

A Non-Technical Approach for Illustrating Item Response Theory

Chong Ho Yu, Angel Jannasch-Pennell, & Samuel DiGangi

Arizona State University

Journal of Applied Testing Technology

Feb 21, 2008

RUNNING HEAD: IRT tutorial

Corresponding author:

Chong Ho Yu, Ph.D.

PO Box 612

Tempe AZ 85280

Email: chonghoyu@gmail.com

Abstract

Since the introduction of the *No Child Left Behind Act*, assessment has become a pre-dominant theme in the US K-12 system. However, making assessment results understandable and usable for the K-12 teachers has been a challenge. While test technology offered by various vendors has been widely implemented, technology of training for test development seems to be under-developed. The objective of this presentation is to illustrate a well-designed interactive tutorial for understanding the complex concepts of Item Response Theory (IRT). The approach of this tutorial is to dissociate IRT from Classical Test Theory (CTT) because it is the belief of the authors that the mis-analogy between IRT and CTT could lead to misconceptions. Initial user feedback is collected as input for further refining the program.

Keywords: Item response theory, assessment, measurement, classical test theory, multimedia

A non-technical approach for illustrating Item Response Theory

Introduction

Since the introduction of the *No Child Left Behind Act* (NCLB) in January 2001, assessment has become a predominant theme in grades K-12. The goal of NCLB is to bring all US students up to a level of academic proficiency within a 15-year period. Schools that fail to make adequate yearly progress for students' achievement for a period of three years will eventually be restructured or even taken over by the state (Goertz & Duffy, 2003). As a result of this high-stakes assessment, many school districts have taken it upon themselves to develop their own assessments in order to identify and provide extra assistance to disadvantaged students. However, most test developers within each school district are typically teachers who are not exposed to or have just minimal experience in measurement theories, and already have a full-time responsibility within their classroom during the academic year. Consequently, the items and tests that are being developed may not be a good indicator of students' performance. While test technology offered by various vendors has been widely implemented, technology of training for test development seems to be under-developed. No matter how sophisticated the item bank technology is, human factors, such as lack of knowledge and motivation among teachers and item authors, are major hindrances to successful deployment of item banking. The purpose of this study is to provide a multimedia tutorial that aids teachers in the interpretation of both the individual students' and the class' performance, as well as to identify problems in test authoring which can improve test items for future assessments. The tutorial is designed to help teachers understand and interpret the psychometric analysis of district tests by teaching item response

theory (IRT) (Embretson & Reise, 2000). The computer-based multimedia program is accessible at <http://www.creative-wisdom.com/multimedia/IRTTHA.htm>, and a PDF document (Yu, 2007) that presents much of the program content can also be viewed at <http://www.creative-wisdom.com/computer/sas/IRT.pdf>. Readers are encouraged to go through at least the multimedia version of the tutorial before reading this article, which explains the rationale of the pedagogical strategy and discussed the feedback from students. It is important to emphasize that the target audience for this tutorial includes K-12 teachers and administrators and thus the focus is on IRT concepts rather than computational procedures.

Mis-analogy between CTT and IRT

It is a common pedagogical strategy to use the classical test theory (CTT) as a metaphor in teaching IRT. However, this tutorial does not adopt this approach for it is the firm belief of the developers that CTT and IRT are too diverse to impose conceptual links on them. Using CTT as a stepping stone to learn IRT would lead to dangerous misconceptions, just like using programming concepts of procedural languages (e.g. Pascal) to learn object oriented programming (e.g. C++, Java) will result in disorientation. Next, we will briefly review CTT and IRT to see why they are fundamentally incompatible. CTT could be traced back to Spearman (1904). The equation of CTT is expressed as:

$$Y = T + E$$

where Y is a total number-right or number-keyed score, T represents a true score, and E is an random error, which is independent of T .

