



Guide to Item Analysis

Introduction

Item Analysis (a.k.a. Test Question Analysis) is a useful means of discovering how well individual test items assess what students have learned. It helps us to answer the following questions.

- Is a particular question as difficult, complex, or rigorous as it should be?
- Does the item do a good job of separating students who know the content from those who guess the right answer or apply test-taking strategies to eliminate the wrong answers?
- Which items should be eliminated or revised before subsequent administrations of the test?

By analyzing item performance over time and making necessary adjustments, you can improve test score validity and reliability. Test items can be systematically analyzed regardless of whether they are administered as a [Canvas](#) assignment or if they are submitted as “bubble sheets” to [Scanning Services](#).

With this guide, you’ll be able to accomplish the following.

- Define and explain the indices related to item analysis.
- Locate each index of interest within Scanning Services’ Exam Analysis reports.
- Identify target values for each index, depending upon your testing intentions.
- Make informed decisions while creating items for new tests, as well as decisions about retaining, revising, or removing items while revising tests.

Anatomy of a Test Item

In this guide, we refer to the following terms to describe the items (or questions) that make up multiple-choice tests. See example item below.

1. **Stem** refers to the portion of the item that presents a problem for the respondents (students) to solve.
2. **Options** refers to the various ways the problem might be solved, from which respondents select their answer.
 - a. **Distractor** is an incorrect option.
 - b. **Key** is a correct option.

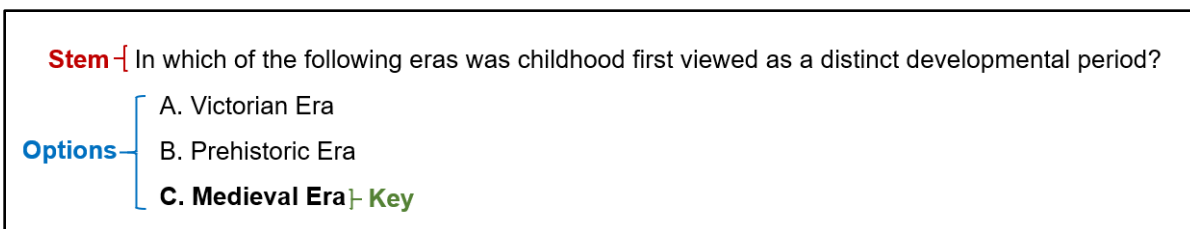


Figure 1: Anatomy of a test item example

Item Analysis in Canvas

By default, the quiz summary function in Canvas shows high score, low score, mean score, standard deviation (how far the values are spread across the entire score range), mean elapsed time of quiz completion, and more. You automatically have access to the results of a quiz once it's been completed by one or more students, and you can sort those results by Student Analysis or Item Analysis. The Canvas Doc Team offers a number of guides on using these functions in that learning management system. Click on [Search the Canvas Guides](#) under the Help menu and enter "Item Analysis" for the most current information.

Item Analysis in Scanning Services

Scanning Services offers an Exam Analysis Report through its [Instructor Tools web site](#). (It would be helpful to refer to the [Exam Analysis Report Example](#) as you continue reading.) Learn how to generate and download the report at [Scanning Services Instructor Tools Help](#).

Four Steps to Item Analysis

Item analysis typically focuses on four major pieces of information: test score reliability, item difficulty, item discrimination, and distractor information. No single piece should be examined independent of the others. In fact, understanding how to put them all together to help you make a decision about the item's future viability is critical.

Reliability

Test Score Reliability is an index of the likelihood that scores would remain consistent over time if the same test was administered repeatedly to the same learners. Both Canvas' Item Analysis and Scanning Services' Exam Analysis Report use the [Cronbach's Alpha](#) measure of internal consistency, which provides reliability information about items scored dichotomously (i.e., correct/incorrect), such as multiple-choice items. A test showing a Chronbach's Alpha score of .80 and higher has less measurement error and is thus said to have very good reliability. A value below .50 is considered to have low reliability.

Item Reliability is an indication of the extent to which your test measures learning about a single topic, such as "knowledge of the battle of Gettysburg" or "skill in solving accounting problems." Measures of internal consistency indicate how well the questions on the test consistently and collectively address a common topic or construct.

In **Scanning Services' Exam Analysis Report**, next to each item number is the percentage of students who answered the item correctly.

To the right of that column, you'll see a breakdown of the percentage of students who selected each of the various **options** (A, B, C, D, etc.), including the **key** (in dark grey) and the **distractors**. Under each **option**, the **Total (TTL)** indicates the total number of students who selected that option. The **Reliability coefficient (R)** value shows the mean score (%) and Standard Deviation of scores for a particular distractor.

How would you use this information?

Score Reliability is dependent upon a number of factors, including some that you can control and some that you can't.

