

Test Scaling and Value-Added Measurement

Dale Ballou
Vanderbilt University
March, 2008

Revised June, 2008

The author thanks Kurt Scheib and Warren Langevin for their expert research assistance.

Abstract

As currently practiced, value-added assessment relies on a strong assumption about the scales used to measure student achievement, namely that these are interval scales, with equal-sized gains at all points on the scale representing the same increment of learning. Many of the metrics in which test results are expressed do not have this property (e.g., percentile ranks, normal curve equivalents). However, this property is claimed for the scales obtained when tests are scored according to Item Response Theory.

This claim requires that examinees and test items constitute, in the terminology of representational measurement theory, a conjoint structure. Unfortunately, it is difficult to confirm that this condition is met. In addition, end users of the data lack access to item-level data to test these assumptions themselves. The best they can do is to check the plausibility of the resulting scales. On this count IRT scales often do poorly. Reasonable rescalings have a substantial impact on students' measured growth.

Methods of ordinal data analysis can be employed instead, on the weaker assumption that IRT scales permit us to rank students. Instead of comparing mean achievement of a teacher's students to the students of a (hypothetical) average instructor, ordinal analysis asks what fraction of the former outperforms the latter. The feasibility of such methods for value-added analysis is demonstrated. It is seen that value-added estimates are sensitive to the choice of ordinal methods over conventional techniques.

Clearly, if IRT scores are an interval-scaled variable, ordinal methods throw away valuable information. Practitioners of value-added measurement should ask themselves, however, whether they are so confident of the metric properties of these scales that they are willing to attribute differences between regression-based estimates of value added and estimates based on ordinal analysis to the superiority of the former.

I. Introduction

Models currently used for value-added assessment of schools and teachers require that the scale on which achievement is measured be one of equal units: the five-point difference between scores of 15 and 20 must represent the same gain as the five-point difference between 25 and 30. If it does not, we will end up drawing meaningless conclusions about such matters as the average level of achievement, relative gains of different groups of students, etc., in that the truth of these conclusions will depend on arbitrary scaling decisions.

A scale that possesses this property is known as an interval scale. It is clear that a simple number-right score is not an interval scale of achievement, as test questions are not of equal difficulty. The same is true of several popular methods of standardizing raw test scores that also fail to account for the difficulty of test items, such as percentiles, normal curve equivalents, or grade-level equivalents normed to a nationally-representative sample.¹

The search for measures of achievement that are independent of the particular items included on a test led to the development in the 1950's of Item Response Theory. IRT is now used to score most of the best-known and most widely administered achievement tests today, such as the CBT/McGraw-Hill Terra Nova series, the SAT, the Stanford Achievement Test, and the National Assessment of Educational Progress. IRT was regarded as a significant advance over earlier scaling methods for the following reasons: (1) The score of an examinee is not dependent on the difficulty of the items on the test, provided the test is not so easy that the examinee answers all items correctly or

so hard that he misses them all; (2) An examinee's score is not dependent on the characteristics of the other students taking the same test; (3) Finally, according to some psychometricians, an examinee's score is an interval-scaled measure of ability². This last claim makes IRT scaling particularly appealing to those practicing value-added assessment.

However, not all psychometricians share these views, and the literature contains many confusing and contradictory statements about the properties of IRT scales. Because many social scientists using test scores to evaluate educational institutions and policies have little or no training in measurement theory, the first objective of this paper is to review the issues. The next section describes IRT. It is followed by a summary of the controversy over scale type in Section III. While in the right conditions, IRT yields interval-scaled measures of achievement, IRT scales are often at odds with common sense notions about the effects of schooling and the distribution of ability as students advance through school. I argue in Section IV that we are rightly suspicious of IRT scales when we see such results. Section V takes up the implications for value-added assessment, with particular attention to methods of ordinal data analysis. Concluding remarks appear in Section VI.

¹ These last methods also have the disadvantage of being dependent on the distribution of ability in the tested population or must rely on arbitrary assumptions about this distribution.

² The psychometric literature uses the term ability. Other social scientists might prefer achievement. It represents the student's mastery of the domain of the test and should not be confused with innate ability as opposed to knowledge and skills acquired through education.

II. Item Response Theory and Ability Scales

In IRT, the probability that student i correctly answers test item j correctly is a function of an examinee trait (conventionally termed ability) and one or more item parameters.³ Thus

$$(1) \quad P_{ij} = \text{Prob}[h(\theta_i, \delta_j) > u_{ij}] = F(\theta_i, \delta_j),$$

where θ_i is ability of person i , δ_j is a characteristic (possibly vector-valued) of item j , and u_{ij} is an idiosyncratic person-item interaction, as a result of which individuals of the same level of ability need not answer a given item alike. The u_{ij} are taken to be independent and identically distributed random errors. The function h expresses how ability and item parameters interact. F is derived from h and assumptions about the distribution of u_{ij} . Common assumptions are that the u_{ij} are normal or logistic. θ_i and δ_j are estimated by maximum likelihood methods.

When there is a single item parameter, the assumption that the u_{ij} are logistic gives rise to the one-parameter logistic model (also known as the Rasch model):

$$(2a) \quad P_{ij} = [1 + \exp(-D(\theta_i - \delta_j))]^{-1},$$

in which D is an arbitrary scaling parameter, invariant over items, that can be chosen by the practitioner. (If $D = 1.7$ it makes essentially no difference whether the model is estimated as a logistic or normal ogive model. Alternatively, D is set to 1.) The estimate of θ is known as the *scale score*. The scale score is the principal measure of performance on the exam, although other measures derived from it, such as percentile ranks, may also

³ There are many lucid expositions of IRT, including Hambleton and Swaminathan (1985), and Hambleton, Swaminathan, and Rogers (1991).

be reported.⁴ The well-known SAT scores, ranging from 200 to 800, are an example of scale scores. The item parameter is conventionally termed difficulty. The functional form of (2a) implies that it is measured on the same scale as ability.

More elaborate models introduce additional item parameters, as in the two- and three-parameter logistic models:

$$(2b) \quad P_{ij} = [1 + \exp(-\alpha_j(\theta_i - \delta_j))]^{-1}$$

$$(2c) \quad P_{ij} = c_j + (1 - c_j)[1 + \exp(-\alpha_j(\theta_i - \delta_j))]^{-1}$$

(2b) contains a second item parameter, α_j , known as the item discrimination parameter because it enters the derivative of P_{ij} with respect to θ_i (thereby determining how well the item discriminates between examinees of different ability). In both (2a) and (2b), the limit of P_{ij} as $\theta_i \rightarrow -\infty$ is 0. This is not appropriate for tests where a student who knows nothing at all can answer an item correctly by guessing. Accordingly, c_j allows for a non-zero asymptote and is conventionally termed the guessing parameter.

The plot of P_{ij} against θ_i is known as the item characteristic curve (ICC). Three item characteristic curves using the 2-parameter logistic IRT model are depicted in Figure 1. Curve II differs from curve I by an increase in the difficulty parameter, holding constant item discrimination. Curve III corresponds to an item with the same level of difficulty as II but a doubling of the discrimination parameter. All three curves have a lower asymptote of 0. Observe that all three curves are steepest where $P_{ij} = .5$ and $\theta = \delta$.⁵ At this point the slope equals α_j .

⁴ Another measure is the so-called “true score,” which is simply the expected number right = $\Sigma P_{ij}(\theta)$ on a test comprising the universe of items. This suffers from the usual defect of a number-right score: the metric depends on the difficulty of the test, though if the universe of items were selected to give appropriate weight to each concept and skill, the true score would arguably be just the metric we require.

⁵ Because these models are additive in θ and δ , ability and difficulty are expressed on the same scale.