Ideally, the true score reflects the exact value of the respondent's ability or attitude. The theory assumes that traits are constant and the variation in observed scores is caused by random errors, which resulted from numerous factors such as guessing and fatigue. These random errors over many repeated measurements are expected to cancel each other out; in the long run the expected mean of measurement errors should be zero. When the error term is zero, the observed score is the true score:

$$Y = T \quad (\because E = 0)$$

Therefore, if an instrument is reliable, the observed scores should be fairly consistent and stable throughout repeated measures. Many popular reliability estimate methods such as split-half, Cronbach coefficient alpha, test-retest and alternate forms are based upon this rationale. In short, reliability in terms of replication plays a central role in CTT. However, in CTT reliability and standard error of measurement (SEM) refer to *test scores*, not to an instrument itself. For this reason the information yielded from CTT is not really about the item attributes, which are supposed to be the focal interest of psychometricians. In "Guidelines for the authors" of *Educational and Psychological Measurement* (EPM), Thompson (1994) asserts that use of wording such as "the reliability of the test" or "the validity of the test" will not be considered acceptable in the journal (p.841). Later, Thompson and Vacha-Haase (2000) went even further to proclaim that "psychometrics is datametrics" (p. 174). Simply stated, reliability attaches to the data rather than the psychological test. However, Thompson and Vacha-Haase's statement is true in CTT only. Item attributes yielded from IRT are not sample or data dependent.

Unlike CTT that emphasizes reliability in terms of a true score that would emerge out of repeated measurements, IRT provides information function and standard errors that describe the precision of a test as an instrument for establishing test-taker ability across the latent trait scale (Doran, 2005). The main reason that CTT is considered a form of datametrics is due to its character of *sample dependence* (whether an item would appear to be difficult or easy is tied to the ability of the examinees). In contrast, the estimation of examinee ability in IRT is item independent, and similarly, the estimation of item attributes is sample independent. Hence, we can assert that the inferences made by IRT are concerned with the instrument rather than the test scores. There are many other incompatible aspects between CTT and IRT. Embretson and Reise (2000) allege that the conventional rules of measurement based on CTT are inadequate and proposed another set of new rules. For example, the conventional theory states that the standard error of measurement applies to all scores in a particular population, but Embretson and Reise found that the standard error of measurement differs across scores but generalizes across populations. When IRT is introduced in the context of CTT, many teachers still associate the item “difficulty” in IRT with the mere percentage of correct answers in CTT. For these reasons, this multimedia tutorial skips CTT altogether in an attempt to avoid any potential misunderstanding.

Program description

This hypermedia tutorial, which is composed of six modules, was developed with the use of Macromedia Captivate®, Macromedia Flash®, Adobe PhotoShop®, SAS®, SPSS®, Microsoft Excel®, and Microsoft PowerPoint®. Hypertext and multimedia are two major features that are commonly found in many computer-assisted tutorials. However, rich media, such as over-use of animation modules, could lead to cognitive overload (Sweller & Chandler,

1991). In addition, improper use of hypertext may interfere with instruction. Without prior knowledge pertaining to the subject matter, non-linear jumping across slides may not lead to a full understanding of the material (Yu, 1993). Hence, contrary to popular practice, the introduction and table of content pages of this tutorial emphasizes the following: “Since some concepts are interrelated, readers are encouraged to go through the tutorial in a sequential manner” (Figure 1).

This tutorial, which is a practical introduction to Item Response Theory (IRT), is composed of six parts:

1	Item calibration and ability estimation
2	Item Characteristic Curve
3	Item Information Function
4	Test Information Function
5	Item-Person Map
6	Misfit

This tutorial is designed for novices, and thus, the orientation of this guide is conceptual and practical. Technical terms and mathematical formulas are omitted as much as possible. Since some concepts are interrelated, readers are encouraged to go through the tutorial in a sequential manner. You can pause and rewind any slide at any moment. Each slide has a navigation bar at the bottom. The function of each button is revealed upon mouse over.




Figure 1. Table of content of the tutorial

The content of the tutorial is based on a guide to IRT (Yu, 2007), which is cross-posted on the lead author’s website and *Scientific Software International* ® website (www.ssicentral.com). The target audience for this program is K-12 teachers who have learned

the basic concepts of statistics. Chapter One is concerned with item calibration and ability estimation whereas Chapter Two pertains to Item Characteristic Curve (ICC). Chapters Three and Four are about Item Information Function and Test Information Function. Chapter Five addresses Item Person Map, while Chapter Six teaches the concept of misfit.

