

New York State Testing Program

English Language Arts Grade 4

Technical Report 2003



Developed and published under contract with New York State Department of Education by CTB/McGraw-Hill LLC, a subsidiary of The McGraw-Hill Companies, Inc., 20 Ryan Ranch Road, Monterey, California 93940-5703. Copyright © 2003 by New York State Department of Education. Only State of New York educators and citizens may copy, download, and/or print the document located online at <http://www.emsc.nysed.gov/ciai/testing/pubs.html>. Any other use or reproduction of this document, in whole or in part, requires written permission of New York State Department of Education.

Foreword

The technical information herein is intended for use by those who evaluate tests, interpret scores, or use test results in making educational decisions. It is assumed that the reader has technical knowledge of test construction and measurement procedures, as described in *Standards for educational and psychological testing* (American Educational Research Association, American Psychological Association, & National Council on Measurement in Education, 1999).

Table of Contents

FOREWORD	1
TABLE OF CONTENTS	2
LIST OF TABLES	3
PART 1: TEST DESIGN.....	4
TEST CONFIGURATION	4
STUDENT PARTICIPATION AND TESTING ACCOMMODATIONS	6
<i>Students to be Tested</i>	6
<i>Testing Accommodations</i>	6
ITEM DEVELOPMENT.....	7
ITEM REVIEW PROCESS.....	7
<i>Documenting Content</i>	7
<i>Minimizing Bias</i>	7
<i>Minimizing Speededness</i>	8
TEST CONSTRUCTION AND PRE-EQUATING	8
<i>Calibration Samples</i>	8
<i>Answer Choice Information</i>	8
<i>Item Response Theory Models</i>	8
<i>Equating Method</i>	10
<i>Item Selection Criteria and Process</i>	10
<i>Procedures for Eliminating Bias and Minimizing Differential Item Functioning</i>	12
PART 2: ITEM STATISTICS FOR THE OPERATIONAL DATA.....	13
DATA CLEANING.....	13
ITEM ANALYSIS	14
DIFFERENTIAL ITEM FUNCTIONING ANALYSIS OF OPERATIONAL DATA	15
PART 3: SCORING AND RELIABILITY.....	17
RAW SCORE TO SCALE SCORE CONVERSION	17
RELIABILITY	17
ESTIMATED CONDITIONAL STANDARD ERRORS OF SCALE SCORES	19
LOWEST AND HIGHEST OBTAINABLE SCALE SCORES	19
INTER-RATER AGREEMENT.....	19
EXPECTED SPI SCORES ON THE STANDARDS AT THE DECISION POINTS.....	20
PART 4: DESCRIPTIVE STATISTICS.....	22
SCALE-SCORE FREQUENCY DISTRIBUTIONS FOR THE STATE AND SUBGROUPS.....	22
G4 ELA SCALE SCORE MEANS AND STANDARD DEVIATIONS.....	23
REFERENCES.....	24

List of Tables

Table 1 New York State Learning Standards for English Language Arts.....	4
Table 2 Points per item type for ELA and RD scores.....	5
Table 3 Condition Codes for the ELA CR items.....	5
Table 4 Steps Involved in Data Clean-up for Analysis Preparation.....	13
Table 5 G4 ELA Item Level Statistics.....	14
Table 6 Number of Students in each Gender or Ethnic Group.....	15
Table 7 The Numbers of Items Flagged for DIF in G4 ELA and Reading Summary.....	16
Table 8 Raw Score to Scale Score with SE for G4 ELA 2003.....	18
Table 9 G4 ELA 2003 Inter-Rater Agreement.....	19
Table 10 Percentages of Inter-Rater Score Differences.....	20
Table 11 Reliability Indices of Hand Scoring.....	20
Table 12 G4 ELA 2003 Standard Performance Index Information.....	21
Table 13 G4 ELA 2003 Summary of Scale Score Information.....	22
Table 14 G4 ELA Statewide Scale Score Information.....	23

Part 1: Test Design

The New York State Learning Standards for English Language Arts

The New York State *Learning Standards for English Language Arts* is available from the New York State Education Department web site, at <http://www.emsc.nysed.gov/ciai/ela/pub/elalearn.pdf>. The four learning standards are listed in Table 1 below. The G4 ELA is written to test students in Standards 1, 2, and 3.

Table 1 New York State Learning Standards for English Language Arts

*Standard 1	Students will read, write, listen, and speak for information and understanding.
*Standard 2	Students will read, write, listen, and speak for literary response and expression.
*Standard 3	Students will read, write, listen, and speak for critical analysis and evaluation.
Standard 4	Students will read, write, listen, and speak for social interaction.

Test Configuration

Similar to the 1999 through 2002 forms, the 2003 G4 ELA test has the following configuration. The test is divided into three sessions. There are 28 multiple choice (MC) items worth a total of 28 points; there are seven constructed response (CR) items worth a total of 14 points. The CR items may be short response (SR) or extended response (ER) items. The total number of items on the test is 36, and the maximum raw score total is 42 points.

Session 1

The 28 MC items comprise session 1. Each item in session 1 addresses one of the three tested New York State Learning Standards for English Language Arts.

Session 2

Session 2 contains a listening passage followed by two SR items and one ER item, and an independent writing prompt linked thematically with the listening prompt. The independent writing prompt is scored as an ER item. The three CR items following the listening passage are scored together as a cluster (the listening cluster). These clusters, together with the independent writing item, address Standard 2.

Session 3

Session 3 contains two linked passages (one informational and one literary, or two literary), accompanied by three SR items and one ER item which are scored as a cluster (reading cluster) under Standard 3. Additionally, each of the three SR items in this session is analytically scored to contribute to the *reading* score. (This score is explained in more detail below.)