In the IRT models in (2a) – (2c), θ and δ are underidentified. Modification of each by an additive constant obviously leaves P_{ij} unchanged. Multiplicative constants can be offset by changes to (or absorbed in) D or the discrimination parameter. Because an additive constant corresponds to a change in the origin of the scale, while a multiplicative constant represents a change in the scale's units (e.g., from meters to yards), many psychometricians have concluded that θ and δ are measured on interval scales, which are characterized by the same conventional choice of origin and unit (see Section III). However, this opinion is by no means universal, nor is it firmly held. Many psychometricians, including some who state that these are interval scales, also regard the θ scale as arbitrary. Others caution that ability scales should be accorded no more than ordinal significance. These conclusions appear to be derived from the following considerations: (1) Because ability is a latent trait, it is impossible to verify by physical means that all 1-unit increments in θ represent the same increase in ability. This (the argument goes) confers an inherent indeterminacy on the scaling of any latent trait.⁶ (2) Replacing θ with a monotonic transformation $g(\theta)$ while making offsetting changes to the function F yields a model that fits the data just as well.⁷ Thus the θ scale rests on arbitrary assumptions regarding functional form.⁸ (3) The notion that θ measures the

⁶ “When the characteristic to be measured cannot be directly observed, claims of equal-interval properties with respect to that characteristic are not testable and are therefore meaningless.” (Zwick, 1992, p. 209)

⁷ An example is Lord's (1975) transformation of the θ scale to eliminate correlation between item difficulty and item discrimination, in which θ_i was replaced by the regression of α_j on δ_j and higher powers of δ_j evaluated at $\delta_j = \theta_i$. The result was a modification of the three-parameter logistic model:

$$P^*_{ij} = F^*_{ij}(\omega) = c_j + (1-c_j)[1 + \exp(-\alpha_j(\omega^{-1}(\omega(\theta_i))-\delta_j))]^{-1}$$

where $\omega(\theta)$ is proportional to $-.27 + 1.1694\theta + .2252\theta^2 + .0286\theta^3$. As a nonlinear transformation of the original ability scale, the ω -scale differs from the θ -scale by more than a change of origin and unit. Lord saw this as no drawback: “There seems to be no firm basis for preferring the θ scale to the ω scale for measuring ability.” (Lord, 1975, p. 205)

⁸ “A long-standing source of dissatisfaction with number-right and percent-correct scores is that their distribution depends on item reliabilities and difficulties. Radically different distributions of true scores (expected percents correct) can be obtained for the same sample of examinees when

amount of something misconceives the entire enterprise. There is no single trait (call it ability, achievement or what-have-you) to be quantified by this or any other means.⁹ As such, the question of scale type is essentially meaningless: there are just various models for reducing the dimensionality of the data, some more convenient than others.¹⁰ I will refer to this view of IRT scales in the ensuing discussion as the *operational perspective*.

To summarize, there are some psychometricians who consider θ to be interval-scaled, others who think it is ordinal, still others who regard the choice of scale as arbitrary, even if it is an interval scale, and finally some who are unsure what it is. Clearly it is disconcerting to find this divergence of views on a question of fundamental importance to value-added assessment. Is the IRT ability trait measured on an interval scale or not? Indeed, how does one tell?

they take different tests. The obvious restriction of such scores to their ordinal properties casts doubt upon their use for problems that require interval scale measurement, such as comparing individuals' gains... Item response theory appeared to offer a general solution, since the same underlying θ scale could explain the different true-score distributions corresponding to any subset of items from a domain. But this line of reasoning runs from model to data, not from data to model as must be done in practice. Suppose that a given dataset can be explained in terms of a unidimensional IRT model with response curves of the form $F_j(\theta)$. Corresponding to any continuous, strictly increasing function h there is an alternative model with curves $F_j^*(\theta) = F_j(h(\theta))$ that fits the data in precisely the same manner (Lord, 1975). That a particular IRT model fits a dataset, therefore, is not sufficient grounds to claim scale properties stronger than ordinal." (Mislevy, 1987, p. 248)

⁹ The claim that a particular unidimensional scaling method is right must be based on the assumption that achievement is unidimensional, that it can be linearly ordered, and that students can be located in this linear ordering independently of performance on a particular achievement test. However, no one has succeeded in identifying or defining a linearly ordered psychological trait in educational achievement, and no one has demonstrated that a particular measurement scale is linearly related to such a trait. A serious obstacle to the establishment of truly (externally verifiable) equal-interval achievement scales is the fact that achievement is multidimensional and qualitatively changing. The nature of what is being learned is constantly being modified. Use of an objective-based approach to achievement highlights the difficulties in hypothesizing and verifying a continuous achievement trait. The student is learning the names of letters of the alphabet one day, associating sounds to those letters another day, and attaching meaning to groups of letters a third day. How can such qualitative changes be hypothesized to fall so many units apart on one particular trait?" (Yen, 1986, p. 311-12)

¹⁰ "It is important for educators and test developers to acknowledge that until the achievement traits are much more adequately defined, it is not possible to develop measurement scales that are linearly related to such traits. In fact, it appears impossible to provide such trait definition. Test users are therefore left to use other criteria to choose the best scale for a particular application; choosing *the* right scale is not an option." (Yen, 1989, p. 314)

III. IRT Models and Scale Type

In the natural sciences, measurement is the assignment of numbers to phenomena in such a way that relations among the numbers represent empirically-given relations among the phenomena. Technically, there is a homomorphism between the empirical relations among objects in the world and the numerical relations on the scale. One such relation is order: if objects can be ordered with respect to the amount of some attribute, that order needs to be reflected in increasing (or decreasing) scale values. If objects A and B can be combined (or concatenated) in such a way that in combination they possess the same amount of some attribute (length, mass) as object C, a scale for that attribute needs to reflect the results of that operation. Thus, in an additive representation, the scaled value of the attribute in A plus the scaled value of the attribute in B equals the scaled value assigned to C.

The importance of convention relative to empirical phenomena turns out to be the key to scale type. At one end of the spectrum are scales in which there is no role for convention, at the other scales that are entirely conventional. Scales of the first type are called absolute. An example is counting: one is not free to change units if the information to be conveyed by the scale is the number of discrete objects under observation. At the other end of the spectrum are nominal scales, where the number assigned to an object is no more than a label and the information conveyed could just as well be represented by a non-numerical symbol. Scales for most physical quantities, such as length and mass, have a degree of freedom for the conventional choice of a unit. Such scales are known as “ratio scales” because the ratio of two lengths or two weights is invariant to the choice of unit: ratios are convention-free. If there is no natural zero, so

that the origin of a scale is also determined conventionally, ratios are no longer convention-free magnitudes. However, ratios of intervals are invariant under change of unit and change of origin. Such scales are therefore known as interval scales: they are characterized by two degrees of freedom. Between interval scales and nominal scales lie ordinal scales. Any increasing function of an ordinal scale conveys the same information.¹¹

With respect to psychological variables, there is less agreement about the nature of measurement. There appear to be three prevalent views (Hand, 1996). On one view (sometimes called “classical measurement”), it is simply assumed that the psychological variable of interest exists and that there is a ratio (or at least interval) scale on which it is measurable. The task of measurement is to discover those values. This appears to be the view of some psychometricians with respect to IRT scales. On a second view, a psychological variable exists only by virtue of its presence in some model. The model effectively defines the variable and, when the model is fit to data, provides a means of determining the scaled value of the variable. This has been called “operational measurement” and is compatible with the operational perspective on IRT scales described above, in which IRT models are devices for reducing the dimensionality of data. Two different models may both contain the term “ability.” There is no basis for deciding between them on grounds that one better represents “true ability”—each is just more or less useful for the purposes to which they are put. Finally, there is a third view that holds that measurement of psychological variables is essentially the same as measurement in

¹¹ Scale type is therefore closely related to the notion of an *admissible* scale transformation (Stevens, 1946). An admissible transformation preserves the empirical information in the original scale by altering only those elements that are purely conventional. For ratio scales, these are the similarity transformations, corresponding to a choice of unit (for example, substituting centimeters for meters). In the case of interval

the physical sciences—the view sketched at the start of this section, known as representational measurement. Psychological attributes are postulated to explain empirically-given relationships (such as the pattern of examinees' responses to the items on a test). It is the structure of those relationships that determines the properties of scales that measure these attributes. Some relationships are so lacking in structure that the attributes may not be scalable at all—the most we can do is to name them. In other cases it may be possible to say that there is a larger quantity of some attribute in one person than another, supporting ordinal scaling. In still other cases there may be sufficient structure to permit interval-scaled comparisons: the difference in the amount of the attribute between A and B equals the difference between C and D.

We can aspire to resolve disputes about scale type only if the third of these views is correct. Classical measurement simply assumes an answer, whereas the question of scale type is either meaningless or of no importance in operational measurement. Virtually all discussions of scale type nod in the direction of representational measurement by invoking Stevens' classification of admissible transformations. To the extent that a case can be made that IRT scales are interval scales, it has to be made in terms of representational measurement theory.

The argument that θ and δ are interval-scaled is found in the analysis of conjoint additive structures. We begin by assuming that the P_{ij} are given—or, more precisely, that we are given the equivalence classes comprising examinees and items with the same P_{ij} . Let $A_1 = \{a, b, \dots\}$ denote the set of examinees and $A_2 = \{p, q, \dots\}$ the set of test items, and let \succeq represent the order induced on $A_1 \times A_2$ by P_{ij} . That is, $(a, p) \succeq (b, q)$ if P_{ap}

scales, the admissible transformations consist of affine transformations, $g = \alpha \cdot f(a) + \beta$.