To avoid possible confusion resulting from prior knowledge of CTT, Chapter One highlights the point that estimation of student ability cannot only rely on the percentage of correct answers. It starts with a scaled down, yet simple example: A data set with five items and five students only. Many instructors use real data sets to illustrate item calibration and ability estimation in a complex simulated environment, and as a result, students may experience cognitive overload. Therefore, the example used in this tutorial was simplified to increase understanding of the material. The example in Figure 2 is ideal as no item parameter can be estimated when all students could answer Item 5 correctly because there is no variation in the distribution.




Figure 2. 5x5 item-person matrix

Nevertheless, many successful scientific “thought experiments” start from “idealization,” in which the conditions do not correspond to the real world. In spite of using hypothetical cases,

insight may still be gained when idealization makes every variable so simple that the user may “mentally manipulate” them without difficulty (Brown, 2004). At the end of this section, the tutorial emphasizes that the example is an ideal case that is too good to be true (Figure 3).

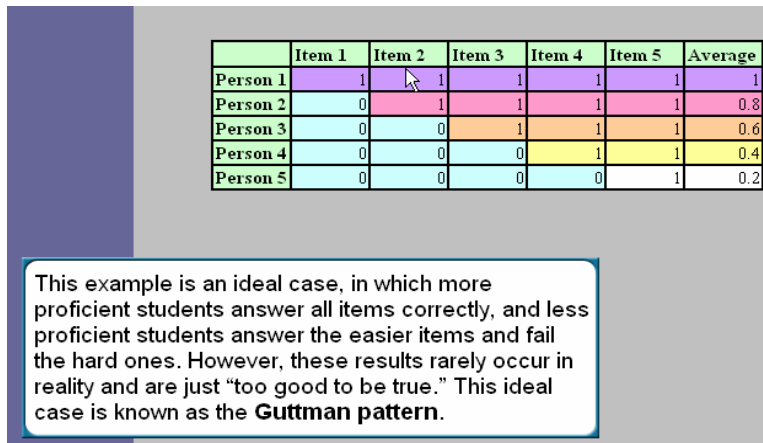


Figure 3. Guttman pattern: Ideal case

In Chapter Two, again we adopt the preceding strategy by presenting theoretical modeling but hiding empirical data. In testing regression, it is a common practice for instructors to overlay the data points and the regression line to offer a visual depiction of residuals. However, it would not work well in this situation, because in IRT there is person misfit and item misfit; in each of these categories there is model fit and individual fit, and these may be further analyzed through infit and outfit. The learner will most likely experience cognitive overload if the model and the data are presented together. Hence, our instructional strategy is to illustrate modeling with simple graphics. For example, Figure 4 shows a typical ICC. The tutorial emphasizes that ICC depicts a theoretical modeling where, for instance, in the actual sample there may be no students with -5 skill level. Nonetheless, these extreme cases in a “what-if”

scenario could clearly illustrate the point that if the person does not know anything about the subject matter, he or she will have zero probability of answering the item correctly.




Figure 4. Item Characteristic Curve (ICC)

The tutorial demonstrates idealization and modeling while also stressing the practical applications of IRT. One of the favorable features of IRT is that the parameter values are centered at zero and thus the visual representation of item difficulty is very easy to interpret. For example, Figure 5 is a screenshot about how IRT can be applied to test construction by selecting items with varying difficulty levels. As you can see, the bars of the average items center around zero, hard items are located at the right side, and easy items are placed on the left side. It is notable that this visually compelling illustration is not used by popular IRT programs, such as Bilog, Winsteps, and RUMM. The bar chart in Figure 5 is generated in a SAS macro code written by Yu (2006) and is imported into the multimedia tutorial.




Figure 5. Item difficulty of all items

Chapter three directs the learners into the realm of item information function (IIF), which is closely related item characteristic curve. As mentioned before, we intend to dissociate IRT from CTT, and thus the typical metaphor between reliability and information function is not mentioned in the chapter. Again, in order to motivate the learners we highlight practical applications of IIF. Figure 6 is a screenshot of how easy items with the most information for the lowest skilled students can be used for diagnostic purposes. Correspondingly, Figure 7 shows how hard items with the most information can be helpful to identify skilled students.

