

PROSETTA STONE[®] ANALYSIS REPORT
A ROSETTA STONE FOR PATIENT REPORTED OUTCOMES

PROMIS FATIGUE AND FACIT-FATIGUE

SEUNG W. CHOI, TRACY PODRABSKY, NATALIE MCKINNEY, BENJAMIN D. SCHALET, KARON F. COOK
& DAVID CELLA

DEPARTMENT OF MEDICAL SOCIAL SCIENCES
FEINBERG SCHOOL OF MEDICINE
NORTHWESTERN UNIVERSITY

This research was supported by an NIH/National Cancer Institute grant PROSETTA STONE (1RC4CA157236-01, PI: David Cella). Authors acknowledge careful reviews, comments, and suggestions from Drs. Robert Brennan, Lawrence Hedges, Won-Chan Lee, and Nan Rothrock.

Table of Contents

1.	Introduction	3
2.	The PRO Rosetta Stone Project.....	3
2.1.	Patient-Reported Outcomes Measurement Information System (PROMIS).....	4
2.2.	The NIH Toolbox for Assessment of Neurological and Behavioral Function (Toolbox).....	4
2.3.	Quality of Life Outcomes in Neurological Disorders (Neuro-QOL)	5
3.	Legacy Instruments	5
3.1.	Functional Assessment of Chronic Illness Therapy (FACIT)	5
4.	Linking Methods.....	6
4.1.	IRT Linking.....	7
4.2.	Equipercentile Linking.....	8
4.3.	Linking Assumptions	8
5.	Linking Results.....	9
5.8.	PROMIS Fatigue and FACIT-F	10
5.8.1.	Raw Summed Score Distribution	10
5.8.2.	Classical Item Analysis.....	11
5.8.3.	Confirmatory Factor Analysis (CFA)	11
5.8.4.	Item Response Theory (IRT) Linking.....	12
5.8.5.	Raw Score to T-Score Conversion using Linked IRT Parameters.....	14
5.8.6.	Equipercentile Linking.....	14
5.8.7.	Summary and Discussion	16
6.	Appendix Table 22: Raw Score to T-Score Conversion Table (IRT Fixed Parameter Calibration Linking) for FACIT-Fatigue to PROMIS Fatigue (PROMIS Study) RECOMMENDED	18
7.	Appendix Table 23: Direct (Raw to Scale) Equipercentile Crosswalk Table – From FACIT-Fatigue to PROMIS Fatigue (PROMIS Study).	19
8.	Appendix Table 24: Indirect (Raw to Raw to Scale) Equipercentile Crosswalk Table - From FACIT-Fatigue to PROMIS Fatigue (PROMIS Study).....	21

PRO Rosetta Stone (*PROsetta Stone*[®]) Analysis

1. Introduction

A common problem when using a variety of patient-reported outcome measures (PROs) for diverse populations and subgroups is establishing the comparability of scales or units on which the outcomes are reported. The lack of comparability in metrics (e.g., raw summed scores vs. scaled scores) among different PROs poses practical challenges in measuring and comparing effects across different studies. Linking refers to establishing a relationship between scores on two different measures that are not necessarily designed to have the same content or target population. When tests are built in such a way that they differ in content or difficulty, linking must be conducted in order to establish a relationship between the test scores. One technique, commonly referred to as equating, involves the process of converting the system of units of one measure to that of another. This process of deriving equivalent scores has been used successfully in educational assessment to compare test scores obtained from parallel or alternate forms that measure the same characteristic with equal precision. Extending the technique further, comparable scores are sometimes derived for measures of different but related characteristics. The process of establishing comparable scores generally has little effect on the magnitude of association between the measures. Comparability may not signify interchangeability unless the association between the measures approaches the reliability. Equating, the strongest form of linking, can be established only when two tests 1) measure the same content/construct, 2) target very similar populations, 3) are administered under similar conditions such that the constructs measured are not differentially affected, 4) share common measurement goals and 5) are equally reliable. When test forms are created to be similar in content and difficulty, equating adjusts for differences in difficulty. Test forms are considered to be essentially the same, so scores on the two forms can be used interchangeably after equating has adjusted for differences in difficulty. For tests with lesser degrees of similarity, only weaker forms of linking are meaningful, such as calibration, concordance, projection, or moderation.