$\geq P_{bq}$. The sets A_1 , A_2 , and the relation \succeq are known as an *empirical relational system*.

If several exacting conditions are met, this system is called a “conjoint additive structure” and the following can be proved: (1) there are functions φ_1 and φ_2 mapping the elements of these sets into the real numbers (i.e., examinees and items can be *scaled*); (2) the relation ordering examinee-item pairs can be represented by an additive function of their scaled values; that is, $(a,p) \succeq (b,q) \Leftrightarrow \varphi_1(a) + \varphi_2(p) \geq \varphi_1(b) + \varphi_2(q)$; (3) only affine transformations of φ_1 and φ_2 preserve this representation. These transformations correspond to a change of origin and a change of units; hence, φ_1 and φ_2 are interval scales (Krantz et al., 1971).

There are two critical steps in the proof. First, from the relation ordering examinee-item pairs we must be able to derive relations ordering the elements of each set separately. For this we require a monotonicity condition (also known as independence): if $P_{aq} > P_{bq}$ for some item q , then $P_{ar} > P_{br}$ for all r . An analogous condition holds for items: if item q is harder for one examinee to answer, it is harder for all examinees. Monotonicity establishes orderings of A_1 and A_2 separately. Without it, not even ordinal scales could be established for examinees and items.

To obtain interval scales we need further conditions, as illustrated in Figure 2, which depicts three “indifference curves” (literally: isoprobability curves) over examinees and items: i.e., three equivalence classes determined by the P_{ij} . Examinees and items are arrayed along their respective axes according to the induced order on each set, but no significance should be attached to the distance between a pair of examinees or a pair of items: the axes are unscaled. The particular conjoint structure depicted in Figure 2 exhibits a property called equal spacing: as we move over one examinee and

down one item, we remain on the same indifference curve. Equal spacing is probably the simplest of all conjoint additive structures but it is not a necessary condition.

To derive an equal-interval metric, we establish that the distance between one pair of examinees is equal to the distance between another pair by relating both to a common interval on the item axis. As shown in Figure 3, the interval between examinees a and b can be said to equal the interval between examinees b and c in that it takes the same increase on the item scale (from q to r) to offset both. Thus the item interval qr functions as a common benchmark for defining a sequence of equal-unit intervals on the examinee axis. The scaled value of any individual i is then obtained by counting the number of such intervals in a sequence from some arbitrarily chosen origin to person i: ab, bc, cd ... hi. This scaling of the heretofore unscaled axes renders the equiprobability contours linear (and, of course, parallel).

In order that the conclusion $ab = bc = cd = \dots$ not be contradicted by other relations in the data, it must be the case that the same conclusion follows no matter which interval on the item axis is selected as the benchmark. A like condition must hold for intervals on the examinee axis to serve as a metric for items. In addition to these consistency conditions, other technical conditions must be met when equal spacing does not obtain.

To summarize, conjoint measurement scales a sequence of intervals in one factor by using differences of the complementary factor as a metric. By moving across indifference curves, these sequences can be concatenated to measure the difference between any two examinees (or items) with respect to the latent factor they embody. Because the measurement of intervals reduces to counting steps in a sequence, there is an

essential additivity in this empirical structure, on the basis of which we obtain additive representations of examinee and item traits.

Conjoint additivity uses only the P_{ij} equivalence classes and not the values of P_{ij} to determine the φ_1 and φ_2 scales. The final step from conjoint additivity to IRT requires a positive increasing transformation F from the positive reals to the interval $[0,1]$, such that $F(\varphi_1(a) + \varphi_2(p)) = P_{ap}$. Because any increasing transformation preserves the representation of $A_1 \times A_2$ by φ_1 and φ_2 , it is clear that a function F satisfying this condition can be found.

Of the three IRT models, only the one-parameter model is consistent with this simple conjoint additive structure. The two- and three-parameter IRT models, in which item discrimination varies, violate the monotonicity assumption. This is easily seen with the aid of Figure 1. The ICC labeled II represents an item with discrimination parameter = 1, while the curve labeled III has a discrimination of 2. The ICC cross where θ equals the common value of the difficulty parameter. An individual whose θ lies to the left of this intersection finds item II more difficult than item III; an individual whose θ is to the right of the intersection ranks the items the other way around. Because different individuals produce inconsistent rankings of items, the ranking of examinee-item pairs on the basis of P_{ij} does not yield orderings of examinees and items and the derivation of the φ scales breaks down. Indeed, it is alleged that because only the Rasch model (among logistic IRT's) produces a consistent ordering of items for all examinees (ICCs in the Rasch model never cross), only the Rasch model's θ can be considered a measure of ability (Wright, 1999).

This goes too far. A straightforward modification of conjoint additivity accommodates structures with a third factor that enters multiplicatively, as does the IRT discrimination parameter. The relevant theorem appears in Krantz et al. (1971, p. 348). The empirical relations between examinees and items are termed a polynomial conjoint structure. The extra conditions on examinees and items ensure that we can obtain separate representations by discrimination classes in the manner just described. When these conditions are met, we find that there are functions φ_1 , φ_2 , and φ_3 such that $(a,p) \succeq (b,q) \Leftrightarrow \varphi_3(p)[\varphi_1(a) + \varphi_2(p)] \geq \varphi_3(q)[\varphi_1(b) + \varphi_2(q)]$. φ_1 and φ_2 are unique to linear transformations and φ_3 is unique to a similarity transformation. Obviously this representation has the structure of the two- and three-parameter IRT models.¹²

To summarize: if the empirical relational system is one of conjoint additivity and the function F correctly specifies the relationship between $\theta_i - \delta_i$ and the P_{ij} , θ and δ are interval-scaled variables. If the empirical relational system is a polynomial conjoint structure and F is again correct (e.g., the two-parameter logistic IRT model, or possibly the three-parameter model, if extra conditions entailing a lower asymptote on P_{ij} are met), θ and δ are again interval-scaled. It is important to stress that the metric properties of the θ scale are established by the features of the empirical relational system, not by the use of a particular IRT model. Applying the Rasch (or any other) model willy-nilly to achievement test data does not of itself confer any particular properties on the scale score

¹² This representation does not include anything that corresponds to the guessing parameter in the three-parameter IRT. While I have not seen an analysis of such a conjoint measurement structure, the additional parameter, c_j , affects only the function mapping $(\varphi_3(q)[\varphi_1(b) + \varphi_2(q)])$ to P_{ij} and not the relations among φ_1 , φ_2 , and φ_3 . Thus it is clear that polynomial conjoint measurement could be extended to the three-parameter IRT model.

metric.¹³

IV. Scale Compression

It is highly unlikely that real test data meet the exacting definition of a conjoint structure. At best, they come close, though it is not easy to say how close. The theory places restrictions on the equivalence classes of $A_1 \times A_2$ (examinees crossed with items), but these classes are not given to us. Instead, the data consist of answers to test items, often binary indicators of whether the answer was correct or not. From these data the membership of the equivalence classes must be inferred before one can ascertain whether these classes can be represented by a set of linear, parallel isoprobability curves. Because the theory stipulates restrictions that hold for every examinee, while the amount of data per examinee is small, there is little power to reject these hypotheses despite anomalies in the data. Many restrictions might be accepted that would be deemed invalid if the actual P_{ij} were revealed.

IRT assumes that conditional on θ and δ , the P_{ij} are independent across items and examinees. Although correlated response probabilities attest to the existence of additional latent traits that affect performance, unidimensional models are fit to the data anyway.

Sometimes it is clear even without statistical tests that the model does not fit the data. Consider multiple choice exams, where the lower asymptote on the probability of a correct response is not zero. The lower asymptote is more important for low ability than

¹³ That this is not always appreciated by researchers is seen in the following statement by the Consortium on Chicago School Research. Having rescaled the Iowa Test of Basic Skills using a Rasch IRT model, they claim to have produced a metric with interval scale properties: “A third major advantage of the Rasch equating is that, in theory, it produces a ‘linear test score metric.’ This is an important prerequisite in studies of quantitative change. This allows us to compare directly the gains of individual students or schools that start at different places on the test score metric.” (Bryk et al., 1998).

high ability examinees. A conjoint structure for these data must take into account the lower asymptote as another factor determining P_{ij} or the scaled values of examinee ability will be wrong. Indeed, if responses to a multiple-choice exam are tested to see whether the definition of conjoint additivity is met, data that appear to meet this definition will no longer do so once guessing is factored in.