Writing Mechanics and Reading Scores

As part of the ELA test, the three ER responses across sessions 2 and 3 are scored as a cluster (writing mechanics). Although the writing mechanics is not linked to any of the New York State Learning Standards for English Language Arts (ELA), it contributes to the overall ELA score.

Separate from the ELA test, the 28 MC items and the three analytically-scored SR items constitute the *reading* (RD) test. The RD test is scaled independently from the ELA test and yields RD scores.

Table 2 shows the numbers of score points by the item type or cluster, and the total numbers of items and clusters, for the grade 4 ELA test.

Table 2 Points per item type for ELA and RD scores

Item type or cluster	Grade 4	
	ELA	Reading
Multiple choice (MC)	28 pts	28 pts
Listening cluster	4 pts	
Independent writing item	3 pts	
Writing mechanics cluster	3 pts	
Reading cluster	4 pts	
Three analytic reading items, 2 points each		6 pts
Total points	42 pts	34 pts
Total items or clusters	32 items	31 items

In scaling and scoring, each of the clusters is treated as a constructed response (CR) item. The following condition codes were used in scoring the responses to the CR items:

Table 3 Condition Codes for the ELA CR items

Condition code	Meaning
A	Blank
B	Refusal
C	Insufficient to score
D	Illegible
E	Other language

Student Participation and Testing Accommodations

Students to be Tested

The New York State Testing Program (NYSTP) Grade 4 English Language Arts test must be administered to all public school students in Grade 4 and all ungraded students who are age-equivalent to students in Grade 4. This includes students who have been retained in Grade 4. Nonpublic schools are strongly encouraged to administer the tests. The exceptions noted below apply to students in public and nonpublic schools participating in the NYSTP.

Testing Accommodations

Accommodations were used in the NYSTP operational tests to provide equal access to assessments for students with disabilities. These accommodations are used to increase the validity of test scores by offsetting behavioral constraints due to the disability and retaining the essential features of the assessment. The following represents the policy of the NYSED for the use of testing accommodations

Students with Disabilities

The Committee on Special Education (CSE) must decide for each student on a case-by-case basis (and document on the student's Individualized Education Program) whether the student will participate in the general State assessment, in a locally selected assessment, or in the New York State Alternate Assessment for Students with Severe Disabilities (NYSAA). The criteria that the CSE must use to determine eligibility for a locally selected assessment is available at <http://www.emsc.nysed.gov/deputy/Documents/disabilities-assess.htm>. The criteria to determine eligibility for the NYSAA is available on <http://www.vesid.nysed.gov/specialed/alterassessment/alterassess.htm>.

It is the responsibility of the principal to ensure that testing accommodations specified in the IEP or 504 Plan are provided to students with disabilities as long as they do not alter a construct being measured by the test. Students who have been declassified may continue to be provided testing accommodations if recommended by the local CSE at the time of declassification and in the student's declassification IEP. Testing accommodations that alter the construct being measured are not permitted on elementary- and intermediate-level State assessments. For more information, see <http://web.nysed.gov/vesid/sped/policy/changeaccomm.htm>.

Principals may modify testing procedures for General Education students who incur an injury (for example, a broken arm) or experience the onset of a short- or long-term disability (for example, epilepsy) sustained or diagnosed within 30 days prior to the administration of State tests. In such cases, when sufficient time is not available for the development of an Individualized Education Program (IEP) or a Section 504 Accommodation Plan (504 Plan), principals may authorize certain accommodations that will not significantly change the skills being tested.

Eligibility for such accommodations is based on the principal's professional discretion, but the principal may confer with members of the Committee for Special Education (CSE) or with other school personnel in making such a determination. Pursuant to Section 100.3 of the Regulations of the Commissioner of Education, building principals are responsible for administering State assessments and for maintaining the integrity of test content and programs in accordance with directions and procedures established by the Commissioner of Education.

Limited English Proficient (LEP) Students

The No Child Left Behind (NCLB) Act requires that the English proficiency of all Limited English Proficient (LEP) students (as defined in Part 154 of the Regulations of the Commissioner of Education) be tested annually. New York State has introduced a new assessment of the English language proficiency of students for whom English is a second language. Effective Spring 2003, all LEP students, regardless of grade, must take the New York State English as a Second Language Achievement Test (NYSESLAT). LEP students must take this assessment even if they take the Grade 4 English Language Arts Test.

Additional information concerning the inclusion of LEP students in State examinations in English Language Arts and Mathematics will be provided on the Department's website <http://www.emsc.nysed.gov/osa>.

Other Considerations

When determining who will participate in the NYSTP and who will participate in the Alternate Assessment, school administrators must consider those students who attend programs operated by the Board of Cooperative Educational Services (BOCES), or who are in approved private school placements, as well as in any other programs located outside the school district. Students who are absent during the testing administrations should be tested during the designated makeup period.

Item Development

A staff of professional item writers researched, collected, and wrote the field test material. All assessment materials were carefully reviewed for content and editorial accuracy. Artists and designers worked with the writers during development to ensure graphic and textual consistency. With assistance from the New York State Department of Education, all test items were developed to align with the content and measure the Learning Standards for English Language Arts. Standards Performance Index (SPI) scores are assigned to students on each of these reporting categories.

Item Review Process

Documenting Content

An integral part of the development process was documentation of content using New York State's Learning Standards. All items used on the New York State tests are reviewed for content by both CTB Development staff and by New York State Department of Education staff and New York State teachers. This procedure ensures that items would be sound in content and format, and targeted appropriately to the courses in which the associated concepts are typically taught.

Minimizing Bias

The developers of the NYSTP tests gave careful attention to questions of ethnic, racial, gender, regional, and age bias. All materials were written and reviewed to conform to the company's editorial policies and guidelines for equitable assessment, as well as NYSED's guidelines for item development.