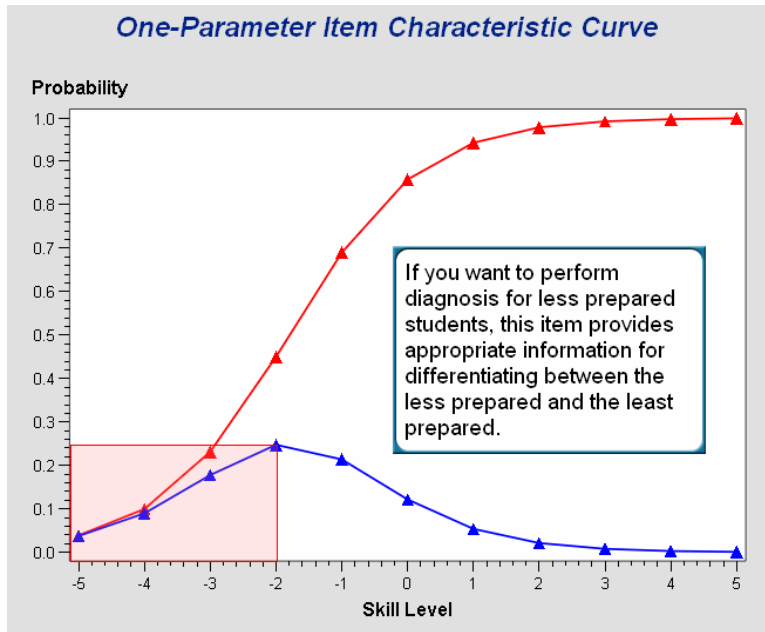


Figure 6. More information for the less skilled students




Figure 7. Less information for the more skilled students

Chapter Five is a logical flow from item level information to exam level information. One of the applications of test information function (TIF) is form balancing, in which alternate test forms are constructed by making the TIFs of all forms equivalent. However, this process is too complicated to be illustrated in this short chapter, and thus only the conceptual aspect of form balancing is introduced (see Figures 8 and 9). Again, it may be tempting to make a comparison between reliability in terms of equivalent alternate form in the classic sense and the balanced forms in the IRT context. But, the developers of this tutorial avoided this metaphor. It is important to point out that the reliability indicator in CTT is for the entire test; there is no individual reliability information for each item. On the other hand, TIF is composed of IFF for each item. Forcing the analogy between equivalent forms and balanced forms may cause misconceptions rather than facilitating learning.




Figure 8. Using TIF for form balancing




Figure 9. Using TIF for form balancing

Chapter Five is concerned with item-person map (IPM), which is a tool to aid teachers to determine the cut score (Figure 10). It may be confusing for K-12 test developers to see that the item and student information are presented on the same scale. Thus, it is inevitable to introduce the concept of logit, a method of data reexpression and rescaling. In this tutorial we use a spreadsheet approach to walk the audience through the simple steps of calculating logit. But, the confusing element is that for person ability rescaling the logit is defined as the natural log of the odd ratio between the number of successes and the failures, while for item difficulty the logit is expressed in terms of the odd ratio between the number of incorrect responses and the number of correct responses (the opposite of the logit in the person dimension). Hence, before introducing IPM, a spreadsheet approach is adopted to illustrate computations of different logits step by step (Figure 11).

The final chapter is arguably the most difficult one because while in CTT, one can use either biserial or point-biserial to determine if the response pattern of a particular item is strongly

correlated to the overall response pattern of all items. In IRT, however, there are many types of misfit indices. These include infit mean square, outfit mean square, infit standardized residual, and outfit standardized residual. Diagnosis of residual-based misfits utilizes distributional information, which is more suitable to examining model fit than item fit. The infit mean-square is a weighted approach, in which a constant is put into the algorithm to dictate the extent to which certain observations are taken into account, whereas the outfit mean-square is unweighted. The infit approach is dominated by unexpected patterns among informative, on-target observations, whereas its outfit counterpart is dominated by unexpected outlying, off-target, low information responses (Linacre & Wright, 1994). On most occasions, test developers aim to develop a test for the on-target population rather than the outliers. To help the audience maintain conceptual clarity, this tutorial addresses only one type of misfit index – infit mean square.