Even when the data do fit the model, the extent to which they have been selected for just this reason should be kept in mind. Indeed, it is not clear that the word data is the right one in this context, as the tests are designed by test makers who decide first on a scaling model and then strive to ensure that their test items meet the assumptions of that model.¹⁴ Even if they are wholly successful in this endeavor, the “data” represent a selected, even massaged, slice of reality and not a world of brute facts.

Thus, notwithstanding the support given by representational measurement theory to IRT methods, IRT can fail to produce satisfactory interval scales. Verifying that the conditions of the theory are met is difficult even for test makers, let alone statisticians and behavioral scientists who use test scores for value added assessments. In these circumstances, it seems reasonable to examine the scales that have been produced.

The strongest evidence that something is amiss appears in the developmental scales used to measure student growth across grades. Students at different grade levels are typically given different exams, but by including a sufficient number of overlapping items on forms at adjacent grade levels, performance on one test can be linked to

¹⁴ For example, in the Rasch model items have a common discrimination parameter (their item characteristic curves do not cross). Test makers using the Rasch model plot ICC's and discard items that violate this assumption. IRT also assumes that conditional on θ , probabilities of a correct response are independent over examinees. Violations of this assumption lead items to be discarded on grounds of potential bias. This is probably a good idea, but it further weakens the sense in which we are dealing here with the fit between a model and “data.”

performance on another test, facilitating the creation of a single scale of ability spanning multiple grade levels.

Table 1 below, reproduced from Yen (1986), displays scores for the norming sample for the Comprehensive Test of Basic Skills (CTBS), Form Q, developed by CTB/McGraw-Hill and scaled using IRT methods for the first time in 1981.¹⁵ In both subjects, mean growth between grades drops dramatically, and almost monotonically, between the lowest elementary grades and secondary grades. In addition, the standard deviation of scale scores declines as the mean score rises.

CTB/McGraw-Hill has since superseded the CTBS with the Terra Nova series. The decline in between-grade gains seen in the CTBS norming sample is still evident, though less pronounced, in Terra Nova results for Mississippi and New Mexico, as shown in Table 2.¹⁶ Mean gains in all subjects tend to decrease with grade level, though there is a break in the pattern between grades 7 and 8.

Between-grade gains can be affected by the differences in the content of tests and by linking error. It is particularly instructive, therefore, to see the patterns that emerge when there are no test forms specific to a grade level and no linking error. Northwest Evaluation Association uses computer-adaptive testing in which items are drawn from a single, large item bank. Results for all examinees in mathematics in the fall of 2005 are presented in Table 3. As in the previous tables, we again find that between-grade

¹⁵ Prior to 1981 the CTBS was scaled using an older method known as Thurstone scaling. Although it has been claimed that Thurstone scaling produces scores on an interval scale, there is no basis for this claim comparable to the proofs provided for conjoint measurement structures in representational measurement theory: scale properties are not based on the structure of empirical relations holding among test items and examinees, but on ad hoc assumptions about the distribution of ability in the population tested.

¹⁶ Data are from www.SchoolData.org maintained by the American Institutes of Research. New Mexico and Mississippi were two of a handful of states in these data that reported mean scale scores for vertically-linked tests produced by test-makers known to employ IRT scaling methods.

differences decline with grade level. Within-grade variance in scores is stable (reading) or increases moderately (mathematics).

Not all achievement tests scaled using IRT methods exhibit such patterns as consistently as these. Table 4 presents results for California from 2000, when the Stanford-9 was used, and 2003, after that exam had been replaced by the California Achievement Test. There is still some tendency for high growth in the early grades to be followed by lower gains. However, the declines are far from monotonic. Instead, the most striking thing about the table is probably how erratic these changes are.

I suspect that many, if not most, educators would find it incredible that the variance of reading ability is no greater among high school students than second-graders. This result may owe something to the fact that in IRT models, the increase in ability required to raise the probability of a correct response by any fixed amount is independent of the difficulty and complexity of the question. If larger increments of ability are in fact required as questions become more difficult, IRT scaling will compress the high end of the scale, diminishing mean gains between the upper grades and reducing the variance in achievement.

Figure 4 contains a sample of mathematics test items posted on the Northwest Evaluation Association web site. The items are drawn from a larger chart providing examples of test questions at eleven different levels of difficulty. The two selections represent items scored at the 171-180 level and the 241-250 level, respectively.¹⁷ The reader is invited to consider the following question: if student A is given the items in the first set, and student B the items in the second set, and if initially each student is able to

¹⁷ According to the data in Table 3, the average second grader tested in 2005 would have answered slightly more than half the questions in the first set correctly, while the average ninth grader would have responded

answer only 2 of 7 items correctly, which student will have to learn more mathematics in order to answer all 7 items correctly? Student A has basics of addition and subtraction to learn, as well, perhaps, as how to read simple charts and diagrams. None of the required calculations are taxing: all could be done by counting on one's fingers. By contrast, student B must make up deficits in several of the following areas: decimal notation, fractions, factoring of polynomials, solving algebraic equations in one unknown, solid geometry, reading box plots, calculating percentages. The calculations required are more demanding. However, the correct answer, according to IRT, is that both require the same increase in mathematics ability.^{18,19}

One might conclude from this evidence that psychometricians should not attempt to construct over-arching developmental scales for mathematics and reading ability that span so many age levels. But other concerns are also raised. Between-grade gains begin to decline at the lowest grade levels in Tables 1-3: gains between third and fourth grades are markedly lower than gains between second and third. Between-grade comparisons may matter little for value-added assessment, if most instructors, particularly in the elementary grades, teach only one grade level. However, there are also implications for within-grade assessments. If third graders are not in fact learning less than second

correctly to slightly less than half the questions in the second set.

¹⁸ Strictly speaking, this is the correct answer only if items have equal discrimination parameters, that is, the one-parameter IRT model fits the data. In fact, NWEA does use the one-parameter model, and has gone to considerable lengths to verify that the items meet the assumptions of that model.

¹² When I put this question to faculty and graduate students of my department in the School of Education at Vanderbilt, 13 of the 108 respondents chose A, 47 chose B, 15 said the amounts were equal, and 33 said the answer was indeterminate. Obviously this was not a scientifically-conducted survey, nor is it clear just what respondents meant by their answers. Conversations with some revealed that they converted the phrase "more mathematics" into something more readily quantified, such as the amount of time a student would need to acquire these skills. Persons who said the answer was indeterminate may have meant it was indeterminate in principle ("more mathematics" is meaningless) or simply that that answer couldn't be determined from the information given. Nonetheless, it is striking how few gave the psychometrically correct response.

graders—if instead IRT methods have compressed the true scale—then the true gains of higher-achieving students within these grades are understated.²⁰

V. Options for Practice

What options are available to the practitioner who wants to conduct value-added assessment but is unwilling to accept at face value claims that the IRT ability trait is measured on an interval scale? Broadly speaking, there are three available courses of action.

- a. Use the θ scale anyway.
- b. Choose another scale.
- c. Adopt analytical methods suited to ordinal data.

I consider these in turn.

- a. Using the θ scale.

²⁰ The phenomenon to which I refer as scale compression has been treated in the psychometric literature under the heading scale shrinkage, with more of an emphasis on decreasing variance of scores at higher level of ability than on declining rates of growth. Yen (1985) shows that shrinkage can arise when a unidimensional model is fit to multidimensional data, i.e., items that test more than one type of ability. In Yen's two-dimensional model, shrinkage occurs when the correlation between difficulties of type 1 and 2 at the item level is lower than the correlation between corresponding abilities 1 and 2 at the individual level. However, hers is a compensatory model in which higher levels of ability 1 can offset lower levels of ability 2. If items are truly multidimensional, a conjunctive model of scoring would presumably be more realistic, in which a deficit in any of the required skills results in an incorrect answer. Thus, while suggestive, her analysis does not settle just what sort of mismatch between model and data explains shrinkage.

Other explanations of shrinkage have focused on measurement error, with the decline in measured variance allegedly a function of increased test reliability (Camilli, 1988). This hypothesis might account for declining variance within a grade level between fall and spring testing, if there is close alignment between the test and the curriculum taught over the year, but has no obvious applicability to the phenomenon of declining variance across grade levels. Indeed, test reliabilities have not been found to increase with higher grade levels, as required by this hypothesis (Yen and Burket, 1997).

