

3. OVERVIEW OF STATISTICAL SUMMARIES

This section provides general information about statistical and psychometric summaries used for the 2008 MSA-Math program. Actual statistical results described in this section appear in section 4 and appendices.

3.1 Classical Descriptive Statistics

Table 4.1 contains the classical descriptive statistics of each form for each grade and includes:

- Form number
- Number of items
- Numbers of students (These numbers were based on a whole population.)
- Means and standard deviations of raw scores
- Stratified Cronbach Alpha
- Standard error of measurement (SEM)

Stratified Cronbach Alpha

The 2008 MSA-Math tests included *SR*, *SPR*, *BCR*, and *ECR* items. Consequently, it was necessary to use an adequate reliability coefficient that addressed the important factor, different item type. The following formula depicts the reliability coefficient, *Stratified Cronbach Alpha*:

$$\text{Stratified } a = 1 - \frac{((\sigma_{SR}^2(1 - \rho_{SR}) + (\sigma_{CR}^2(1 - \rho_{CR})))}{\sigma_t^2}$$

where

σ_{SR}^2 = variance of score on SR and SPR items

σ_{CR}^2 = variance of score on BCR and ECR items

σ_t^2 = variance of total score

ρ_{SR} = reliability coefficient of score on SR and SPR items, and

ρ_{CR} = reliability coefficient of score on BCR and ECR items.

Standard Error of Measurement (Based on Classical Test Theory)

The *standard error of measurement (SEM)* is commonly used in interpreting and reporting individual test scores and score differences on tests (Harvill, 1991).

Classical test theory is based on the following assumptions (Andrich & Luo, 2004):

- Each person v has a true score on the construct, usually denoted by the variable T_v
- The best overall indicator of the person's true score is the sum of the scores on the items and is usually denoted by the variable X_v
- This observed score will have an error for each person which is usually denoted by E_v
- These errors are not correlated with the true score
- Across a population of people, the errors sum to 0 and they are normally distributed.

From these assumptions, the following equations can be derived:

$$X_v = T_v + E_v.$$

Therefore,

$$\sigma_x^2 = \sigma_t^2 + \sigma_e^2$$

where

σ_x^2 = the variance of the observed score in a population of persons,

σ_t^2 = the variance of their true score variance, and

σ_e^2 = the error variance.

The reliability coefficient of the test can be calculated by the following formula:

$$\rho_x = \frac{\sigma_t^2}{\sigma_x^2} = \frac{\sigma_x^2 - \sigma_e^2}{\sigma_x^2}.$$

Thus, the *SEM* is calculated by the following formula:

$$\sigma_e = \sigma_x \sqrt{1 - \rho_x}.$$

For example, consider a student with a score of 90 from a sample of students with a mean score of 60 and variance of 225 on a test with reliability of 0.80. According to the formulas provided above, the obtained score is 90, and its *SEM* is 6.71. Thus, an approximate 68% score band for estimating this student's true score is from 83.29 (90 - 6.71) to 96.71 (90 + 6.71).

Note that this equation is only useful to estimate true score when the test reliability is reasonably high and the obtained score for the examinee is not an extreme deviate from the mean of the appropriate reference group. When we use this equation, consequently, we should be careful with statements so that they do not imply greater precision than is actually involved (Harvill, 1991).

Conditional Standard Error of Ability Estimate (Based on the Rasch Model)

Under the Rasch (i.e., 1-PL IRT) model, the *SE* for each person is as follows (Andrich & Luo, 2004):

$$\sigma_{\hat{\beta}} = \frac{1}{\sqrt{\sum_{i=1}^L p_{vi}(1-p_{vi})}}$$

where

v = subscript for a person,

i = subscript for an item,

L = length of the test,

$\hat{\beta}$ = ability estimate, and

p_{vi} = the probability that a person answers an item correctly and defined as follows:

$$p_{vi} = \frac{e^{\beta_v - \delta_i}}{1 + e^{\beta_v - \delta_i}} \text{ where } \beta_v \text{ is person's ability and } \delta_i \text{ is item's difficulty.}$$