Factor	Why it's important
Length of the test	Reliability improves as more items are included.
Proportion of students responding correctly and incorrectly to each item	Helps determine item reliability.

Factor	Why it's important
Item difficulty	Very easy and very difficult items do not discriminate well and will lower the reliability estimate.
Homogeneity of item content	Reliability on a particular topic improves as more items on that topic are included. This can present a challenge when a test seeks to assess a lot of topics. In that case, ask questions that are varied enough to survey the topics, but similar enough to collectively represent a given topic.
Number of test takers	Reliability improves as more students are tested using the same pool of items.
Factors that influence any individual test taker on any given day	Distraction, lack of sleep, physical illness, emotional distress, test anxiety, etc. can affect even the most prepared students' ability to choose the correct option.

What should you aim for?

Reliability coefficients range from 0.00 to 1.00. Ideally, score reliability should be above 0.80. Coefficients in the range 0.80-0.90 are considered to be very good for course and licensure assessments.

Difficulty

Item Difficulty represents the percentage of students who answered a test item correctly. Low item difficulty values (e.g., 28, 56) indicate difficult items, since only a small percentage of students answer correctly. Conversely, high item difficulty values (e.g., 84, 96) indicate easier items, as a greater percentage of students answer correctly.

In Scanning Services' Exam Analysis Report, there are two numbers in the **Item** column: item number and the percentage of students who answered the item correctly. A higher percentage indicates an item that is easier for students to answer; a lower percentage indicates a more difficult item. That percentage helps to gauge this difficulty index against what you expect and how difficult you'd like the item to be. You should find a higher percentage of students correctly answering items you think should be easy and a lower percentage correctly answering items you think should be difficult. Low difficulty items should not be immediately abandoned or adjusted if your goal is to test recall, particularly if you intend for those items to help students to [retrieve information](#) to answer more complex questions later in the test.

Item difficulty is also important as you try to determine how well an item separates students who know the content from those who do not (see **Item Discrimination** below). Certain items do not discriminate well. Very easy questions and very difficult questions, for example, are poor discriminators. That is, when most students get the answer correct, or when most answer incorrectly, it is difficult to ascertain who really knows the content, versus those who are guessing.

How should you use this information?

As you examine the difficulty of the items on your test, consider the following.

1. Which items did students find to be easy; which did they find to be difficult? Do those items match the items you anticipated would be easy/difficult for students? Sometimes, for example, an instructor may place an item

on a test believing it to be one of the easier items on the exam when, in fact, even the best performing students find it to be challenging, which may indicate that the item or its options are confusing or misleading.

2. Very easy items and very difficult items don't do a good job of discriminating between students who know the content and those who do not. (The section on **Item Discrimination** discusses this further.) However, you may have very good reason for using either type of question on your exam. For example, some instructors deliberately start their exam with an easy question or two to reduce test anxiety levels or to help students recall foundational information, like data and definitions, for use in subsequent, more challenging items.

What should you aim for?

Popular consensus suggests that the best approach is to aim for a mix of difficulties, i. e. a few very difficult, some difficult, some moderately difficult, and a few easy. However, the level of difficulty should be consistent with the degree of difficulty of the concepts being assessed. We offer the following as rough estimates.

- 20% very difficult items
- 40% difficult items
- 30% moderately difficult items
- 10% easy items

Discrimination

Item Discrimination is the degree to which students with high overall exam scores also got a particular item correct. It is often referred to as Item Effect, since it is an index of an item's effectiveness at discriminating those who know the content from those who do not.

The **Point Biserial correlation coefficient (PBS)** provides this discrimination index. Its possible range is -1.00 to 1.00. A strong and positive correlation suggests that students who get a given question correct also have a relatively high score on the overall exam. Theoretically, this makes sense. Students who know the content should be able to perform well on the test overall. There's a problem, however, if students who don't actually know the content perform well on the overall test.

One would expect that the students who did well on the exam selected the correct response, thus generating a higher mean score and a higher PBS which shows the correlation between a high overall exam score to a given correct response to an item. Conversely, cases where an incorrect response distracted students who otherwise did well on the exam, exhibited by a high R value, should result in a lower PBS score.

PBS score ranges from -1.0 to 1.0, with a minimum desired score greater than 0.15. If a single test is weighted heavily as part of students' grades, reliability must be high. Low score reliability is an indication that, if students took the same exam again, they might get a different score. Optimally, we would expect to see consistent scores on repeated administrations of the same test.

In Scanning Services' Exam Analysis Report, you'll find the PBS final column color-coded, so you can easily distinguish the items that may require revision. Likewise, the key for each item is color-coded grey.

How should you use this information?

As you examine item discrimination, there are a number of things you should consider.