In summary, it seems fair to say that while the question of declining test variance has received some attention from psychometricians, the underlying reasons are not perfectly understood. I call attention to it in this paper not to suggest otherwise, but because such a feature of test scales strains our credence in claims that ability is measured on an interval scale. It is worth noting that the psychometricians who have examined shrinkage have themselves expressed strong doubts about such claims: “[I]t is generally acknowledged that IRT methods do not yield interval scales” (Camilli, Yamamoto, and Wang, 1993, p. 387); “[F]or real tests, the ‘true scale is unknowable, and an objective criterion does not exist for evaluating how scales should behave (Yen and Burket, 1997, p. 293).

Even if the IRT ability scale is ordinal, one might continue to use it if reasonable transformations of θ all yield essentially the same estimates of value added. This view finds some support in the claim that statistics calculated from ordinal variables are generally robust to all but the most grotesque transformations of the original scale (Cliff, 1996).²¹ The practice of many social scientists who are aware that scales may not be of the interval type but proceed with value-added assessment anyway suggests that this view may be widespread.

I have tested this conjecture using data from a sample of districts in a Southern state.²² Mathematics tests were administered in grades 2-8 during the 2005-06 school year. The tests were scaled using the one-parameter IRT model. Results were similar to those we have seen in Tables 1-3, with near-monotonic declines in between-grade gains. A transformation $\psi = g(\theta)$ was sought that would equalize between-grade gains for the median student.²³ It turned out that this could be closely approximated by a quadratic function increasing over the range of observed scores. Figure 5 depicts the relationship between the transformed and original scaled values for the median examinee. Obviously there is a perfect mathematical equivalence between the IRT model $F(\theta, \delta)$ and the alternative $F(g^{-1}(\psi), \delta)$ that uses the transformed variable; on grounds of fit, there is nothing to recommend one model over the other. Whatever objections one might make

²¹ Cliff attributes the remark to Abelson and Tukey (1959), but it does not appear in that paper.

²² Anonymity has been promised to both the state and test maker. Data are available for only a portion of the state. In addition, because between-grade growth is central to the analysis, only those districts that tested at least 90 percent of students in each grade are included. The final sample comprised 98,760 students in grades 2 through 8 during the 2005-06 school year.

²³ To accomplish this, the original median scores were replaced with a new series in which the between-grade gain was set to the overall sample median gain across all grades. This left unchanged the second grade values but altered the values in subsequent grades, as shown in Figure 5. A quadratic function of θ was then fit to the new series. The fit, as shown, was exceedingly close, though on the resulting scale, median gains can vary by ± 0.5 points.

to the ψ -scale, it would be a considerable stretch to regard it as a grotesque transformation of the data.

Table 5 shows the consequences of this transformation for measured between-grade growth. At each grade level, I have calculated the change in θ (alternatively, ψ) required of a student who remains at the 10th, 25th, 50th, 75th, and 90th percentiles of the achievement distribution. Comparisons of absolute changes in θ and ψ are affected by the magnitude of the median student's assumed growth. However, comparisons of relative changes are invariant to this assumption. Accordingly, column 1 presents the growth in θ , then ψ , of a student who remains at the 25th percentile, relative to the growth of a student remaining at the 10th percentile. Comparable ratios for the 50th, 75th, and 95th percentiles appear in the remaining columns.

In the original scale, differences in growth at various points of the distribution are not pronounced. The ratios are greatest in the fifth and sixth grades, but even here, there is not much difference between a student at the median and one at the 90th percentile. By contrast, the ratios are much greater using the transformed scale. They continue to widen between the 50th and 90th percentiles. The direction of the difference is what we would have expected on the basis of the preceding discussion: there is less compression at the high end of the ψ -scale. However, the magnitude of the difference may be surprising. The impact on value-added assessment depends on how students are distributed among schools and teachers. Clearly, transformation of the scale can make a great difference when that distribution is not uniform.

b. Changing the scale.

Clearly nothing is gained by substituting percentiles, normal curve equivalents, or standardized scores for θ . However, two approaches merit discussion.

The first, which I will refer to as binning, assigns every examinee to a group defined by prior achievement (e.g., deciles of the distribution of prior scores). Gain scores are normalized by mean gains within bins, either through dividing by the mean gain or standardizing with respect to the mean and standard deviation within each bin. The normalized gain scores are then used as raw data for further analysis, which could include value-added assessment of schools and teachers. Examples of this approach are Springer (2008) and Hanushek et al. (2005). In the latter study, binning is explicitly motivated by concern that the metric in which test results are reported does not represent true gains uniformly at all points of the achievement distribution.²⁴

Binning does not solve the problem of scale dependence. Binning merely declares that a normalized gain in one bin is to count the same as one that takes the same normalized value in another bin. The declaration does not ensure that these two gains are equal when measured on the true scale (if such a thing exists). Rather than solving the problem, binning is simply the substitution of a particular transformation for the θ scale.

The notion that binning represents a solution may derive from the fact that normalizing within bins removes much of the effect of any prior transformation of scale (for example, the substitution of the ψ scale for the θ scale). Differences between bins have no effect on value-added measures. Only differences within bins matter (though

²⁴ A practice similar to binning has been used in the Educational Value-Added Assessment System (EVAAS) of the SAS Institute, wherein gain scores are divided by the gain required to keep an examinee at the same percentile of the post-test distribution that he occupied in the pre-test distribution (Ballou, 2005). Thus, if the distribution in terms of the original scale exhibits increasing variance at higher grade levels, the transformation pulls up gains of examinees whose pretest scores were below the mean and reduces gains of examinees whose pretest scores were above the mean. It is unclear why this should be regarded as superior to using the original scale.

these are still affected by the transformation whenever the slope of $g(\theta) = \psi$ at a bin mean differs from 1). Hence, the results of the binned analysis are less sensitive to prior transformations of the original scale. This may have led some to believe that scale no longer matters as much. This is not so, given that the normalized-within-bin scores are themselves just another transformation of θ . Moreover, if invariance of this kind were the sole desideratum, percentiles could be used in place of θ . Percentiles are invariant to any increasing transformation of θ , but that does not make percentiles an interval scale of achievement.

Suppose we decide that the problem of finding interval scales for academic achievement is intractable. We still might be able to conduct value-added assessment if achievement—by whatever metric—could be mapped into other variables whose measurement poses no such difficulties. This mapping could go backward to inputs or forward to outcomes. In the former, academic achievement would be related to measurable inputs required to produce that achievement. Thus, instead of worrying whether one student's 5 point gain was really the equal of another student's 5 point gain, we would concern ourselves with measuring the educational inputs required to produce either of these gains. If those inputs (e.g., teacher time) turn out to be equal, then for all purposes that matter, the two gains are equivalent. If those inputs turn out not to be equal, the teacher or school that has produced the gain requiring the greater inputs has contributed more and should be so recognized by value-added analysis.

Forward mapping treats test scores as an intermediate output. Value-added assessment would proceed by tying scores to long-term consequences.

“An important methodological issue...is the problem of choosing the correct metric with which to measure academic growth. Because the metric issue is so perplexing, almost all

researchers simply use the particular test at their disposal, without questioning how the test's metric affects the results... The only solution I see to the problem of determining whether gains from different points on a scale are equivalent is to associate a particular test with an outcome we want to predict (say, educational attainment or earnings), estimate the functional form of this relationship, and then use this functional form to assess the magnitude of gains. For example, if test scores are linearly related to years of schooling, then gains of 50 points can be considered equal, regardless of the starting point. If the log of scores is linearly related to years of schooling, however, then a gain of 50 points from a lower initial score is worth more than a gain of 50 points from a higher initial score. This 'solution' is, of course, very unsatisfactory, because the functional form of the relationship between test scores and outcomes undoubtedly varies across outcomes." (Phillips, 2000, p. 127)

We are very far from being able to carry out either of these programs. Only some educational inputs are easily quantified. For such inputs as the clarity of a teacher's explanations or the capacity to inspire students, the challenges to quantification are at least as great as for academic achievement. Indeed, the low predictive power of those inputs that are easily quantified is largely responsible for the current interest in value-added assessment.

The practical difficulty mentioned in the last sentence of the quoted passage is not the only problem facing the forward mapping of test scores to long-term educational outcomes. Given the variation in test results in Tables 1-4, it would almost certainly be the case that the functional form of the relationship between test scores and outcomes would vary across tests as well as outcomes. The introduction of each new test would require additional analysis to determine how scores on its metric were related to long-term objectives like educational attainment and earnings. In many cases, the data for such analysis would not be available for years to come, if ever. In the interim we would have to make do with very imperfect efforts to equate the new tests with tests already in use (for which we would hope this mapping had already been done).