A confidence band can be found for use in interpreting the ability estimate. For example, an approximate 68% confidence interval for $\hat{\beta}$ is given by

$$\hat{\beta} \pm SE$$

3.2 Scale Score Descriptive Statistics

Table 4.2 provides information about scale score descriptive statistics of each form for each grade and includes:

- Form number
- Number of items
- Numbers of students
- Mean and standard deviation of scale scores
- 10% quantile (P10), 25% quantile (Q1), median (P50), 75% quantile (Q3), 90% quantile, and IQR (Interquantile Range= Q3-Q1)
- Conditional *standard errors* (*SE*) for the proficient and advanced cut scores

In addition, Appendix A provides frequency distributions and histograms of the scale scores of the 2008 MSA-Math as well as the 2006 MSA-Math.

3.3 Classical and Rasch (1-Parameter Logistic IRT) Item Parameters

Appendix C provides both classical and Rasch item parameters and includes:

- Item type (*SR*, *SPR*, *BCR*, or *ECR*)
- *P*-value: in order for *p*-values of *BCR* and *ECR* items to be comparable with *p*-values of the *SR* and *SR* items they were calculated as modified proportions of the maximum obtainable domain scores.
- Point-biserial correlation: a Pearson's *r* between the scored item and the total score
- Rasch item difficulty estimate (D_i)
- Conditional standard error of Rasch item difficulty estimate
- Rasch step difficulty estimate (or structure calibration estimate, F_{ij})
- Mean-square infit
- Mean-square outfit

First of all, it should be noted that all the Rasch item and step difficulty parameters were placed on a common scale (i.e., the 2003 scale for grades 3, 5, and 8; the 2004 scale for grades 4, 6, and 7).

Second, the following formula shows how structure measure estimate (D_{ij}) is calculated from both D_i and F_{ij} directly obtained from a run of Winsteps:

$$D_{ij} = D_i + F_{ij},$$

where D_{ij} = structure measure estimate

D_i = item difficulty estimate,

F_{ij} = structure calibration estimate (i.e., step difficulty estimate).

Finally, the following formulas show how conditional standard error (SE) of item difficulty estimate (D_i) and structure measure estimate (F_{ij}) were driven (Wright & Masters, 1982):

$$SE(D_i) = 1 / \sqrt{\sum_{n=1}^N [\sum_k^{m_i} k^2 p_{nik} - (\sum_k^{m_i} k p_{nik})^2]}$$

$$SE(F_{ij}) = 1 / \sqrt{\sum_{n=1}^N (\sum_{k=0}^j p_{nik} - (\sum_{k=j+1}^{m_i} p_{nik})^2)}$$

$$\text{where } P_{nix} = \exp \sum_{j=0}^x (\theta_n - D_{ij}) / \sum_{k=0}^{m_i} \left[\exp \sum_{j=0}^k (\theta_n - D_{ij}) \right]$$

$x = 0, 1, \dots, m_i$, and

$k = 1, 2, \dots, m_i$.

Fit Statistics for the Rasch Model

Fit statistics are used for evaluating the goodness-of-fit of a model to the data. Fit statistics are calculated by comparing the observed and expected trace lines obtained for an item after parameter estimates are obtained using a particular model. *WINSTEPS* provides two kinds of fit statistics called *mean-squares* that show the size of the randomness or amount of distortion of the measurement system.

Outfit mean-squares are influenced by outliers and are usually easy to diagnose and remedy. *Infit* mean-squares, on the other hand, are influenced by response patterns and are harder to diagnose and remedy. Table 3.1 provides a guideline for evaluating mean-square fit statistics (Linacre & Wright, 2000).