In addition, educators and other stakeholders from different parts of the state reviewed the items from their perspective as members of various ethnic groups. They identified assessment materials that might reflect possible bias in language, subject matter, or representation of people. Their comments and suggestions were considered carefully during the revision and selection of items for the calibration tests. All materials were written to SED specifications and carefully checked by groups of trained New York community participants.

Minimizing Speededness

Test developers also considered speededness in the development of the NYSTP tests. CTB believes that achievement tests should not be speeded; little or no useful instructional information can be obtained from the fact that a student did not finish a test, while a great deal can be learned from student responses to questions. For that reason, sufficient administration time limits were set for the NYSTP tests.

The Research Department at CTB routinely conducts additional speededness analyses based on actual test data. Table 5 shows the omit rates for items on the G4 ELA test. These results provide little evidence of speededness on these tests.

Test Construction and Pre-equating

Calibration Samples

Three field test forms for the NYSTP tests were administered to students in public and private schools across the State in 2002. Efforts were made to ensure that the sample of students was representative of the state tested population. The 2002 field test items were calibrated and equated to the existing scale. Thus, parameters for these items were already on the appropriate New York State scale (one each for grade 4 ELA, grade 4 Mathematics, grade 8 ELA, and grade 8 Mathematics).

Since these items are calibrated and on a common scale, the pool of available grade 4 English Language Arts items can be used to construct a test form and to produce a raw-score-to-scale-score table for that form. The 2003 operational NYSTP tests were constructed using the items in the pool. What follows is an overview of the analysis of field test data which results in the calibration of items.

Answer Choice Information

Statistical information about student performance is produced for each multiple choice item. Specifically, three statistics are examined for each item: (1) the proportion of students choosing each answer, (2) the point-biserial correlation between the answer choice and the number-correct score on the rest of the test, and (3) omit rates. For each constructed response item, the proportion of students at each score level, omit rates, and p-values (mean item score divided by the total number of points possible) are examined.

Item Response Theory Models

Although useful, the differences in proportion of points received (p-values) limit the degree to which one can compare important characteristics of the test items. Item-response theory (IRT) allows one to make better comparisons among items, even those from different test forms, by using a common scale for all items (i.e., as if there were a hypothetical test that contained items from all forms). The three-parameter logistic (3PL) model (Lord & Novick, 1968; Lord, 1980) was used to analyze item responses on the

multiple choice items. For analysis of the constructed response items, the two-parameter partial credit model (2PPC) (Muraki, 1992; Yen, 1993) was used.

Item response theory is a statistical procedure that takes into account the fact that not all test items are alike and that all items do not provide the same amount of information in determining how much a student knows or can do. Computer programs that implement IRT models use actual students' data to estimate the characteristics of the items on a test -- called "parameters." The parameter estimation process is called "item calibration."

IRT models typically vary according to the number of parameters estimated. For the New York State tests, three parameters are estimated: The discrimination parameter, the difficulty parameter(s), and, for multiple choice items, the guessing parameter. The discrimination parameter is an index of how well an item differentiates between high-performing and low-performing students. An item that low-performing students cannot answer correctly, but high-performing students can, will have a high discrimination value. The difficulty parameter is an index of how easy or difficult an item is. The higher the difficulty parameter, the harder the item. The guessing parameter is the probability that a student with very low ability will answer the item correctly.

The scale score (SS) is the basic score for the New York State tests. It is used to derive other scores that describe test performance, such as the four performance levels and the standard-based performance index scores (SPIs). Scale scores can be obtained by one of two scoring methods: IRT item-pattern scoring, or number-correct scoring. Starting in 2002, scores on the New York State tests are determined using number-correct scoring.

Because the characteristics of MC and CR items are different, two IRT models were used in item calibration. The three-parameter logistic (3PL) model (Lord & Novick, 1968; Lord, 1980) was used in the analysis of MC items. In this model, the probability that a student with ability θ responds correctly to item i is

$$P_i(\theta) = c_i + \frac{1 - c_i}{1 + \exp[-1.7a_i(\theta - b_i)]},$$

where a_i is the item discrimination, b_i is the item difficulty, and c_i is the probability of a correct response by a very low-scoring student.

For analysis of the constructed response items, the two-parameter partial credit (2PPC) model (Muraki, 1992; Yen, 1993) was used. The 2PPC model is a special case of Bock's (1972) nominal model. Bock's model states that the probability of an examinee with ability θ having a score $(k - 1)$ at the k -th level of the j -th item is

$$P_{jk}(\theta) = P(X_j = k - 1 | \theta) = \frac{\exp Z_{jk}}{\sum_{i=1}^{m_j} \exp Z_{ji}}, \quad k = 1 \dots m_j,$$

where

$$Z_{jk} = A_{jk}\theta + C_{jk}.$$

The m_j denotes the number of score levels for the j -th item, and typically the highest score level is assigned $(m_j - 1)$ score points. For the special case of the 2PPC model used here, the following constraints were used:

$$A_{jk} = \alpha_j (k - 1),$$

and

$$C_{jk} = -\sum_{i=0}^{k-1} \gamma_{ji}, \text{ where } \gamma_{j0} = 0,$$

where α_j and γ_{ji} are the free parameters to be estimated from the data. Each item has $(m_j - 1)$ independent γ_{ji} parameters and one α_j parameter; a total of m_j parameters are estimated for each item.

The IRT model parameters were estimated using CTB's PARDUX software (Burket, 1991). PARDUX estimates parameters simultaneously for MC and CR items using marginal maximum likelihood procedures implemented via the EM (expectation-maximization) algorithm (Bock & Aitkin, 1981; Thissen, 1982).