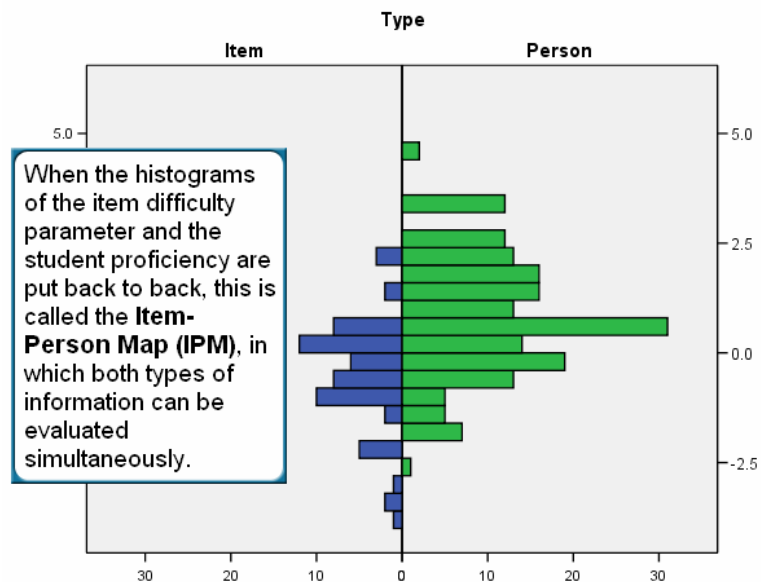


Figure 10. Item person map.

Odd Ratio and Logit for Persons

	Item 1	Item 2	Item 3	Item 4	Item 5	Odd	Prob	Logit
Person 1	1	1	1	1	0	4.00	0.80	0.60
Person 2	1	1	1	0	0	1.50	0.60	0.18
Person 3	1	1	1	0	0	1.50	0.60	0.18
Person 4	1	1	0	0	0	0.67	0.40	-0.17
Person 5	1	1	0	0	0	0.67	0.40	-0.17
Person 6	1	0	0	0	0	0.25	0.20	-0.60
Odd	0.00	0.20	1.00	5.00	#DIV/0!			
Prob	1.00	0.83	0.50	0.17	0.00			
Logit	#NUM!	-0.70	0.00	0.70				

The remaining cells can be filled in using the same approaches. Please keep in mind that it is a simplified example. Usually a real test has more than six students and five items. In the following you will see a graphical illustration of item and student attributes, but more students and items will be used.

Figure 11. A spreadsheet approach to illustrate logit.

Student feedback

To evaluate the efficacy of this multimedia tutorial, 48 K-12 teachers in a Southwestern school district who were involved in item authoring were recruited to take the tutorial and to give feedback in a survey. The survey was developed by our team to measure user attitude towards the tutorial and no sophisticated latent is involved. In addition, no performance test is given because in this initial stage the focus is on collecting feedback for enhancing the tutorial. Table 1 shows the composition of the survey respondents in terms of their prior knowledge. This variable is derived from two questions: “Have you learned the classical true score theory for assessment before?” and “Have you learned the item response theory before?” It indicates that the majority of the participants do not have any background in either CTT or IRT.

Table 1. Prior knowledge of participants.

Prior knowledge	Frequency	Percentage
Classical only	1	2.08
IRT only	2	4.16
classical & IRT	6	12.50
Neither	35	72.91
No response	4	8.33

The authors realized that this is a self-selected sample and also a control group is absent; hence, this evaluation is exploratory in nature and only descriptive statistics (frequency and percentage) of each item is reported, which is shown in Appendix A (Table 2-12 and Figure 11-21). Overall, the respondents tended to agree that the multimedia features of the program are helpful in learning IRT and indeed they felt that after viewing the tutorial they understand IRT better and are motivated to learn more in the future. In response to the open-ended item “What do you think is the best approach to help teachers learn the information provided in the tutorial?” most of the comments are positive, as shown in Appendix B. Many users like the multimedia components (e.g. Observation 2, 5, 21, 26) while there is a major concern with respect to enabling item authors by hand-on experience (e.g. Observation 3, 9, 10, 16, 22, 25).

Conclusion

The multimedia program reflects our pursuit to provide a timely training tool for educational assessment while also enhancing statistical education. We believe that instructional design, especially the conceptual aspect, plays an important role in improving instructional effectiveness. This approach is manifested in disconnecting IRT from CTT. Although the limitations of this small self-selected sample hinder the developers from making a broader generalization and a stronger inference, the positive user feedback encourages the team to further

enhance the tutorial using the current pedagogical approach. Last, the major advantage of using online multimedia tutorials is the ease of constant revision based upon feedback. While evaluation of this program is currently conducted at Arizona State University, use of the application by other institutions and dialogue on this topic are highly encouraged.