In general, mean-squares near 1.0 indicate little distortion of the measurement system, while values less than 1.0 indicate observations are too predictable (redundancy, model overfit). Values greater than 1.0 indicate unpredictability (unmodeled noise, model underfit).

Table 3.1 Criteria to Evaluate Mean-Square Fit Statistics

Mean-Square	Interpretation
> 2.0	Distorts or degrades the measurement system
1.5 – 2.0	Unproductive for construction of measurement, but not degraded
0.5 – 1.5	Productive for measurement
< 0.5	Unproductive for measurement, but not degrading. May produce misleadingly good reliabilities and separations

3.4 Inter-Rater Reliability

Tables 4.39 through 4.44 (pages 214-219) contain information about the scoring agreement between two ratings received for each item. When the two readers assigned the same score to a student's answer, the scores were in perfect agreement. Scores differed by one score point were adjacent, and scores differing by two or more score points were in discrepancy. For further information about inter-rater agreement, please see chapter 1.6, *Scoring Procedures of the 2008 MSA-Math*. For the 2008 MSA-Math, the adjacent agreement rates were above 99%, and perfect agreement rates were above 96% for Step A and above 80% for Step B for all items across all grades.

3.5 Correlations among Mathematics Standards

The 2008 MSA-Math consisted of five subscore reporting standards: *Algebra, Geometry and Measurement, Statistics and Probability, Numbers and Computations, and Process*. Tables 4.3 through 4.8 (pages 161-163) contain correlation coefficients among these math standards.

3.6 Decision Accuracy and Consistency at the Cut Scores

Tables 4.9 through 4.14 (pages 164-166) contain the results of analyses performed to estimate the accuracy and consistency of the decisions for passing (proficient) on the 2008 MSA-Math. The analyses make use of the methods outlined and implemented in Livingston and Lewis (1995), Haertel (1996), and Young and Yoon (1998).

The *accuracy* of a decision is the extent to which it would agree with the decisions that would be made if each student could somehow be tested with all possible parallel forms of the assessments. The *consistency* of a decision is the extent to which it would agree with the decisions that would be made if the students had taken a different form of the examination, equal in difficulty and covering the same content as the form they actually took.

Students can be misclassified in one of two ways. Students who were below the proficiency cut score, but were classified (on the basis of the assessment) as being above a cut score, are considered to be *false positives*. Students who were above the proficiency cut score, but were classified as being below a cut score, are considered to be *false negatives*.

For the 2008 MSA-Math, Tables 4.9 through 4.14 include:

- Performance level
- Accuracy classifications
- False positives
- False negatives
- Consistency classifications

The tables illustrate the general rule that decision consistency was less than decision accuracy.

3.7 Differential Item Functioning

This section provides information about *differential item functioning (DIF)* analyses used for the 2008 MSA-Math. While the *reference* group was either male or Caucasian students, the *focal* group was either female or African-American students. It should be noted that DIF analyses on the 2008 operational items indicated that all the items were satisfactory. All the DIF results were archived in the 2008 Maryland item bank.

Since the 2008 MSA-Math was a mixed-format examination, comprised of *SR*, *SRP*, *BCR*, and *ECR* items, the *DIF* procedure used consists of the Mantel Chi-square (Mantel, 1963) for the *BCR* and *ECR* items and the Mantel-Haenszel procedure (Mantel & Haenszel, 1959) for the *SR* and *SRP* items.

Brief Constructed Response (BCR) and Extended Constructed Response (ECR) Items

To help interpret the Mantel Chi-square (Mantel χ^2), the Educational Testing Service (ETS) *DIF* procedure uses the Mantel statistic in conjunction with the *standardized mean difference (SMD)*.

Mantel Statistic

The Mantel χ^2 is simply a conditional mean comparison of the ordered response categories for reference and focal groups combined over values of the matching variable score. By “ordered” we mean that a response of 1 on an item is higher than 0, a response of 2 is higher than 1, and so on. “Conditional,” on the other hand, refers to the comparison of members from the two groups who received the same score on the matching variable, i.e., the total test score in our analysis.