Simulation studies have compared PARDUX with MULTILOG (Thissen, 1991), PARSCALE (Muraki & Bock, 1991), and BIGSTEPS (Wright & Linacre, 1992). PARSCALE, MULTILOG, and BIGSTEPS are among the most widely known and used IRT programs. PARDUX was found to perform at least as well as these other programs (Fitzpatrick, 1990; Fitzpatrick, 1994; Fitzpatrick and Julian, 1996).

Equating Method

After the item calibration, all of the Grade 4 English Language Arts field test items were placed on the NYS G4 ELA scale using the 2002 operational items as anchors. This was possible because the operational items were taken by the same students who took the field test items within the same testing window. The equating was performed using the test characteristic curve method (Stocking & Lord, 1983) implemented by PARDUX. In previous years, operational data were used to re-calibrate items and re-equate them. NYSED, however, made a decision in 2002 to use the pre-equating model, which is similar to what is done for the New York State Regents program. This allows the production of scoring tables (see Part 3) ahead of the operational administration, once the operational form is selected.

Item Selection Criteria and Process

Item selection for the NYSTP tests was based on the classical and IRT statistics of the test items. Selection was conducted by content experts from CTB and NYSED and reviewed by psychometricians at CTB. Final approval of the items selected was given by NYSED. Two criteria governed the item selection process. The first of these was to meet the content specifications specified by the New York State Department of Education. Within the limits set by these requirements, developers selected from the pool of field test items those with the best psychometric characteristics. Developers selected items that minimized measurement error throughout the range of expected achievement as indicated by the reciprocal of the square root of the IRT information function (Lord, 1980, p. 71). Developers aimed to create forms with the content and psychometric properties of previous operational forms.

Item selection for the calibration tests was facilitated using the Windows version of the program ITEMSYS (Burket, 1988). ITEMSYS creates an interactive connection between the developer selecting the test items and the item database. This program monitors the impact of each decision made during the item selection process and offers a variety of options for grouping, classifying, sorting, and ranking items to highlight key information as it is needed (see Green, Yen, & Burket, 1989).

ITEMSYS has three parts. The first part selects a working item pool of manageable size from the larger tryout pool. The second part of the program uses this selected item pool to perform the final test selection. In the third part of the program a table shows both expected number correct and the standard error of ability estimate (a function of scale score), as well as statistical and graphic summaries on bias, fit, and the standard error of the final test. Any fault in the final selection becomes immediately apparent as the final statistics are generated. Examples of possible faults that may occur are when the test is too easy or difficult, contains items demonstrating differential item functioning or DIF (see below), does not meet the requirements to match a parallel form, or does not adequately measure part of the range of performance. A developer detecting any such problems can then return to the second stage of the program and revise the selection. The flexibility and utility of the program encourages multiple attempts at fine-tuning the item selection.

Procedures for Eliminating Bias and Minimizing Differential Item Functioning

As part of the testing, the students reported their gender and ethnic background information. Using this self-reported information, statistical differential item functioning (DIF) analyses were conducted for gender groups and such ethnic groups as African-American, Hispanic-American and Asian-American in the sample.

Three procedures were used to eliminate bias and minimize differential item functioning (DIF) in the New York State Tests.

The first was based on the premise that careful editorial attention to validity is an essential step in keeping bias to a minimum. Bias can occur only if the test is differentially valid for a given group of test takers. If the test entails irrelevant skills or knowledge (however common), the possibility of DIF is increased. Thus, preserving content validity is essential.

The second step was to follow the item writing guidelines established by NYSED. Developers reviewed NYSTP materials with these guidelines in mind. These internal editorial reviews were done by at least four separate people: the content editor, who directly supervises the item writers; the project director; a style editor; and a proofreader. The final test built from the tryout materials was reviewed by at least these same people.

In the third procedure, New York State educational community professionals who represent various ethnic groups reviewed all tryout materials. These professionals were asked to consider and comment on the appropriateness of language, subject matter, and representation of people.

It is believed that these three procedures improved the quality of the New York State Tests and reduced bias. However, current evidence suggests that expertise in this area is no substitute for data; reviewers are often wrong about which items work to the disadvantage of a group, apparently because some of their ideas about how students will react to items may be faulty (Sandoval & Mille, 1979; Jensen, 1980). Thus, an empirical approach is desirable.

Part 2: Item Statistics for the Operational Data

Data Cleaning

Item analyses were conducted once CTB received data that met the following requirements established by NYSED:

- Comprises at least 85% of the estimated number of students in the State,
- Includes New York City and Buffalo,
- Includes at least one of the cities of Rochester, Syracuse, or Yonkers, and
- Includes at least two of the cities of Mount Vernon, Albany, Binghamton, Schenectady, or New Rochelle.

Initially, the state data set contained 247,209 cases. Table 4 below shows the data cleaning steps and the resulting size of the 85% sample used for conducting item analyses.

Table 4 Steps Involved in Data Clean-up for Analysis Preparation

Steps Taken	# Cases Deleted	Ending N
Original Data		247,209
Duplicate Records	74	247,135
Grade Not Equal to 4	0	247,135
LEP3 Data	886	246,249
Non-LEP3 Data	0	246,249
Non-LEP3 Data after Exclusion Rules	8,318	237,931

Students whose LEP status = 3 are not required to take the test.

As Table 4 shows, the following records were eliminated, in the order listed:

- Duplicated records,
- Students whose limited English proficient (LEP) status was "3," indicating that they scored below the thirtieth percentile on a norm-referenced English reading test or the publisher's recommended score on an approved measure of English as a second language in reading,
- All students whose LEP status is not 3, and
- Students who did not have a valid attempt in each of three sections as determined by the application of CTB's Invalidation / Omission / Suppression rules (approved by NYSED).

Item Analysis

Table 5 presents the results of item analyses conducted using the scaling sample for the G4 ELA test. The labels for the variables denote the following:

ITEM Item number.