References

- Brown, J. R. (2004). Why thought experiments transcend empiricism? In Christopher Hitchcock (Ed.), *Contemporary debates in philosophy of science* (pp. 21-43). MA: Blackwell.
- Doran, H. C. (2005). The information function for the one-parameter logistic model: Is it reliability? *Educational and Psychological Measurement*, 65, 759-769.
- Embretson, S., & Reise, S. (2000). *Item response theory for psychologists*. Mahwah, N.J.: L. Erlbaum Associates.
- Goertz, M., & Duffy, M. (2003). Mapping the landscape of high-stakes testing and accountability programs. *Theory into Practice*, 42, 4-11.
- Linacre J. M., & Wright, B. D. (1994) Chi-Square Fit Statistics. *Rasch Measurement Transactions*, 8(2), 350.
- Spearman, C. (1904). General intelligence: Objectively determined and measured. *American Journal of Psychology*, 15, 201-293.
- Sweller, J., & Chandler, P. (1991). Evidence for cognitive load theory. *Cognition and Instruction*, 8, 351-362.
- Thompson, B. (1994). Guidelines for the authors. *Educational and Psychological Measurement* 54, 837-847.

- Thompson, B., & Vacha-Haase, T. (2000). Psychometrics is datametrics: The test is not reliable. *Educational and Psychological Measurement, 60*, 174–195.
- Yu, C. H. (1993). *Use and effectiveness of navigational aids in a hypertext system*. Unpublished master's thesis, University of Oklahoma, Norman, OK.
- Yu, C. H. (2006). SAS programs for generating Winsteps control files and web-based presentations. *Applied Psychological Measurement, 30*, 247-248.
- Yu, C. H. (2007). *A simple guide to the Item Response Theory*. Retrieved May 1, 2007, from <http://www.ssicentral.com/irt/resources.html>

Appendix A.

Descriptive statistics of the responses to survey items.

Table 2. Descriptive statistics of Q1: Learning on a computer keeps me motivated and interested in the material.

	Frequency	Percentage
Strongly disagree (1)	1	2.08
Disagree (2)	4	8.33
Neutral (3)	13	27.08
Agree (4)	22	45.83
Strongly agree (5)	8	16.67

Figure 11. Bar chart of descriptive statistics of Q1: Learning on a computer keeps me motivated and interested in the material.




Table 3. Descriptive statistics of Q2: The definition of terminology in the tutorial is clear and easy to understand.

	Frequency	Percentage
Strongly disagree (1)	0	0.00
Disagree (2)	1	2.08
Neutral (3)	7	14.58
Agree (4)	30	62.50
Strongly agree (5)	10	20.83

Figure 12. Bar chart of descriptive statistics of Q2: The definition of terminology in the tutorial is clear and easy to understand.




Table 4. Descriptive statistics of Q3: The graphs and tables are easy to interpret and make the explanations easier to understand.

	Frequency	Percentage
Strongly disagree (1)	0	0.00
Disagree (2)	0	0.00
Neutral (3)	3	6.25
Agree (4)	34	70.83
Strongly agree (5)	11	22.92

Figure 13. Bar chart of descriptive statistics of Q3: The graphs and tables are easy to interpret and make the explanations easier to understand.




Table 5. Descriptive statistics of Q4: The information presented in this tutorial is easy to understand.

	Frequency	Percentage
Strongly disagree (1)	0	0.00
Disagree (2)	2	4.17
Neutral (3)	6	12.50
Agree (4)	33	68.75
Strongly agree (5)	7	14.58

Figure 14. Bar chart of descriptive statistics of Q4: The information presented in this tutorial is easy to understand.




Table 6. Descriptive statistics of Q5: Having the text to read along helped me learn.

	Frequency	Percentage
Strongly disagree (1)	0	0.00
Disagree (2)	0	0.00
Neutral (3)	5	10.42
Agree (4)	23	47.92
Strongly agree (5)	20	41.67

Figure 15. Bar chart of descriptive statistics of Q5: Having the text to read along helped me learn.




Table 7. Descriptive statistics of Q6: Having the audio to listen to helped me learn.