2. The PRO Rosetta Stone Project

The primary aim of the PRO Rosetta Stone (*PROsetta Stone*[®]) project (1RC4CA157236-01, PI: David Cella) is to develop and apply methods to link the Patient-Reported Outcomes Measurement Information System (PROMIS) measures with other related “legacy” instruments to expand the range of PRO assessment options within a common, standardized metric. The project identifies and applies appropriate linking methods that allow scores on a range of PRO instruments to be expressed as standardized T-score metrics linked to the PROMIS. This preliminary report encompasses the first wave of 20 linking studies based on available PRO data from PROMIS (aka, PROMIS Wave I), Toolbox, and Neuro-QOL.

2.1. Patient-Reported Outcomes Measurement Information System (PROMIS)

In 2004, the NIH initiated the PROMIS¹ cooperative group under the NIH Roadmap² effort to re-engineer the clinical research enterprise. The aim of PROMIS is to revolutionize and standardize how PRO tools are selected and employed in clinical research. To accomplish this, a publicly-available system was developed to allow clinical researchers access to a common repository of items and state-of-the-science computer-based methods to administer the PROMIS measures. The PROMIS measures include item banks across a wide range of domains that comprise physical, mental, and social health for adults and children, with 12-124 items per bank. Initial concepts measured include emotional distress (anger, anxiety, and depression), physical function, fatigue, pain (quality, behavior, and interference), social function, sleep disturbance, and sleep-related impairment. The banks can be used to administer computerized adaptive tests (CAT) or fixed-length forms in these domains. We have also developed 4 to 20-item short forms for each domain, and a 10-item Global Health Scale that includes global ratings of five broad PROMIS domains and general health perceptions. As described in a full issue of *Medical Care* (Cella et al., 2007), the PROMIS items, banks, and short forms were developed using a standardized, rigorous methodology that began with constructing a consensus-based PROMIS domain framework.

All PROMIS banks have been calibrated according to Samejima's (1969) graded response model (based on large data collections including both general and clinical samples) and re-scaled (mean=50 and SD=10) using scale-setting subsamples matching the marginal distributions of gender, age, race, and education in the 2000 US census. The PROMIS Wave I calibration data included a small number of full-bank testing cases (approximately 1,000 per bank) from a general population taking one full bank and a larger number of block-administration cases (n= ~14,000) from both general and clinical populations taking a collection of blocks representing all banks with 7 items each. The full-bank testing samples were randomly assigned to one of 7 different forms. Each form was composed of one or more PROMIS domains (with an exception of Physical Function where the bank was split over two forms) and one or more legacy measures of the same or related domains.

The PROMIS Wave I data collection design included a number of widely accepted "legacy" measures. The legacy measures used for validation evidence included Buss-Perry Aggression Questionnaire (BPAQ), Center for Epidemiological Studies Depression Scale (CES-D), Mood and Anxiety Symptom Questionnaire (MASQ), Functional Assessment of Chronic Illness Therapy-Fatigue (FACIT-F), Brief Pain Inventory (BPI), and SF-36. In addition to the pairs for validity (e.g., PROMIS Depression and CES-D), the PROMIS Wave I data allows for the potential for linking over a dozen pairs of measures/subscales. Furthermore, included within each of the PROMIS banks were items from many other existing measures. Depending on the nature and strength of relationship between the measures, various linking procedures can be used to allow for cross-walking of scores.

2.2. The NIH Toolbox for Assessment of Neurological and Behavioral Function (Toolbox)

Developed in 2006 with the NIH Blueprint funding for Neuroscience Research, four domains of assessment central to neurological and behavioral function were created to measure cognition, sensation, motor functioning, and emotional health. The NIH Toolbox for Assessment of Neurological

¹ www.nihpromis.org

² www.nihroadmap.nih.gov

and Behavioral Function³ provides investigators with a brief, yet comprehensive measurement tool for assessment of cognitive function, emotional health, sensory and motor function. It provides an innovative approach to measurement that is responsive to the needs of researchers in a variety of settings, with a particular emphasis on measuring outcomes in clinical trials and functional status in large cohort studies, e.g. epidemiological studies and longitudinal studies. Included as subdomains of emotional health were negative affect, psychological well-being, stress and self-efficacy, and social relationships. Three PROMIS emotional distress item banks (Anger, Anxiety, and Depression) were used as measures of negative affect. Additionally, existing “legacy” measures, e.g., Patient Health Questionnaire (PHQ-9) and Center for Epidemiological Studies Depression Scale (CES-D), were flagged as potential candidates for the Toolbox battery because of their history, visibility, and research legacy. Among these legacy measures, we focused on those that were available without proprietary restrictions for research applications. In most cases, these measures had been developed using classical test theory.