OMIT Proportion of students who had blank response or double marks on MC items, or condition codes on the CR items.

PCTSEL* For MC items, this is the percentage of students who chose the first through the fourth answer option. For CR items, it is the percentage of students who received a score of 0 through the maximum number of points possible. Asterisked numbers indicate values for the correct response option.

PTBIS* Point-biserial correlations for each response option. Asterisked numbers indicate values for the correct response option.

P_VAL Item difficulty after omitted responses are converted to 0s (wrong). For MC items, p-value is the proportion of students responding correctly. For a CR item, p-value is the mean raw score divided by the maximum number of score points for an item.

Table 5 G4 ELA Item Level Statistics

Raw Score Data		Test Administration Data						Reliability Feldt-Raju				P-Value Mean
Mean	SD	Number of Items			Number of Students			ELA	RD			
29.12	7.64	35			237,931			0.904	0.884			0.71
Item	Omit	Pctsel0	Pctsel1	Pctsel2	Pctsel3	Pctsel4	p_bis1	p_bis2	p_bis3	p_bis4	Key	p-value+
1	0.0003		14.46%	*78.75%	4.95%	1.82%	-0.496	*0.472	-0.105	-0.246	2	0.7875
2	0.0004		1.27%	32.48%	1.29%	*64.92%	-0.189	-0.309	-0.152	*0.265	4	0.6492
3	0.0005		1.31%	*87.53%	3.86%	7.24%	-0.202	*0.329	-0.264	-0.236	2	0.8753
4	0.0004		*76.00%	6.99%	13.33%	3.63%	*0.412	-0.271	-0.282	-0.298	1	0.76
5	0.0007		3.66%	5.82%	*83.67%	6.78%	-0.331	-0.334	*0.464	-0.247	3	0.8367
6	0.0009		13.25%	*71.49%	5.76%	9.41%	-0.495	*0.542	-0.205	-0.257	2	0.7149
7	0.0009		2.78%	3.53%	*89.90%	3.70%	-0.252	-0.262	*0.435	-0.324	3	0.899
8	0.0014		13.22%	9.89%	10.38%	*66.37%	-0.242	-0.282	-0.291	*0.416	4	0.6637
9	0.0012		*79.77%	5.78%	4.81%	9.52%	*0.535	-0.351	-0.330	-0.332	1	0.7977
10	0.0011		2.76%	21.43%	*64.59%	11.12%	-0.234	-0.390	*0.524	-0.336	3	0.6459
11	0.0011		5.74%	*81.56%	7.51%	5.09%	-0.302	*0.447	-0.243	-0.335	2	0.8156
12	0.0013		8.07%	*78.15%	4.10%	9.55%	-0.281	*0.476	-0.289	-0.345	2	0.7815
13	0.0013		11.21%	*70.90%	7.87%	9.89%	-0.351	*0.473	-0.307	-0.237	2	0.709

Table 5 continues

Table 5 G4 ELA Item Level Statistics (continued)

Item	Omit	Pctsel0	Pctsel1	Pctsel2	Pctsel3	Pctsel4	p_bis1	p_bis2	p_bis3	p_bis4	Key	p-value+
14	0.0012		*87.28%	3.42%	4.11%	5.07%	*0.401	-0.227	-0.272	-0.287	1	0.8728
15	0.0012		4.74%	6.90%	*79.73%	8.51%	-0.271	-0.343	*0.499	-0.327	3	0.7973
16	0.0014		5.75%	8.08%	*84.07%	1.96%	-0.301	-0.340	*0.474	-0.278	3	0.8407
17	0.0021		6.67%	5.93%	*82.86%	4.33%	-0.311	-0.313	*0.431	-0.200	3	0.8286
18	0.0027		1.92%	2.99%	10.42%	*84.40%	-0.221	-0.251	-0.286	*0.368	4	0.844
19	0.0027		6.48%	*77.23%	3.30%	12.73%	-0.234	*0.384	-0.217	-0.311	2	0.7723
20	0.0031		4.30%	*74.60%	3.84%	16.95%	-0.267	*0.330	-0.232	-0.236	2	0.746
21	0.0032		*68.62%	12.39%	3.41%	15.26%	*0.218	-0.155	-0.234	-0.176	1	0.6862
22	0.0059		*71.72%	12.04%	5.28%	10.37%	*0.305	-0.240	-0.222	-0.183	1	0.7172
23	0.0074		*56.25%	11.03%	16.05%	15.94%	*0.359	-0.228	-0.308	-0.155	1	0.5625
24	0.0073		16.27%	8.88%	*68.80%	5.32%	-0.241	-0.243	*0.413	-0.351	3	0.688
25	0.0086		10.39%	*72.39%	5.24%	11.12%	-0.252	*0.451	-0.274	-0.325	2	0.7239
26	0.0105		*48.52%	16.62%	8.37%	25.45%	*0.357	-0.360	-0.299	-0.063	1	0.4852
27	0.0111		8.65%	*70.34%	12.73%	7.18%	-0.323	*0.498	-0.313	-0.275	2	0.7034
28	0.0029	0.93%	12.93%	35.76%	36.57%	13.51%					CR	0.6205
29	0.0065	3.59%	23.19%	43.74%	28.84%						CR	0.6573
30	0.0030	1.40%	21.71%	49.11%	27.48%						CR	0.6746
31	0.0042	1.46%	14.71%	35.74%	33.46%	14.22%					CR	0.6086

+ average score divided by maximum score

Differential Item Functioning Analysis of Operational Data

To assess DIF for the New York State tests, students were identified as African-American, White, Hispanic, or Asian-American. For grade 4, students bubble in this information. These ethnic groups were chosen for DIF analyses because these populations are the largest in the State. Gender analyses were also conducted.