	Frequency	Percentage
Strongly disagree (1)	1	2.08
Disagree (2)	2	4.17
Neutral (3)	2	4.17
Agree (4)	19	39.58
Strongly agree (5)	24	50.00

Figure 16. Bar chart of descriptive statistics of Q6: Having the audio to listen to helped me learn.




Table 8. Descriptive statistics of Q7: There is too much information presented in this program.

	Frequency	Percentage
Strongly disagree (1)	2	4.17
Disagree (2)	18	37.50
Neutral (3)	19	39.58
Agree (4)	7	14.58
Strongly agree (5)	2	4.17

Figure 17. Bar chart of descriptive statistics of Q7: There is too much information presented in this program.




Table 9. Descriptive statistics of Q8: The information presented in this program will be useful for test item developers in the school.

	Frequency	Percentage
Strongly disagree (1)	0	0.00
Disagree (2)	2	4.17
Neutral (3)	6	12.50
Agree (4)	31	64.58
Strongly agree (5)	9	18.75

Figure 18. Bar chart of descriptive statistics of Q8: The information presented in this program will be useful for test item developers in the school




Table 10. Descriptive statistics of Q9: I'm bored with the tutorial.

	Frequency	Percentage
Strongly disagree (1)	5	10.42
Disagree (2)	18	37.50
Neutral (3)	21	43.75
Agree (4)	4	8.33
Strongly agree (5)	0	0.00

Figure 19. Bar chart of descriptive statistics of Q9: I'm bored with the tutorial.




Table 11. Descriptive statistics of Q10: I now have a good understanding of the information presented in the program.

	Frequency	Percentage
Strongly disagree (1)	0	0.00
Disagree (2)	6	12.50
Neutral (3)	10	20.83
Agree (4)	31	64.58
Strongly agree (5)	1	2.08

Figure 20. Bar chart of descriptive statistics of Q10: I now have a good understanding of the information presented in the program.





Table 12. Descriptive statistics of Q11: I would like to learn more about IRT after viewing the tutorial.

	Frequency	Percentage
Strongly disagree (1)	0	0.00
Disagree (2)	8	16.67
Neutral (3)	12	25.00
Agree (4)	23	47.92
Strongly agree (5)	5	10.42

Figure 21. Bar chart of descriptive statistics of Q11: I would like to learn more about IRT after viewing the tutorial.



Appendix B

Open-ended responses to “What do you think is the best approach to help teachers learn the information provided in the tutorial?”

Obs.	Feedback
1	I think this worked well.
2	I think this is the best approach. Even people who have trouble with computers can learn this way. It's visual and also auditory which is good for different types of learners.
3	I think you need to give the teachers actual questions and see how you decide the questions are too difficult or too easy.
4	More than one example and clearer explanation.
5	Charts and audio are valuable.
6	More examples and reminders of the terminology, which may be unfamiliar to many teachers.
7	The tutorial seems ok.
8	Completing the tutorials.
9	Hands on work with the information and what it means.
10	I think for teachers it would work best if they had the actual data for a test item on the SAMS and then applied this knowledge to the item. Experiential vs. theory.
11	The tutorial is great but would also require some follow-up interpersonal discussion with colleagues.
12	The tutorial is good.
13	Hands on through the tutorial.
14	Having the teachers take the tutorials - an on-going basis.
15	Not sure there is a best/better approach. If teachers do not buy into the whole SAMs testing approach and the need to battle with outdated technologies at school.... you can present this information all you want, you will still get back blank stares.

16	Give them hands-on practicing it while introducing it. Have them use their own data or use made up data and practice the usage.
17	This is a good way.
18	buy in.....I've been in the business of math education and assessment practices for so long... and watched different assessments come and go....
19	I like the computer tutorial, so I can learn at my own pace.
20	Simplify it to what is needed.
21	The audio was great.
22	Look at their own test data.
23	The approach you have taken is good.
24	A generalized overview of IRT is really the only concept that needs to be presented. The real benefit of the program is the data bank and analysis of test responses.
25	Use one of the taskforce meetings to work through the tutorials and have the taskforce talk about the information and then have them decide on who wants to analyze the data.
26	I liked seeing the words on the screen as I heard it with the headphones.

Acknowledgments

Special thanks to Lori Long, Zeynep Kilic, Wen-juo Lo, Victoria Stay, Chang Kim, Kristina Zeller, Seok Kang, and Charles Kaprolet for their valuable input to the development of this tutorial.