2.3. Quality of Life Outcomes in Neurological Disorders (Neuro-QOL)

The National Institute of Neurological Disorders and Stroke sponsored a multi-site project to develop a clinically relevant and psychometrically robust Quality of Life (QOL) assessment tool for adults and children with neurological disorders. The primary goal of this effort, known as Neuro-QOL³, was to enable clinical researchers to compare the QOL impact of different interventions within and across various conditions. This resulted in 13 adult QOL item banks (Anxiety, Depression, Fatigue, Upper Extremity Function - Fine Motor, Lower Extremity Function - Mobility, Applied Cognition - General Concerns, Applied Cognition - Executive Function, Emotional and Behavioral Dyscontrol, Positive Affect and Well-Being, Sleep Disturbance, Ability to Participate in Social Roles and Activities, Satisfaction with Social Roles and Activities, and Stigma).

3. Legacy Instruments

The following instruments are widely accepted “legacy” measures that have been used as part of the initial validation work for PROMIS and Toolbox. Data were collected on a minimum of 500 respondents (for stable item parameter estimation) along with at least one other conceptually similar scale or bank.

3.1. Functional Assessment of Chronic Illness Therapy (FACIT)

The Functional Assessment of Chronic Illness Therapy (FACIT) Measurement System is a collection of QOL questionnaires targeted to the management of chronic illness including cancer. The FACT-G (now in Version 4) is a 27-item compilation of general questions divided into four subscales: Physical Well-Being, Social/Family Well-Being, Emotional Well-Being, and Functional Well-Being. It is considered appropriate for use with patients with any form of cancer, and has also been used and validated in other chronic illness conditions (e.g., HIV/AIDS, multiple sclerosis) and in the general population (using a slightly modified version). Validation of a core measure allowed for the evolution of multiple disease, treatment, condition, and non-cancer-specific subscales. FACIT subscales are constructed to complement the FACT-G, addressing relevant disease-, treatment-, or condition-related issues not

³ www.nihtoolbox.org

already covered in the general questionnaire. Each is intended to be as specific as necessary to capture the clinically-relevant problems associated with a given condition or symptom, yet general enough to allow for comparison across diseases, and extension, as appropriate, to other chronic medical conditions. For the current analysis, we used the Fatigue scale. The Functional Assessment of Chronic Illness Therapy-Fatigue Scale (FACIT-Fatigue scale) is a 13-item questionnaire that assesses self-reported fatigue and its impact upon daily activities and function (Yellen, Cella, Webster, Blendowsky, & Kaplan, 1997). It was developed to meet a growing demand for the precise evaluation of fatigue associated with anemia in cancer patients. Subsequently, it has been employed in over 70 published studies including over 20,000 people, including cancer patients receiving chemotherapy (Berndt et al., 2005; Quirt et al., 2001), cancer patients not receiving chemotherapy (Quirt et al., 2001; Quirt et al., 2002), long term cancer survivors (Ng et al., 2005), childhood cancer survivors (Mulrooney et al., 2008), rheumatoid arthritis (Cella et al., 2005; Mease et al., 2008; Mittendorf et al., 2007), psoriatic arthritis (Chandran, Bhella, Schentag & Gladman, 2007), paroxysmal nocturnal hemoglobinuria (Brodsky et al., 2008), and Parkinson's disease (Hagell et al., 2006). It has also been validated in the general United States population (Brucker, Yost, Cashy, Webster & Cella., 2005; Cella, Lai, Chang, Peterman & Slavin, 2002). In all cases, the FACIT-Fatigue scale has been found to be reliable and valid.

4. Linking Methods

PROMIS full-bank administration allows for single group linking. This linking method is used when two or more measures are administered to the same group of people. For example, two PROMIS banks (Anxiety and Depression) and three legacy measures (MASQ, CES-D, and SF-36/MH) were administered to a sample of 925 people. The order of measures was randomized so as to minimize potential order effects. The original purpose of the full-bank administration study was to establish initial validity evidence (e.g., validity coefficients), not to establish linking relationships. Some of the measures revealed severely skewed score distributions in the full-bank administration sample and the sample size was relatively small, which might be limiting factors when it comes to determining the linking method. Another potential issue is related to how the non-PROMIS measures are scored and reported. For example, all SF-36 subscales are scored using a proprietary scoring algorithm and reported as normed scores (0 to 100). Others are scored and reported using simple raw summed scores. All PROMIS measures are scored using the final re-centered item response theory (IRT) item parameters and transformed to the T-score metric (mean=50, SD=10).