Developers strive to produce tests that minimize DIF. The DIF results reported here are those obtained when scoring students on the operational test using the pre-equated field test parameters. Thus, they may differ from DIF results obtained at the time of the field test administration.

Using demographic information, statistical DIF analyses were conducted for various ethnic groups and for males and females. A random sample was drawn from the final state GRT. Next, the sample was augmented by randomly selecting additional cases for any group of students whose count in the sample was less than 500 in an attempt to enhance the reliability of the DIF analyses. The numbers of cases for the groups are reported in Table 6 below.

Table 6 Number of Students in each Gender or Ethnic Group

Test	Female	Male	African-American	Asian-American	Hispanic-American
Grade 4 ELA / RD	116563	121301	48177	13283	41756

The standardized mean difference (SMD) statistic (Zwick, Donoghue, & Grima, 1993) was used to examine DIF on the operational data. The SMD statistics can provide DIF information for both multiple choice and constructed response items. The SMD takes into account the natural ordering of the response levels of the items and has the desirable property of being based on those ability levels where members of the focal group are present. The standardized mean difference output results in a single statistic for each item.

$$SMD = \sum p_{Fk} m_{Fk} - \sum p_{Rk} m_{Rk},$$

where p_{Fk} is the proportion of focal group members who are at the kth level of the matching variable,

m_{Fk} is the mean item score for the focal group at the kth level, and

m_{Rk} is the analogous value for the reference group.

The matching variable is raw score and the kth level refers to the each successive raw score point.

A moderate amount of practically significant DIF, for or against the focal group, is represented by an SMD with an absolute value between .10 and .19, inclusive. A large amount of practically significant DIF is represented by an SMD with an absolute value of .20 or greater. SMD DIF results using operational data for G4 English Language Arts are summarized below.

Table 7 The Numbers of Items Flagged for DIF in G4 ELA and Reading Summary

Focal Group	Direction of DIF	G4 ELA	G4 RD
Female	In favor of	2 ¹	0
	Against	0	0
African-American	In favor of	0	0
	Against	0	0
Asian-American	In favor of	0	0
	Against	1 ²	1 ³
Hispanic-American	In favor of	0	0
	Against	0	0

¹ Items #31, #32 (D = .11, and .11)

² Item #2 (D = -.12)

³ Item #2 (D = -.12)

Part 3: Scoring and Reliability

Raw Score to Scale Score Conversion

To facilitate ease of interpretation and implementation, number-correct scoring was used on the New York State Tests in 2003. In number-correct scoring, a student's scale score is derived directly from his or her raw, or number-correct, score. The relationship between raw scores and their corresponding scale scores is expressed in a raw-score-to-scale-score (RS-SS) table.

In IRT, all the item characteristic curves for the items on a test can be added together to yield a function - the test characteristic curve (TCC) - that shows the expected raw score for each given scale score. By inverting the TCC, an expected scale score can be computed for each raw score. This new function - the inverse of the TCC - can be summarized in an RS-SS table. An advantage of RS-SS tables is that they make scoring relatively straightforward: With number-correct scoring, it is sufficient to know how many raw score points a student obtained on the test to determine a student's scale score. The RS-SS conversion tables for both content areas appear in Table 8.

Reliability

The reliability of measurement refers to the reproducibility or consistency of an individual's tests scores. The two most frequently reported indices of reliability are the standard error of measurement and the reliability coefficient.

The standard error of measurement is a measure of the extent to which an individual's scores vary over numerous parallel tests. We computed a *conditional* error – the standard error (SE) for each scale score for G4 ELA and these are reported below in Table 8. See also the section on estimated conditional standard errors of scale scores, below.

The reliability coefficient is the correlation coefficient between scores on parallel tests and is an index of how well scores on one parallel test predict scores from another parallel test. Among several ways to estimate the reliability of a test, Cronbach's alpha (Cronbach, Schönemann, & McKie, 1965) probably is the most frequently used. Cronbach's alpha is a measure of internal consistency (i.e., how homogeneous test items are) that is appropriate for a test containing only MC items. Since the G4 ELA test contains MC and CR items, Cronbach's alpha would underestimate reliability because of the effect of variance attributable to item types. A more appropriate index of internal consistency, the Feldt-Raju index, was used to estimate the reliability of the G4 ELA test. It was 0.904, a value comparable to that for 2002.

Table 8 Raw Score to Scale Score with SE for G4 ELA 2003

No. Correct (RS)	ELA		Reading	
	Scale Score	SE	Scale Score	SE
0	455	126	455	133
1	455	126	455	133
2	455	126	455	133
3	455	126	455	133
4	455	126	455	133
5	505	76	492	97
6	538	43	540	48
7	554	28	557	31
8	564	22	568	23
9	572	18	577	19
10	579	16	583	16
11	585	14	589	15
12	590	13	594	13
13	595	12	599	12
14	599	11	603	12
15	603	11	607	11
16	607	10	611	11
17	610	10	615	10
18	614	10	619	10
19	617	9	622	10
20	620	9	626	10
21	624	9	629	10
22	627	9	633	10
23	630	9	637	10
24	633	9	641	10
25	637	9	645	10
26	640	9	649	10
27	644	9	654	11
28	647	9	660	12
29	651	10	667	14
30	655	10	675	16
31	660	11	687	20
32	665	11	709	30
33	671	12	800	121
34	677	13		
35	684	14		
36	693	15		
37	702	16		
38	713	18		
39	727	21		
40	748	28		
41	800	70		

Estimated Conditional Standard Errors of Scale Scores

Each student's scale score is based on a sample of student performance at a given time and inherently contains some measurement error. The classical SEM presumes the amount of measurement error is constant throughout the range of student ability. However, this is not realistic. Measurement error is less, and reliability greater, where more items exist and items are more informative. Item response theory lends itself to the calculation of a standard error for each scale score.

