.70
Item 19	0.37	0.81
Item 20	0.08	0.63
Item 21	0.55	0.76
Item 22	0.28	0.72
Item 23	0.40	0.37
Item 24	0.59	0.75
Item 25	0.22	0.52
Item 26	0.29	0.84
Item 27	0.19	0.52
Item 28	0.39	0.79
Item 29	0.60	0.40
Item 30	0.64	0.64

A.5. 4 MATGP Accommodated Form 2 Statistics

Item	P-value	Polyserial
Item 1	0.19	0.80

Item	P-value	Polyserial
Item 2	0.39	0.48
Item 3	0.11	0.73
Item 4	0.48	0.82
Item 5	0.22	0.70
Item 6	0.38	0.43
Item 7	0.25	0.80
Item 8	0.32	0.64
Item 9	0.35	0.71
Item 10	0.28	0.44
Item 11	0.57	0.65
Item 12	0.37	0.47
Item 13	0.37	0.41
Item 14	0.04	0.75
Item 15	0.66	0.73
Item 16	0.25	0.43
Item 17	0.50	0.76
Item 18	0.49	0.71
Item 19	0.51	0.76
Item 20	0.05	0.80
Item 21	0.06	0.71
Item 22	0.03	0.76
Item 23	0.08	0.63
Item 24	0.38	0.67
Item 25	0.24	0.48
Item 26	0.11	0.62
Item 27	0.70	0.71
Item 28	0.43	0.64
Item 29	0.29	0.82

Item	P-value	Polyserial
Item 30	0.36	0.51

A.5. 5 MATGP Spanish Form Statistics

Item	P-value	Polyserial
Item 1	0.04	0.27
Item 2	0.29	0.37
Item 3	0.01	0.47
Item 4	0.19	0.41
Item 5	0.38	0.44
Item 6	0.28	0.41
Item 7	0.31	0.24
Item 8	0.31	0.50
Item 9	0.12	0.25
Item 10	0.24	0.26
Item 11	0.40	0.40
Item 12	0.48	0.36
Item 13	0.02	0.41
Item 14	0.04	0.51
Item 15	0.14	0.29
Item 16	0.01	0.19
Item 17	0.02	0.55
Item 18	0.01	0.32
Item 19	0.19	0.48
Item 20	0.01	0.34
Item 21	0.39	0.71
Item 22	0.07	0.39
Item 23	0.35	0.22
Item 24	0.32	0.40

Item	P-value	Polyserial
Item 25	0.14	0.29
Item 26	0.10	0.63
Item 27	0.14	0.21
Item 28	0.10	0.42
Item 29	0.49	0.42
Item 30	0.53	0.42

Appendix 7: Item Response Theory Parameters

IRT parameters are presented for each item used on any operational test form.

A7.1 ELAGP IRT Parameters for Unique Items

Item	Max Points	a	b	d1	d2	d3	d4	d5
Item 1	2	0.39	-0.13	0.00	-3.45	3.45		
Item 2	2	0.38	1.52	0.00	1.17	-1.17		
Item 3	2	0.48	-0.52	0.00	-0.64	0.64		
Item 4	2	0.39	0.70	0.00	-0.18	0.18		
Item 5	2	0.42	1.26	0.00	0.91	-0.91		
Item 6	2	0.42	0.40	0.00	0.30	-0.30		
Item 7	2	0.39	1.16	0.00	0.99	-0.99		
Item 8	16	0.87	0.69	0.00	1.41	0.83	-0.54	-1.71
Item 9	3	0.92	0.07	0.00	0.63	0.18	-0.81	
Item 10	2	0.33	0.97	0.00	1.09	-1.09		
Item 11	2	0.32	0.53	0.00	-3.28	3.28		
Item 12	2	0.48	-0.53	0.00	-0.56	0.56		
Item 13	2	0.28	0.47	0.00	-3.89	3.89		
Item 14	2	0.65	0.79	0.00	-0.07	0.07		
Item 15	2	0.23	-0.61	0.00	-5.69	5.69		
Item 16	2	0.60	0.40	0.00	0.49	-0.49		
Item 17	16	0.99	1.50	0.00	1.71	0.45	-0.66	-1.49