PROMIS's T-score distributions are standardized such that a score of 50 represents the average (mean) for the US general population, and the standard deviation around that mean is 10 points. A high PROMIS score always represents more of the concept being measured. Thus, for example, a person who has a T-score of 60 is one standard deviation higher than the general population for the concept being measured. For symptoms and other negatively-worded concepts like pain, fatigue, and anxiety, a score of 60 is one standard deviation worse than average; for functional scores and other positively-worded concepts like physical or social function, a score of 60 is one standard deviation better than average, etc.

In order to apply the linking methods consistently across different studies, linking/concordance relationships will be established based on the raw summed score metric of the measures. Furthermore, the direction of linking relationships to be established will be from legacy to PROMIS. That is, each raw summed score on a given legacy instrument will be mapped to a T-score of the corresponding PROMIS

instrument/bank. Finally, the raw summed score for each legacy instrument was constructed such that higher scores represent higher levels of the construct being measured. When the measures were scaled in the opposite direction, we reversed the direction of the legacy measure in order for the correlation between the measures to be positive and to facilitate concurrent calibration. As a result, some or all item response scores for some legacy instruments will need to be reverse-coded.

4.1. IRT Linking

One of the objectives of the current linking analysis is to determine whether or not the non-PROMIS measures can be added to their respective PROMIS item bank without significantly altering the underlying trait being measured. The rationale is twofold: (1) the augmented PROMIS item banks might provide more robust coverage both in terms of content and difficulty; and (2) calibrating the non-PROMIS measures on the corresponding PROMIS item bank scale might facilitate subsequent linking analyses. At least, two IRT linking approaches are applicable under the current study design; (1) linking separate calibrations through the Stocking-Lord method and (2) fixed parameter calibration.

Linking separate calibrations might involve the following steps under the current setting.

- First, simultaneously calibrate the combined item set (e.g., PROMIS Depression bank and CES-D).
- Second, estimate linear transformation coefficients (additive and multiplicative constants) using the item parameters for the PROMIS bank items as anchor items.
- Third, transform the metric for the non-PROMIS items to the PROMIS metric.

The second approach, fixed parameter calibration, involves fixing the PROMIS item parameters at their final bank values and calibrating only non-PROMIS items so that the non-PROMIS item parameters may be placed on the same metric as the PROMIS items. The focus is on placing the parameters of non-PROMIS items on the PROMIS scale. Updating the PROMIS item parameters is not desired because the linking exercise is built on the stability of these calibrations. Note that IRT linking would be necessary when the ability level of the full-bank testing sample is different from that of the PROMIS scale-setting sample. If it is assumed that the two samples are from the same population, linking is not necessary and calibration of the items (either separately or simultaneously) will result in item parameter estimates that are on the same scale without any further scale linking. Even though the full-bank testing sample was a subset of the full PROMIS calibration sample, it is still possible that the two samples are somewhat disparate due to some non-random component of the selection process. Moreover, there is some evidence that linking can improve the accuracy of parameter estimation even when linking is not necessary (e.g., two samples are from the same population having the same or similar ability levels). Thus, conducting IRT linking would be worthwhile.

Once the non-PROMIS items are calibrated on the corresponding PROMIS item bank scale, the augmented item bank can be used for standard computation of IRT scaled scores from any subset of the items, including computerized adaptive testing (CAT) and creating short forms. The non-PROMIS items will be treated the same as the existing PROMIS items. Again, the above options are feasible only when the dimensionality of the bank is not altered significantly (i.e., where a unidimensional IRT model is suitable for the aggregate set of items). Thus, prior to conducting IRT linking, it is important to assess dimensionality of the measures based on some selected combinations of PROMIS and non-PROMIS measures. Various dimensionality assessment tools can be used including a confirmatory factor analysis, disattenuated correlations, and essential unidimensionality.