Table 8 lists standard errors for selected scale scores. These standard errors are "constrained" so that the upper and lower limits of one standard error band around a scale score are below the upper and lower limits of the band for the next higher scale score. Typically, only standard errors on extreme ends are constrained. Because more items exist in the middle range of scale scores, the standard error is typically the smallest in the middle. A SS plus and minus one SE constitutes a 68% confidence interval. For example, for a student whose grade 4 ELA SS is 630, we are 68% confident that his or her true score lies within the range 630 plus or minus 9, that is, between 621 and 639.

Lowest and Highest Obtainable Scale Scores

A maximum likelihood procedure cannot produce scale score estimates for students with zero or perfect scores. Scale score estimates below the level expected by guessing are unreliable and subsequently not reported. Also, while maximum likelihood estimates may be available for students with extreme scores other than a perfect score, occasionally these estimates have standard errors that are very large, and differences between these extreme values have little meaning. Therefore, scores are established for these students based on a rational but necessarily non-maximum likelihood procedure. These values are called the lowest obtainable scale score (LOSS) and the highest obtainable scale score (HOSS). The same LOSS and HOSS values are used for either number-correct or item-pattern scoring. For the New York State G4 ELA test, LOSS and HOSS values were set at 455 and 800.

Inter-Rater Agreement

In order to monitor the reliability of scoring among the teachers who scored the student responses, approximately 10% of the student papers were submitted to a second group of raters provided by Measurement Incorporated. Note that the teachers were trained by Measurement Incorporated. The results of the inter-rater agreement analyses for public schools and outside of New York City are provided in Tables 9-11. Additional results for public schools in New York City and non-public school will be reported as they become available.

Table 9 G4 ELA 2003 Inter-Rater Agreement

Inter-Rater Agreement (Read 1 : Non-NYC public school teachers; Read 2 : MI readers)								
CR Item	Score Points	Agreement (%)			RS Mean		RS SD	
		Exact	Approx.	TOTAL	Read 1	Read 2	Read 1	Read 2
Listening	4	48.4	45.6	94.0	2.6	2.2	0.88	0.82
Independent Writing	3	57.2	40.6	97.8	2.0	1.8	0.81	0.75
Writing Mechanics	3	57.8	40.5	98.3	2.1	1.9	0.74	0.70
Reading	4	49.4	44.6	94.0	2.5	2.2	0.95	0.88

Approximate agreement (%) is the percent of pairs of reads that differ by one score point.
Total agreement (%) is the sum of exact and approximate agreement percents

Table 10 Percentages of Inter-Rater Score Differences

Reader 1 (Non-NYC public school teachers) minus Reader 2 (MI readers)								
CR Item	-3	-2	-1	0	1	2	3	4
Listening		0.54	10.61	48.40	34.97	5.27	0.20	
Ind Writing	0.02	0.85	13.54	57.21	27.04	1.27	0.07	
Writ Mech		0.27	11.06	57.81	29.41	1.45		
Reading	0.02	0.87	13.05	49.41	31.51	4.96	0.13	0.04

Table 11 Reliability Indices of Hand Scoring

CR Item	Intra-Class Correlation ¹	Weighted Kappa ²
Listening	0.77	0.55
Ind Writing	0.80	0.59
Writ Mech	0.78	0.56
Reading	0.80	0.60

1 Agresti, A. (1990). Categorical data analysis (pp.366-367). New York: Wiley. Intra-class correlation is the percent of overall score variance accounted for by the variance of mean response scores.

2 Weighted kappa is a measure of association in contingency tables, and is 1 when agreement is perfect and 0 when agreement is what would be expected by chance.

Expected SPI Scores on the Standards at the Decision Points

The current New York State Grades 4 and 8 Score Reports for students report a Standard Performance Index (SPI) score for each of the standards or Key Ideas. The SPI for a student – for a given Key Idea – is an estimate of the percent of maximum raw score that the student would get if he or she took a large sample of items in that Key Idea. The SPI is a diagnostic tool in the sense that it provides a profile of the student's relative strengths and weaknesses in terms of the content standards. However, just because a student has a high SPI on one key idea and a low SPI on another key idea does not necessary mean that she or he is strong on the former standard and weak on the latter. This can occur if items measuring one key idea tend to be easy, while items measuring another key idea tend to be hard.

What teachers and students seem to need in order to better understand the SPIs are the SPIs expected of students who are just at each of the New York State decision points. These expected SPIs at the decision points can be used as "reference points" against which each student's SPIs are compared. For example if a student's SPI on Standard 3 is 60 and the expected SPI for the Level 3 Student is 57, the student's 60, although seemingly low compared with the perfect 100, is still higher than what is expected for the Level 3 Student on the Standard. Expected SPIs for the 2003 Grade 4 English Language Arts exam are listed in Table 12.

Table 12 G4 ELA 2003 Standard Performance Index Information

Standard	Expected Percent of the Max. Raw Score at each of the Cut Points.				
	# Items	Max Pts.	Level 2	Level 3	Level 4
			At SS=603	At SS=645	At SS=692
1	14	14	39	76	96
2	11	16	40	66	85
3	5	8	27	57	82

Part 4: Descriptive Statistics

Scale-Score Frequency Distributions for the State and Subgroups

Table 13 summarizes the scale-score frequency distributions for the state and the following groups of students:

- public schools,
- non-public schools,
- two groups of limited-English-proficient (LEP) students,
- non-disabled students, and
- students with disabilities.