These are the technical issues. There is in addition the difficult normative question of how to value various outcomes for different students in order to assign a unique social value to each θ_i . It is not obvious how we would come by these weights. Even if future earnings were the only outcome that mattered, we would require the relative values of a marginal dollar of future income for all examinees (whose future incomes are, of course, unknown at the time the assessment is done). Attaching a price to the non-pecuniary benefits of an education is still more difficult. Education is a transformative enterprise. The ex ante value placed on acquiring an appreciation of great literature is doubtless very different from the ex post value. In these circumstances there is likely to be a great discrepancy between the compensating and equivalent variations associated with a given educational investment. Accordingly, the weights in our index would have to reflect the values of “society” rather than the still-unformed persons to be educated.

It is not clear that we should subject decisions about education to this kind of utilitarian calculus, as it fails to respect the autonomy of individuals. There is a strong tradition in our polity of regarding educational opportunity as a right. Individuals have a claim on educational resources not because distributing resources in this manner maximizes a social welfare function, but because they are entitled to the chance to realize their potential as individuals. If we take this seriously, then the notion that teachers and schools are to be evaluated by converting test scores into the outputs that matter and weighting these outputs according to their social value is wrong-headed. The point of education is to provide students with skills and knowledge that as autonomous persons they can make of what they will. Teachers should therefore be judged on how

successfully they equip their students with these tools—regardless of anyone’s views of the merits of the final purposes, within wide limits, for which students use them.

c. Analyzing test scores as ordinal data

The final option for researchers uncertain of the metric properties of ability scales is to treat such scales as ordinal, thus foregoing any analysis based on the distance between two scores. On the assumption that θ scales contain valid ordinal information about examinees, statistics based on the direction of this distance remain meaningful. There are a variety of closely-related statistics of this kind (known generally as measures of concordance/discordance) employed in the analysis of ordinal data. In this discussion I will not attempt to identify a particular approach as best. Rather, my objective is to demonstrate the feasibility of such methods for value-added assessment.

Suppose we want to compare the achievement of teacher A’s class to the achievement of other students at the same grade level in the school system. If A has n students, and teachers elsewhere in the system have m students, there are nm possible pairwise comparisons of achievement. As only ordinal statements are meaningful, each pair is examined to determine which student has the higher score. If it is A’s student, we count this as one in A’s favor (+1); if it is the student from elsewhere in the system, we count this as one against A (-1). Ties count as zeros. The sum of these counts, divided by the number of pairs, is known as Somers’ d statistic. Somers’ d can be considered an estimate of the difference in two probabilities: the probability that a randomly selected student from A’s class outperforms a randomly selected student from elsewhere in the system, less the probability that A’s student scores below the outsider.

This procedure suffers from the obvious defect that no adjustment has been made for other influences on achievement. In conventional value-added analyses, this might be accomplished through the introduction of prior test scores as covariates in a regression model or by conditioning on prior scores in some other manner. An analogous procedure in the ordinal framework would be to divide students into groups on the basis of one or more prior test scores, compute separate values of Somers' d by group, and aggregate the resulting statistics using the share of students in each group as weights. In principle there is no reason to restrict the information used to define groups to test scores. Any student characteristic could be used to define a group. Only data limitations prevent the construction of ever-finer groups.

Results from an application of this approach are presented Table 6. Because the data used in the previous example do not contain teacher identifiers, for this application I have used a different data set provided by a single, large district that contains student-teacher links. Two sets of value-added estimates are presented for teachers of fifth grade mathematics. The first is the weighted Somers' d statistic, based on pairwise comparisons of each teacher's students to other fifth graders in the district. To control for prior achievement, students were grouped by decile of the fourth grade mathematics score. Students without fourth grade scores were dropped from the analysis. The second value-added measure is obtained from a regression model in which fifth grade scores were regressed on fourth grade scores and a dummy variable for the teacher in question. Separate regressions were run for each teacher so that each teacher was compared to a hypothetical counterpart representing the average of the rest of the district, preserving the parallel with the Somers' d statistic. The coefficients on the dummy variables represent

teachers' value added. Statistical significance was assessed using the conventional t statistics in the case of the regression analysis and jackknifed standard errors in the case of the ordinal analysis.

How much difference does it make to teacher to be evaluated by the one method rather than the other? The hypothesis that teachers are ranked the same by both methods is rejected by the Wilcoxon signed rank test ($p = .0078$). The proportion of statistically significant estimates is higher using the ordinal measure (which is less sensitive to noisiness in test scores): 86 of the 237 teachers are significant by this measure, compared to 64 by the other. The maximum discrepancy in ranks is 229 positions (out of 237 teachers in all). In ten percent of cases, the discrepancy in ranks is 45 positions or more.

Depending on the uses to which value-added assessment is put, the question of greatest importance to teachers may be whether they fall at one end of the distribution or the other. There are 59 teachers in each quartile of the distribution. The two measures agree in 47 cases on the teachers in the top quartile, and in 48 cases on teachers in the bottom quartile. Thus, if falling in the top quartile qualifies a teacher for a reward, 24 teachers (more than a third of the number of awardees) will qualify or not depending on which measure is used. A comparable figure applies to teachers placing in the bottom quartile, if that event is used to determine sanctions. In a small number of cases (3), the effect of using one measure rather than the other is great enough to move a teacher from the top quartile to the bottom quartile.

In principle it is possible to condition on multiple variables (additional prior test scores, student demographic characteristics or SES) by defining groups as functions of several covariates. In practice this is apt to exceed the capacity of the data. Consider, for

example, a data set containing prior test scores in two subjects, plus indicators of race and participation in the free and reduced-price lunch program. If groups are defined by deciles of the two test scores plus two binary indicators, the total number of groups is 400 ($10 \times 10 \times 2 \times 2$). In a district of moderate size, there could be many cells with only one observation and therefore no matching pair.

A two-stage method circumvents this difficulty. In the first stage, multiple covariates are used to predict $\text{Prob}(Y_i > Y_j) - \text{Prob}(Y_j < Y_i)$ for all pairs of students i and j , where Y is the dependent variable of ultimate interest (e.g., end-of-year test scores in the current year). The prediction is of the form $\hat{\pi}_i = \sum w_k (X_{ki} - X_{kj})$ or $\sum w_k \text{sgn}(X_{ki} - X_{kj})$, depending on whether the X_k are themselves interval-scaled or ordinal. The w_k are weights that reflect how informative the different covariates are about $\text{sgn}(Y_i - Y_j)$. (For details, see Cliff, 1996.)

In the second stage, the n students of one teacher are compared to the m students elsewhere in the same system, using $\hat{\pi}_i$ (or grouped values of $\hat{\pi}_i$) as the covariate. $\hat{\pi}_i$ is therefore analogous to a propensity score: it is a summary measure of the effects of the stage-one covariates on the probability that a student outranks other students, *before controlling for teachers*. The resulting measure of a teacher's value added is based on her students' performance relative to this prediction (in the ordinal sense, of course).

Given that many achievement tests are now administered to thousands of students throughout a state, it is worth noting that all the data can be used in the first stage to form an ordinal prediction based on prior achievement, demographics, and SES, while

continuing to rely on within-district comparisons in the second stage for the final measures of value added.²⁵

VI. Conclusion

Are IRT ability traits measured on an interval scale? It seems hazardous to assume so. Whether examinees and test items constitute a conjoint structure depends on the make-up of equivalence classes defined by the P_{ij} , but those are not given. Statistical testing can reveal whether the data are strongly inconsistent with this hypothesis, but moderate departures from the conjoint structure almost certainly go undetected. Moreover, even if these conditions hold in the norming samples used by test-makers to calibrate item difficulties, this provides no assurance they will hold in the population of students to whom the test is finally administered.

End users of the data, such as practitioners of value-added assessment, typically have no access to item-level data to test these assumptions themselves. The best they can do is to check the plausibility of the resulting scales. On this count IRT ability scales often do poorly. Gain scores frequently fall from one grade to the next. While some of this may reflect adolescents' declining interest in academic achievement, the patterns set in as early as third grade and the drop between second and third grade is often the largest. Worse, IRT ability often shows a diminishing or constant variance from lower to higher grades.