Table 3.2 shows a $2 \times T \times K$ contingency table, where T is the number of response categories and K is the number of levels of the matching variable. The values, y_1, y_2, \dots, y_T are the T scores that can be gained on the item. The values, n_{Ftk} and n_{Rtk} , represent the numbers of focal and reference groups who are at the k^{th} level of the matching variable and gain an item score of y_t . The “+” indicates total number over a particular index (Zwick, Donoghue, & Grima, 1993).

Table 3.2 $2 \times T$ Contingency Table at the k^{th} level

Group	Item Score				Total
	y_1	y_2	...	y_T	
Reference	n_{R1k}	n_{R2k}	...	n_{RTk}	n_{R+k}
Focal	n_{F1k}	n_{F2k}	...	n_{FTk}	n_{F+k}
Total	n_{+1k}	n_{+2k}	...	n_{+Tk}	n_{++k}

Note. This table was cited from Zwick, et al. (1993)

The Mantel statistic is defined as the following formula:

$$\text{Mantel } \chi^2 = \frac{(\sum_k F_k - \sum_k E(F_k))^2}{\sum_k \text{Var}(F_k)}$$

where

F_k = the sum of scores for the focal group at the k^{th} level of the matching variable and is defined as follows:

$$F_k = \sum_t y_t n_{Ftk}$$

The expectation of F_k under the null hypothesis is

$$E(F_k) = \frac{n_{F+k}}{n_{++k}} \sum_t y_t n_{+tk}.$$

And, the variance of F_k under the null hypothesis is as follows:

$$\text{Var}(F_k) = \frac{n_{R+k} n_{F+k}}{n_{++k}^2 (n_{++k} - 1)} \left[(n_{++k} \sum_t y_t^2 n_{+tk}) - (\sum_t y_t n_{+tk})^2 \right].$$

Under H_0 , the Mantel statistic has a chi-square distribution with one degree of freedom. In *DIF* applications, rejecting H_0 suggests that the students of the reference and focal groups who are similar in overall test performance tend to differ in their mean performance. In the case of dichotomous items, on the other hand, the statistic is identical to the Mantel-Haenszel (1959) statistic without the continuity correction (Zwick, Donoghue, & Grima, 1993).

Standardized Mean Difference (SMD)

A summary statistic to accompany the Mantel approach is the *standardized mean difference (SMD)* between the reference and focal groups proposed by Dorans and Schmitt (1991). This statistic compares the means of the reference and focal groups, adjusting for differences in the distribution of the reference and focal group members across the values of the matching variable.

$$SMD = \sum_k p_{Fk} m_{Fk} - \sum_k p_{Fk} m_{Rk}$$

where

$p_{Fk} = \frac{n_{F+k}}{n_{F++}}$, the proportion of the focal group members who are at the k^{th} level of the matching variable,

$m_{Rk} = \frac{1}{n_{F+k}} \times (\sum_t y_t n_{Ftk})$, the mean item score of the focal group members at the k^{th} level,

and

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

As can be seen from the equation above, the *SMD* is the difference between the unweighted item mean of the focal group and the weighted item mean of the reference group. The weights for the reference group are applied to make the weighted number of the reference group students the same as in the focal group within the same ability. A negative *SMD* value implies that the focal group has a lower mean item score than the reference group, conditional on the matching variable.

DIF classification for BCR and ECR items

The *SMD* is divided by the total group item standard deviation to obtain an effect-size value for the *SMD*. This effect-size *SMD* is then examined in conjunction with the Mantel χ^2 to obtain *DIF* classifications that are depicted in Table 3.3 below.