The public vs. non-public distinction was identified by the 9th character of the BEDs LEA code for each school. The non-disabled vs. disabled distinction was identified in the final state dataset. Additionally, two groups of LEP students are defined as those who have either "2" or "3" in the appropriate column of the final state dataset. The "LEP2" group is identified as having limited English proficiency and scored at or above either the 30th percentile on a norm-referenced English reading test or the publisher's recommended score on an approved measure of English as a Second Language (ESL) in reading. Similarly, the "LEP3" group is identified as having limited English proficiency and scored below either the 30th percentile on a norm-referenced English reading test or the publisher's recommended score on an approved measure of English as a Second Language (ESL) in reading.

A summary table of the scale score frequency distributions containing the SSs at the 10th, 25th, 50th, 75th, and 90th percentiles is provided below. No interpolation was employed in computing the percentiles. As an example, in the row of Statewide Inclusive at the 25th percentile the number 637 represents the highest scale score achieved by the lowest 25 percent of the population.

Table 13 G4 ELA 2003 Summary of Scale Score Information

Sub Groups - Percentages	10th	25th	50th	75th	90th
Statewide Inclusive	614	637	660	684	713
LEP = 2	607	630	651	677	702
LEP = 3	455	455	617	647	677
Public	610	633	660	684	713
Non-Public	620	640	660	684	713
Disabled	579	599	620	644	665
Visually Impaired	579	603	633	665	693
Non-Disabled	620	640	665	684	713

G4 ELA Scale Score Means and Standard Deviations

The total number of students and the percent of students in each performance level in the statewide final general research file are shown in the table below. Statistics for the three previous years are also included.

Table 14 G4 ELA Statewide Scale Score Information

Year	Population Sub Grouping	Number of Students (N)	PCT in PL1	PCT in PL2	PCT in PL3	PCT in PL4
2003	All Students	238545	5.63	29.58	42.87	21.92
2002	All Students	242172	7.72	30.00	41.57	20.70
2001	All Students	245019	9.94	29.43	43.65	16.98
2000	All Students	247301	9.37	31.40	43.63	15.60
1999	All Students	240539	10.84	40.50	43.65	5.01

References

- Agresti, A. (1990) *Categorical Data Analysis*. Wiley, New York, NY pp.366-367.
- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (1999). *Standards for educational and psychological testing*. Washington, DC: American Psychological Association, Inc.
- Bock, R. D. (1972). Estimating item parameters and latent ability when responses are scored in two or more nominal categories. *Psychometrika*, 37, 29-51.
- Bock, R. D., & Aitkin, M. (1981). Marginal maximum likelihood estimation of item parameters: An application of an EM algorithm. *Psychometrika*, 46, 443-459.
- Burket, G. R. (1988). *ITEMSYS* [Computer program]. Unpublished.
- Burket, G. R. (1991). *PARDUX* [Computer program]. Unpublished.
- Cronbach, L. J., Schönemann, P., & McKie, D. (1965). Alpha coefficients for stratified-parallel tests. *Educational and Psychological Measurement*, 25, 291-312.
- Fitzpatrick, A. R. (1990) *Status Report on the results of Preliminary Analysis of Dichotomous and Multi-Level Items Using the PARMATE Program*. Unpublished manuscript
- Fitzpatrick, A. R. (1994) *Two studies comparing parameter estimates produced by PARDUX and BIGSTEPS*. Unpublished manuscript.
- Fitzpatrick, A. R., & Julian, M. W. (1996) *Two studies comparing the parameter estimates produced by PARDUX and PARSCALE*. Monterey, CA:CTB/McGraw-Hill.
- Green, D. R., Yen, W. M., & Burket, G. R. (1989). Experiences in the application of item response theory in test construction. *Applied Measurement in Education*, 2, 297-312.
- Jensen, A. R. (1980). *Bias in mental testing*. New York: Free Press.
- Lord, F. M. (1980). *Applications of item response theory to practical testing problems*. Hillsdale, NJ: Lawrence Erlbaum.
- Lord, F. M., & Novick, M. R. (1968). *Statistical theories of mental test scores*. Menlo Park, CA: Addison-Wesley.
- Muraki, E. (1992). A generalized partial credit model: Application of an EM algorithm. *Applied Psychological Measurement*, 16, 159-176.
- Muraki, E., & Bock, R. D. (1991). *PARSCALE: Parameter Scaling of Rating Data* [Computer program]. Chicago, IL: Scientific Software, Inc.

New York State Department of Education. (1995). *Test Access and Modification for Individuals with Disabilities*. Available at <ftp://unix2.nysed.gov/pub/education.dept.pubs/vesid/oses/test.access.mod/testacce.txt>.

New York State Department of Education. (2003). *Learning Standards for English Language Arts*. Available at <http://www.emsc.nysed.gov/ciai/ela/pub/elalearn.pdf>.

Sandoval, J. H., & Mille, M. P. (1979, August). *Accuracy of judgments of WISC-R item difficulty for minority groups*. Paper presented at the annual meeting of the American Psychological Association, New York.

Stocking, M. L., & Lord, F. M. (1983). Developing a common metric in item response theory. *Applied Psychological Measurement*, 7, 201-210.

Thissen, D. (1982). Marginal maximum likelihood estimation for the one-parameter logistic model. *Psychometrika*, 47, 175-186.

Thissen, D. (1991). *MULTILOG* [Computer program]. Chicago, IL: Scientific Software, Inc.

Wright, B. D., & Linacre, J. M. (1992). *BIGSTEPS Rasch Analysis* [Computer program]. Chicago, IL: MESA Press.

Yen, W. M. (1993). Scaling performance assessments: Strategies for managing local item dependence. *Journal of Educational Measurement*, 30, 187-213.

Zwick, R., Donoghue, J. R., & Grima, A. (1993). Assessment of differential item functioning for performance tasks. *Journal of Educational Measurement*, 36, 233-25.