What, then, is the practitioner of value-added assessment to do? It is no good hoping that the choice of scale makes little difference to estimates of value added. We have seen that reasonable transformations of the θ scale can have a substantial effect on

²⁵ Within district comparisons are preferred, given the impact of curriculum and other district policies on achievement.

relative gains across the distribution of achievement. No other scales with superior metric properties are at hand. We can, however, use methods of ordinal data analysis on the assumption that IRT scales (or any of their monotonic transformations) at least permit us to rank students.

Ordinal analysis changes the question we ask in value-added assessment. Instead of measuring mean achievement of a teacher's students vis-à-vis the students of a (hypothetical) average instructor, we ask what fraction of the former outperform the latter. In ordinal analysis, as in regression-based methods, it is possible to control for other influences in order to isolate the teacher or school's contribution. Clearly, if θ is an interval-scaled variable, ordinal methods throw away valuable information. Practitioners should ask themselves, however, whether they are so confident of the metric properties of θ scales that they are willing to attribute differences between conventional estimates of value added and estimates based on ordinal analysis to the superiority of the former.

Ordinal methods have other advantages. They are likely to be more robust to measurement error in test scores and to various model misspecifications (though the question of robustness is a complicated one; see Cliff, 1996). The question they answer may be a more sensible way to evaluate educators, given that it attaches more value to spreading gains over a wider number of students, compared to larger but more concentrated gains. However, this paper has considered ordinal methods from one standpoint only—that of finding appropriate value-added models when test scores are not expressed on an interval scale. Numerous questions have been raised about the assumptions and methods of more conventional regression-based analyses. Some of these concerns can be addressed through modifications of those models. It remains to be

seen whether the same concerns arise with respect to ordinal methods and, if so, how readily they can be accommodated within the ordinal framework.

References

- Abelson, Robert, and John Tukey. 1959. Efficient conversion of nonmetric information into metric information. Proceedings of the Social Statistics Section, American Statistical Association, 226-230.
- Ballou, Dale. 2005. "Value-Added Assessment: Lessons from Tennessee." Lissitz, Robert, ed., Value Added Models in Education: Theory and Applications. Maple Grove, MN: JAI Press.
- Bryk, Anthony et al. 1998. Academic Productivity of Chicago Public Elementary Schools. Chicago: Consortium on Chicago School Research.
- Camilli, Gregory. 1988. "Scale Shrinkage and the Estimation of Latent Distribution Parameters." Journal of Educational Statistics, 13(3), 227-241.
- Camilli, Gregory, Kentaro Yamamoto and Ming-wei Wang. 1993. "Scale Shrinkage in Vertical Equating." Applied Psychological Measurement, 17(4), 379-388.
- Cliff, Norman. 1996. Ordinal Methods for Behavioral Data Analysis. Mahwah, NJ: Lawrence Erlbaum.
- Hambleton, Ronald and H. Swaminathan. 1985. Item Response Theory: Principles and Applications. Boston: Kluwer-Nijhoff.
- Hambleton, Ronald, H. Swaminathan and Jane Rogers. 1991. Fundamentals of Item Response Theory. Newbury Park: Sage.
- Hand, D.H. 1996. "Statistics and the Theory of Measurement." Journal of the Royal Statistical Society. Series A (Statistics in Society). 159(3), 445-492.
- Hanushek, Eric A. et al. 2005. Charter School Quality and Parental Decision-Making with School Choice. Cambridge, MA: NBER Working Paper 11252.
- Krantz, David H., et al. 1971. Foundations of Measurement. Vol. 1, Additive and Polynomial Representations. NY: Academic Press.
- Lord, Frederic. 1975. "The 'Ability' Scale in Item Characteristic Curve Theory." Psychometrika, 40(2), 205-217.
- Mislevy, Robert. 1987. "Recent Developments in Item Response Theory with Implications for Teacher Certification." Review of Research in Education, 14. Washington DC: American Educational Research Association.
- Northwest Evaluation Association. 2008. Math_Single.pdf. Retrieved February 20, 2008, from www.nwea.org/assessments/ritcharts.asp.

Phillips, Meredith. 2000. "Understanding Ethnic Differences in Academic Achievement: Empirical Lessons." Grissmer and Ross, eds., Analytic Issues in the Assessment of Student Achievement. Washington DC: US Department of Education, National Center for Education Statistics.

Springer, Matthew G. 2008. Accountability Incentives: Do Schools Practice Educational Triage? Education Next. (Winter), 75-79.

Stevens, S.S. 1946. "On the Theory of Scales of Measurement." Science, 103, 677-680.

Wright, Benjamin D. 1999. "Fundamental Measurement for Psychology." Embretson and Hershberger, eds., The New Rules of Measurement. Mahwah, NJ: Lawrence Erlbaum.

Yen, Wendy. 1985. "Increasing Item Complexity: A Possible Cause of Scale Shrinkage for Unidimensional Item Response Theory." Psychometrika, 50(4), 399-410.

Yen, Wendy. 1986. "The Choice of Scale for Educational Measurement: An IRT Perspective." Journal of Educational Measurement, 23(4), 299-325.

Yen, Wendy and George R. Burket. 1997. "Comparison of Item Response Theory and Thurstone Methods of Vertical Scaling." Journal of Educational Measurement, 34(4), 293-313.

Zwick, Rebecca. 1992. "Statistical and Psychometric Issues in the Measurement of Educational Achievement Trends: Examples from the National Assessment of Educational Progress." Journal of Educational Statistics, 17(2), 205-218.

Table 1: Scale Scores for Comprehensive Test of Basic Skills, 1981 Norming Sample

Grades	Reading/Vocabulary		Mean Between- Grade Change	Mathematics Computation		Mean Between- Grade Change
	Mean Score	Std. Dev.		Mean Score	Std. Dev.	
1	488	85	----	390	158	---
2	579	78	91	576	77	186
3	622	65	43	643	44	67
4	652	60	30	676	35	33
5	678	59	26	699	24	23
6	697	59	19	713	20	14
7	711	57	14	721	23	6
8	724	54	13	728	23	7
9	741	52	17	736	17	8
10	758	52	17	739	16	3
11	768	53	10	741	18	2
12	773	55	5	741	20	0

Source: Yen (1986).

Table 2: Mean Between-Grade Differences, Terra Nova

Subject_Year	2nd to 3rd	3rd to 4th	4th to 5th	5th to 6th	6th to 7th	7th to 8th	8th to 9th
<i>Mississippi, 2001</i>							
Language Arts	30.2	20.4	17.8	6.9	9.9	10.9	----
Reading	23.8	21.5	15.5	10.5	9.2	12.4	----
Math	47.3	24.7	19.0	21.7	10.7	17.0	----
<i>New Mexico, 2003</i>							
Language _Arts_2003	----	14.8	12.6	5.1	3.7	4.5	8.3
Reading_2003	----	14.4	13.8	5.0	4.7	11.4	4.9
Math_2003	----	21.1	14.6	19.0	5.5	16.5	5.4
Science_2003	----	20.2	13.9	9.0	11.6	11.9	7.1
History/SS_2003	----	13.3	6.2	11.2	11.5	3.6	3.6

Source: Author's calculations from school-level data posted on www.SchoolData.org, maintained by American Institute of Research. Data for each school are weighted by enrollment and aggregated to the state level.

Table 3: Scale Scores, Northwest Evaluation Association, Fall, 2005

Grade	Reading			Mathematics		
	Mean	Std	Change	Mean	Std	Change
2	175.57	16.22	----	179.02	11.81	----
3	190.31	15.56	14.74	192.96	12.06	13.95
4	199.79	14.95	9.48	203.81	12.80	10.85
5	206.65	14.60	6.86	212.35	13.92	8.53
6	211.49	14.76	4.84	218.79	15.00	6.44
7	215.44	14.82	3.96	224.59	15.99	5.80
8	219.01	14.76	3.56	229.38	16.79	4.79
9	220.93	15.28	1.92	231.76	17.42	2.38

Source: Author's calculations from data provided by NWEA.

Table 4: Mean Between-Grade Differences in California

Subject	2nd to 3rd	3rd to 4th	4th to 5th	5th to 6th	6th to 7th	7th to 8th	8th to 9th	9th to 10th	10th to 11th
2000 SAT-9									
Language Arts	18.6	23.3	8.8	3.6	17.7	3.0	13.9	-3.0	15.0
Reading	29.7	27.8	16.7	2.5	22.6	12.1	2.7	2.3	12.6
Math	32.0	22.6	15.4	12.1	8.8	6.8	18.5	3.3	12.0
2003 CAT									
Language Arts	15.8	7.4	19.7	-1.6	6.7	7.2	9.5	16.1	9.7
Reading	16.2	4.1	18.0	0.0	8.5	11.1	11.9	8.4	10.5
Math	51.0	10.3	16.6	11.3	1.9	19.7	16.5	12.2	12.4

Source: Author's calculations from school-level data posted on www.SchoolData.org, maintained by American Institute of Research. Data for each school are weighted by enrollment and aggregated to the state level.