Table 3.3 DIF Classification for BCR and ECR Items

Category	Description	Criterion
AA	No <i>DIF</i>	Non-significant Mantel χ^2 or Significant Mantel χ^2 and $ SMD/SD \leq .17$
BB	Weak <i>DIF</i>	Significant Mantel χ^2 and $.17 < SMD/SD \leq .25$
CC	Strong <i>DIF</i>	Significant Mantel χ^2 and $.25 < SMD/SD $

Note. SD is the total group standard deviation of the item score in its original metric.

Selected Response (SR) and Student-Produced Response (SPR) Items

For the *SR* and *SPR* items, the Mantel-Haenszel Chi-square (M-H χ^2) is used in conjunction with the M-H odds ratio transferred to what ETS calls the *delta scale* (D).

The Odds Ratio

The odds of a correct response are P/Q or $P/(1-P)$. The odds ratio, on the other hand, is simply the odds of a correct response of the reference group divided by the odds of a correct response of the focal group.

For a given item, the odds ratio is defined as follows:

$$\alpha_{M-H} = \frac{P_r / Q_r}{P_f / Q_f}$$

The corresponding null hypothesis is that the odds of getting the item correct are equal for the two groups. Thus, the odds ratio is equal to 1:

$$H_0: \alpha_{M-H} = \frac{P_r / Q_r}{P_f / Q_f} = 1.$$

The Delta Scale

In order to make the odds ratio symmetrical around zero with its range being in the interval $-\infty$ to $+\infty$, the odds ratio is transformed into a log odds ratio as per the following:

$$\beta_{M-H} = \ln(\alpha_{M-H}).$$

The simple natural logarithm transformation of this odds ratio is symmetrical about zero in which zero has the interpretation of equal odds. This *DIF* measure is a signed index where a positive value signifies *DIF* in favor of the reference group while a negative value indicates *DIF* in favor of the focal group. β_{M-H} also has the advantage of being transformed linearly to other interval scale metrics (Camilli & Shepard, 1994). This fact is utilized by ETS in creating their delta scale (D), which is defined as follows:

$$D = -2.35 \cdot \beta_{M-H}.$$

DIF classification for *SR* and *SPR* items

The following table depicts *DIF* classifications for *SR* items to examine the M-H χ^2 in conjunction with the delta scale (D):

Table 3.4 *DIF* Classification for *SR* and *SPR* Items

Category	Description	Criterion
A	No <i>DIF</i>	Non-significant M-H χ^2 or $ D < 1.0$
C	Strong <i>DIF</i>	Significant M-H χ^2 and $ D \geq 1.5$
B	Weak <i>DIF</i>	Otherwise classified as B

3.8 Equating and Scaling

Tables 4.15 through 4.38 contain the 2008 MSA-Math total and subtotal raw score to scale score (RS/SS) conversion tables. Conditional standard errors for the total and subtotal scale scores are also included.

The Rasch and Partial Credit Models

The most basic expression of the Rasch model is in the *item characteristic curve* (ICC). It shows the probability of a correct response to an item as a function of the ability level. The probability of a correct response is bounded by 1 (certainty of a correct response) and 0 (certainty of an incorrect response).




Figure 3.1 Item Characteristic Curve

As an example, consider Figure 3.1 which depicts an item that falls at approximately 0.85 on the ability (horizontal) scale. When a person answers an item at the same level as their ability, then that person has a probability of roughly 50% of answering the item correctly. Another way of expressing this is that if we have a group of 100 people, all of whom have an ability of 0.85, we would expect about 50% of them to answer the item correctly. A person whose ability was above 0.85 would have a higher probability of getting the item right, while a person whose ability is below 0.85 would have a lower probability of getting the item right. This makes intuitive sense and is the basic formulation of Rasch measurement for test items having only 2 possible categories (i.e., wrong or right).