1. Very easy or very difficult items are not good discriminators. If an item is so easy (e.g., difficulty = 98) that nearly everyone gets it correct or so difficult (e.g., difficulty = 12) that nearly everyone gets it wrong, then it becomes very difficult to discriminate those who actually know the content from those who do not.
2. That does not mean that all very easy and very difficult items should be eliminated. In fact, they are viable as long as you are aware that they will not discriminate well and if putting them on the test matches your intention to either really challenge students or to make certain that everyone knows a certain bit of content.
3. Nevertheless, a poorly written item will have little ability to discriminate.

What should you aim for?

It is typically recommended that item discrimination be at least 0.15. It's best to aim even higher. Items with a negative discrimination are theoretically indicating that either the students who performed poorly on the test overall got the question correct or that students with high overall test performance did *not* get the item correct. Thus, the index could signal a number of problems.

1. There is a mistake on the scoring key.
2. Poorly prepared students are guessing correctly.
3. Well prepared students are somehow confused or misled by the wrong answer.

In those cases, action must be taken. Items with negative item difficulty must be addressed. Also, reliability improves when students are provided with a total of 4 options that include only 1 key and 3 distractors. Be certain that there is only one possible answer, that the question and options are written clearly and are grammatically correct, and that your answer key is correct.

Distractors

Distractors are the multiple-choice response options that are *not* the correct answer. They are plausible but incorrect options that are often developed based upon students' common misconceptions or miscalculations to see if they've moved beyond them. As you examine distractors, there are a number of things you should consider.

1. Are there at least some respondents for each distractor? If you have 4 possible options for each item but students are selecting from between only one or two of them, it is an indication that the other distractors are ineffective. Even low-knowledge students can reduce the "real" options to one or two, so the odds are now good that they will choose correctly. In short, this situation increases the possibility of students without knowledge guessing correctly.
2. It is not necessary to revisit every single "0" in the response table. Instead, be mindful, and responsive, where it looks as if distractors are ineffective. Typically, this is where there are two or more distractors selected by no one.
3. Are the distractors overly complex, vaguely worded, or contain obviously wrong, "jokey" or "punny" content? Distractors should not be mini-tests in themselves, nor should they be a waste of effort. Ineffective distractors can be a strong demotivating factor for students that last beyond their initial appearance, since they often lead to questions of fairness.
4. Theoretically, distractors should not be selected by students who have a good understanding of the material. It's important to determine whether the best performing students on the exam are, for some reason, being drawn

to an incorrect answer. Either the distractor is ambiguous or unclear, or the topic deserves more attention during instruction.

What should you aim for?

Distractors should be plausible options. Test writers often use students' misconceptions, mistakes on homework, or missed quiz questions as fodder for crafting distractors. When this is the approach to distractor writing, information about student understanding can be gleaned even from their selection of *wrong* answers.

For best practices on developing effective multiple-choice items, [schedule a consultation](#) or see the following resources.

[Multiple Choice Item Writing Rules](#)

[Writing Multiple-Choice Questions that Demand Critical Thinking](#)

Conclusion

Item analysis is an empowering process. Knowledge of score reliability, item difficulty, item discrimination, and crafting effective distractors can help an instructor make decisions about whether to retain items for future administrations, revise them, or eliminate them from the test item pool. Item analysis can also help an instructor to determine whether a particular portion of course content should be revisited. In any case, all indices should be considered together before making decisions or revisions. One important thing to always keep in mind is that decisions about item revision should be based on the extent to which item performance matches your intent for the item and your intent for the overall exam.

References

- Suskie, L. (2017). "Making Multiple Choice Tests More Effective." Schreyer Institute for Teaching Excellence. The Pennsylvania State University.
- Understanding Item Analyses. (2018). Office of Educational Assessment. University of Washington. Retrieved from washington.edu/assessment/scanning-scoring/scoring/reports/item-analysis.
- Advanced Research Computing Statistical Methods and Data Analytics. (2024). What does Cronbach's alpha mean? University of California Los Angeles. Retrieved from stats.oarc.ucla.edu/spss/faq/what-does-cronbachs-alpha-mean.
- The Learning Scientists. (2017). Episode 2: Retrieval Practice. Retrieved from learningscientists.org/learning-scientists-podcast/2017/9/6/episode-2-retrieval-practice?rq=retrieval%20practice.
- Office of Educational Assessment. (2024). Understanding Item Analyses. University of Washington. Retrieved from washington.edu/assessment/scanning-scoring/scoring/reports/item-analysis.
- Boston University School of Public Health. Writing Multiple-choice Questions that Demand Critical Thinking. Retrieved from sphweb.bumc.bu.edu/otlt/teachingLibrary/Assessment/WritingMultiple.pdf.