Item	Max Points	a	b	d1	d2	d3	d4	d5
Item 18	3	1.02	0.96	0.00	1.17	-0.05	-1.12	
Item 19	2	0.43	-0.33	0.00	-2.81	2.81		
Item 20	2	0.37	1.43	0.00	0.25	-0.25		
Item 21	2	0.22	0.17	0.00	0.72	-0.72		
Item 22	2	0.27	0.35	0.00	-4.69	4.69		
Item 23	2	0.16	1.83	0.00	-1.51	1.51		
Item 24	2	0.35	0.19	0.00	-2.38	2.38		
Item 25	2	0.24	1.81	0.00	-2.98	2.98		
Item 26	2	0.29	0.45	0.00	-0.95	0.95		
Item 27	16	0.93	1.24	0.00	1.41	0.56	-0.40	-1.57
Item 28	3	0.99	0.54	0.00	0.82	0.10	-0.92	
Item 29	2	0.46	0.04	0.00	0.56	-0.56		
Item 30	2	0.28	0.32	0.00	-2.60	2.60		
Item 31	2	0.41	0.40	0.00	-1.03	1.03		
Item 32	2	0.29	1.26	0.00	-1.99	1.99		
Item 33	2	0.22	0.13	0.00	-1.82	1.82		
Item 34	2	0.25	0.12	0.00	-2.76	2.76		
Item 35	2	0.29	1.22	0.00	-0.12	0.12		
Item 36	2	0.40	-0.09	0.00	-0.98	0.98		
Item 37	2	0.45	0.77	0.00	0.61	-0.61		
Item 38	2	0.58	0.12	0.00	0.08	-0.08		
Item 39	16	1.10	1.17	0.00	1.53	0.74	-0.65	-1.62
Item 40	3	1.18	0.59	0.00	0.92	0.18	-1.10	
Item 41	2	0.55	0.58	0.00	0.20	-0.20		
Item 42	2	0.29	1.22	0.00	-0.94	0.94		
Item 43	2	0.64	-1.08	0.00	-1.07	1.07		
Item 44	2	0.66	-0.06	0.00	-0.38	0.38		

A7. 2 MATGP IRT Parameters For Unique Items

Item	Max Points	a	b	d1	d2	d3	d4	d5	d6	d7
Item 1	6	0.40	1.63	0.00	-1.74	3.06	-1.30	2.26	-2.67	0.39
Item 2	2	0.16	1.40	0.00	2.60	-2.60				
Item 3	3	0.70	2.19	0.00	0.75	-0.17	-0.57			
Item 4	6	0.42	2.17	0.00	1.08	-0.11	0.28	0.09	-0.59	-0.76
Item 5	2	0.71	-0.39	0.00	0.39	-0.39				
Item 6	1	0.64	1.49							
Item 7	1	0.35	0.87							
Item 8	1	0.94	0.79							
Item 9	2	0.43	0.70	0.00	1.01	-1.01				
Item 10	2	0.48	0.20	0.00	1.35	-1.35				
Item 11	3	0.79	2.13	0.00	-0.22	1.68	-1.46			
Item 12	1	0.97	0.05							
Item 13	2	0.33	0.93	0.00	1.07	-1.07				
Item 14	1	0.66	0.17							
Item 15	1	0.48	-0.17							
Item 16	1	0.49	0.40							
Item 17	2	0.73	-0.19	0.00	1.05	-1.05				
Item 18	1	0.64	1.09							
Item 19	1	0.44	0.91							
Item 20	2	0.60	-0.37	0.00	0.76	-0.76				
Item 21	1	0.36	-1.08							
Item 22	6	0.59	1.71	0.00	-0.74	1.25	-0.28	0.91	-0.98	-0.15
Item 23	6	0.57	2.01	0.00	1.53	1.40	0.79	-0.53	-2.19	-1.00
Item 24	1	0.44	0.65							
Item 25	2	0.35	0.39	0.00	1.34	-1.34				
Item 26	3	0.61	2.12	0.00	-0.17	0.39	-0.23			
Item 27	2	0.36	2.33	0.00	1.27	-1.27				

Item	Max Points	a	b	d1	d2	d3	d4	d5	d6	d7
Item 28	2	0.45	3.74	0.00	-0.75	0.75				
Item 29	1	1.05	-1.50							
Item 30	1	0.90	1.07							
Item 31	4	0.64	2.10	0.00	0.20	0.23	0.36	-0.78		
Item 32	1	0.32	1.96							
Item 33	1	1.12	0.10							
Item 34	3	0.63	2.63	0.00	1.57	-1.51	-0.06			
Item 35	1	0.79	-0.38							
Item 36	1	0.64	-0.60							
Item 37	1	1.05	-0.62							
Item 38	3	1.03	2.50	0.00	0.61	-0.09	-0.52			
Item 39	4	0.60	2.22	0.00	1.03	0.08	-0.86	-0.25		
Item 40	3	0.97	2.23	0.00	0.47	-0.28	-0.18			
Item 41	1	0.72	2.48							
Item 42	1	0.84	0.60							
Item 43	2	0.35	1.71	0.00	1.53	-1.53				
Item 44	1	1.00	1.57							
Item 45	1	0.93	0.57							
Item 46	1	1.02	-0.95							
Item 47	1	0.33	0.68							
Item 48	1	0.99	-0.72							
Item 49	1	0.54	1.34							
Item 50	1	1.34	0.36							
Item 51	2	0.47	1.57	0.00	-0.09	0.09				
Item 52	1	0.45	0.24							
Item 53	1	0.82	0.06							
Item 54	1	0.28	-1.37							

Item	Max Points	a	b	d1	d2	d3	d4	d5	d6	d7
Item 55	1	0.46	-1.29							
Item 56	1	0.31	0.54							