Figure 3.2 Category Response Curves for a One-Step Item

Figure 3.2 extends this formulation to show the probabilities of obtaining a wrong answer or a right answer. The curve on the left ($j = 0$) shows the probability of getting a score of “0” while the curve on the right ($j = 1$) shows the probability of getting a score of “1”. The point at which the two curves cross indicates the transition point on the ability scale where the most likely response changes from a “0” to a “1”. Here, the probability of answering the item correctly is 50%.

The key step in the formulation, and the point at which the Rasch dichotomous model merges with the PCM, requires us to assume an additional response category. Suppose that, rather than scoring items as completely wrong or completely right, we add a category representing answers that, though not totally correct, are still clearly not totally incorrect. These relationships are shown in Figure 3.3.

The left-most curve ($j = 0$) in Figure 3.3 represents the probability for all examinees getting a score of “0” (completely incorrect) on the item, given their ability. Those of very low ability (i.e., below -2) are very likely to be in this category and, in fact, are more likely to be in this category than the other two. Those receiving a “1” (partial credit) tend to fall in the middle range of abilities (the middle curve, $j = 1$). The final, right-most curve ($j = 2$) represents the probability for those receiving scores of “2” (completely correct). Very high-ability people are clearly more likely to be in this category than in any other, but there are still some of average and low ability that can get full credit for the item.




Figure 3.3 Category Response Curves for a Two-Step Item

Although the actual computations are quite complex, the points at which lines cross each other have a similar interpretation as for the dichotomous case. Consider the point at which the $j = 0$ line crosses the $j = 1$ line, indicated by the left arrow. For abilities to the left of (or less than) this point, the probability is greatest for a “0” response. To the right of (or above) this point, and up to the point at which the $j = 1$ and $j = 2$ lines cross (marked by the right arrow), the most likely response is a “1”. For abilities to the right of this point, the most likely response is a “2”.

Note that the probability of scoring a “1” response ($j = 1$) declines in both directions as ability decreases to the low extreme or increases to the high extreme. These points then may be thought of as the difficulties of crossing the *thresholds* between categories.

An important implication of the formulation can be summarized as follows: If the commonly used Rasch model applied to dichotomously (right/wrong) scored items can be thought of as simply a

special case of the PCM, then the act of scaling multiple-choice items together with polytomous items, whether they have three or more response categories, is a straightforward process of applying the measurement model. The quality of the scaling can then be assessed in terms of known procedures.

One important property of the PCM is its ability to separate the estimation of item/task parameters from the person parameters. With the PCM, as with the Rasch model, the total score given by the sum of the categories in which a person responds is a sufficient statistic for estimating person ability (i.e., no additional information need be estimated). The total number of responses across examinees in a particular category is a sufficient statistic for estimating the step difficulty for that category. Thus with PCM, the same total score will yield the same ability estimate for different examinees.

The PCM is a direct extension of the dichotomous one-parameter logistic *IRT* model developed by Rasch (Rasch, 1980). For an item/task involving m_i score categories, one general expression for the probability of scoring x on item/task i is given by

$$P_{mix} = \exp \sum_{j=0}^x (\theta_n - D_{ij}) / \sum_{k=0}^{m_i} \left[\exp \sum_{j=0}^k (\theta_n - D_{ij}) \right] \quad x = 0, 1, \dots, m_i,$$

where $\sum_{j=0}^0 (\theta - D_{ij}) = 0$ and

$$\exp \sum_{j=0}^0 (\theta - D_{ij}) = 1.$$

The above equation gives the probability of scoring x on the i -th test item as a function of ability (θ) and the difficulty of the m_i steps of the task (Masters, 1982).

According to this model, the probability of an examinee scoring in a particular category (step) is the sum of the logit (log-odds) differences between θ and D_{ij} of all the completed steps, divided by the sum of the differences of all the steps of a task. Thissen and Steinberg (1986) refers to this model as a divide-by-total model. The parameters estimated by this model are (1) an ability estimate for each person (or ability estimate at each raw score level) and (2) m_i threshold (difficulty) estimates for each task with $m_i + 1$ score categories.