Table 5: Effect of Scale Transformations on Between-Grade Growth

Relative to students at the 10th percentile, growth by students at the:

	25th percentile	Median	75th percentile	90th percentile
<i>Original Scale</i>				
2nd to 3rd	0.97	1.06	1.03	0.95
3rd to 4th	1.10	1.03	1.16	1.34
4th to 5th	0.96	1.15	1.35	1.23
5th to 6th	1.61	2.12	2.17	2.24
6th to 7th	1.21	1.39	1.43	1.62
7th to 8th	1.16	1.17	1.16	1.13
<i>Transformed Scale</i>				
2nd to 3rd	2.63	5.06	6.83	7.90
3rd to 4th	1.69	2.10	2.93	3.95
4th to 5th	1.30	1.94	2.74	2.87
5th to 6th	2.12	3.48	4.25	4.91
6th to 7th	1.60	2.29	2.79	3.50
7th to 8th	1.51	1.89	2.17	2.35

Data: Mathematics test results furnished to author from a Southern state, 2005-06.

Table 6: Comparison of Conventional and Ordinal Measures of Teacher Value Added

Total number of teachers	237
Wilcoxon signed ranks test, p-value	0.0078
Maximum discrepancy in ranks	229
Absolute discrepancy in ranks, 90th percentile	45
Number of statistically significant teacher effects, fixed effect estimate	64
Number of statistically significant teacher effects, ordinal estimate	86
Number of significant effects by both estimates	55
Number ranked in the top quartile	59
Number ranked in the top quartile by both estimates	47
Number ranked in the bottom quartile by both estimates	48
Number ranked above the median by one estimate, below the median by the other	14
Number ranked in the bottom quartile by one, top quartile by the other	3

Data: 5th grade mathematics teachers, large Southern district, 2005-06, author's calculations.

Figure 1: Item Characteristic Curves





Figure 2: A Conjoint Additive Structure, Prior to Scaling



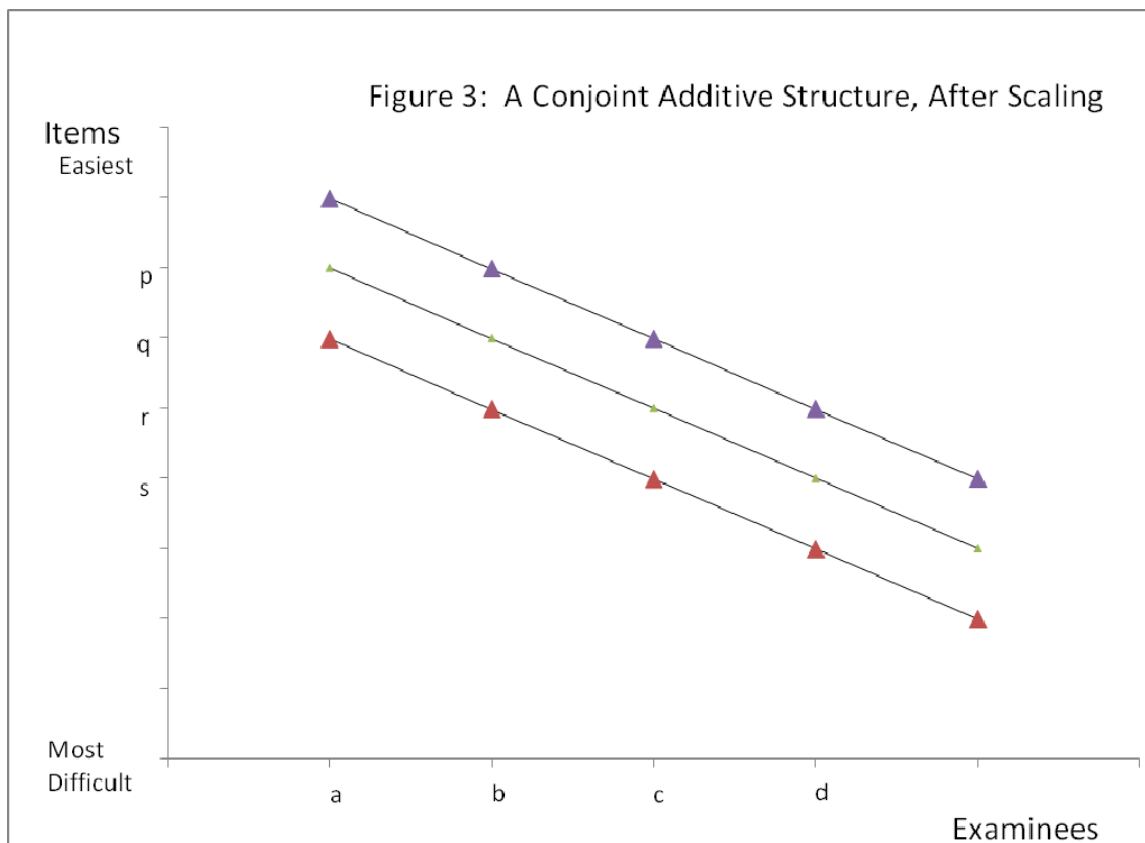



Figure 4: Math Items, Two Hypothetical Students

Student A	Student B
68 equals	43,000 equals:
<ul style="list-style-type: none"> ✓ A. $60 + 8$ B. $60 + 80$ C. $6 + 8$ D. $600 + 8$ E. $6 + 80$ 	<ul style="list-style-type: none"> ✓ A. 4.3×10^3 B. 4.3×10^4 C. 4.3×10^5 D. 43×10^4 E. 43×10^5
$\begin{array}{r} 99 \\ - 56 \\ \hline \end{array}$	$\frac{31}{6} - \left(-3 \frac{3}{8}\right) =$
<ul style="list-style-type: none"> A. 34 B. 155 C. 53 ✓ D. 43 E. 42 	<ul style="list-style-type: none"> ✓ A. $8 \frac{13}{24}$ B. $4 \frac{1}{7}$ C. $1 \frac{19}{24}$ D. $17 \frac{7}{16}$ E. $17 \frac{7}{16}$
$14 \square 6 = 8$ $\square =$	Factor $x^2 - 5x - 36$
<ul style="list-style-type: none"> A. + ✓ B. - C. + D. < E. > 	<ul style="list-style-type: none"> A. $(x - 6)(x - 6)$ B. $(x + 9)(x - 4)$ C. $(x - 36)(x + 1)$ ✓ D. $(x - 9)(x + 4)$ E. $(x + 6)(x + 6)$
What makes us think of a circle?	Using the Pythagorean Theorem, $a^2 + b^2 = c^2$, when $a = 9$ and $b = 12$, then $c = ?$
<ul style="list-style-type: none"> A. block B. pen C. door D. a football field ✓ E. bicycle wheel 	<ul style="list-style-type: none"> A. 8 B. 21 ✓ C. 15 D. $\sqrt{21}$ E. 225


What is the area of the figure?

- A. 18 square units
- B. 9 square units
- ✓ C. 20 square units
- D. 16 square units
- E. 5 square units



Calculate the surface area of this rectangular solid.

- A. 79 cm^2
- B. 110 cm^2
- C. 120 cm^2
- D. 128 cm^2
- ✓ E. 158 cm^2



Student Council Election Results	
Student	Number of Votes
Ann	
Mark	
Sue	

How many votes did Mark get?


- A. 20
- ✓ B. 16
- C. 22
- D. 17
- E. 19

10 students will play basketball after school. They ask 5 more students to play with them.

If you want to find out how many students will play in all, what method should you use?

- A. simplify
- B. subtract
- ✓ C. add
- D. multiply
- E. divide

Ages of the First 37 People to Enter the Museum



What percentage of these people are less than 10 years old? (round to whole percent)

- A. 12%
- B. 33%
- ✓ C. 23%
- D. 40%
- E. 19%

A \$30.00 pair of jeans is discounted 20%. If sales tax is 5%, what will be the final price for the jeans?

- A. \$22.80
- B. \$24.00
- C. \$24.20
- ✓ D. \$25.20
- E. \$28.35

Source: Northwest Evaluation Association RIT Charts, Sample Mathematics Items, 2008.

**Figure 5: Median Between-Grade Gains,
OriginalScale and Transformed Scale**