4.2. Equipercentile Linking

The IRT Linking procedures described above are permissible only if the traits being measured are not significantly altered by aggregating items from multiple measures. One potential issue might be creating multidimensionality as a result of aggregating items measuring different traits. For two scales that measure distinct but highly related traits, predicting scores on one scale from those of the other has been used frequently. Concordance tables between PROMIS and non-PROMIS measures can be constructed using equipercentile equating (Lord, 1982; Kolen & Brennan, 2004) when there is insufficient empirical evidence that the instruments measure the same construct. An equipercentile method estimates a nonlinear linking relationship using percentile rank distributions of the two linking measures. The equipercentile linking method can be used in conjunction with a presmoothing method such as the loglinear model (Hanson, Zeng, & Colton, 1994). The frequency distributions are first smoothed using the loglinear model and then equipercentile linking is conducted based on the smoothed frequency distributions of the two measures. Smoothing can also be done at the backend on equipercentile equivalents and is called postsmoothing (Brennan, 2004; Kolen & Brennan, 2004). The cubic-spline smoothing algorithm (Reinsch, 1967) is used in the LEGS program (Brennan, 2004). Smoothing is intended to reduce sampling error involved in the linking process. A successful linking procedure will provide a conversion (crosswalk) table, in which, for example, raw summed scores on the PHQ-9 measure are transformed to the T-score equivalents of the PROMIS Depression measure.

Under the current context, equipercentile crosswalk tables can be generated using two different approaches. First is a direct linking approach where each raw summed score on non-PROMIS measure is mapped directly to a PROMIS T-score. That is, raw summed scores on the non-PROMIS instrument and IRT scaled scores on the PROMIS (reference) instrument are linked directly, although raw summed scores and IRT scaled score have distinct properties (e.g., discrete vs. continuous). This approach might be appropriate when the reference instrument is either an item bank or composed of a large number of items and so various subsets (static or dynamic) are likely to be used but not the full bank in its entirety (e.g., PROMIS Physical Function bank with 124 items). Second is an indirect approach where raw summed scores on the non-PROMIS instrument are mapped to raw summed scores on the PROMIS instrument; and then the resulting raw summed score equivalents are mapped to corresponding scaled scores based on a raw-to-scale score conversion table. Because the raw summed score equivalents may take fractional values, such a conversion table will need to be interpolated using statistical procedures (e.g., cubic spline).

Finally, when samples are small or inadequate for a specific method, random sampling error becomes a major concern (Kolen & Brennan, 2004). That is, substantially different linking relationships might be obtained if linking is conducted repeatedly over different samples. The type of random sampling error can be measured by the standard error of equating (SEE), which can be operationalized as the standard deviation of equated scores for a given raw summed score over replications (Lord, 1982).

4.3. Linking Assumptions

In Section 5, we present the results of a large number of linking studies using secondary data sets. In each case, we have applied all three linking methods described in sections 4.1 and 4.2. Our purpose is to provide the maximum amount of useful information. However, the suitability of these methods depends upon the meeting of various linking assumptions. These assumptions require that the two instruments

to be linked measure the same construct, show a high correlation, and are relatively invariant in subpopulation differences (Dorans, 2007). The degree to which these assumptions are met varies across linking studies. Given that different researchers may interpret these requirements differently, we have taken a liberal approach for inclusion of linkages in this book. Nevertheless, we recommend that researchers diagnostically review the classical psychometrics and CFA results in light of these assumptions prior to any application of the cross-walk charts or legacy parameters to their own data.

5. Linking Results

Table 5.1 lists the linking analyses included in this report, which have been conducted based on samples from two different studies: PROMIS and Toolbox (see Section 2 for more details). In all cases, PROMIS instruments were used as the reference (i.e., scores on non-PROMIS instruments are expressed on the PROMIS score metric); however, shorter versions of PROMIS were used in Toolbox.

Table 5.1. Linking by Study

Section	Study	PROMIS Instrument	Non-PROMIS Instrument to Link
5.8	PROMIS Wave1	Fatigue	Functional Assessment of Chronic Illness Therapy – Fatigue Scale (FACIT-F)

5.8.PROMIS Fatigue and FACIT-F

In this section we provide a summary of the procedures employed to establish a crosswalk between two measures of Fatigue, namely the PROMIS Fatigue item bank (95 items) and FACIT-F (13 items). PROMIS Fatigue was scaled such that higher scores represent higher levels of Fatigue; for the FACIT-F, higher scores represent lower levels of Fatigue. We created raw summed scores for each of the measures separately and then for the combined. Summing of item scores assumes that all items have positive correlations with the total as examined in the section on Classical Item Analysis.

5.8.1. Raw Summed Score Distribution

The maximum possible raw summed scores were 475 for PROMIS Fatigue and 53 for FACIT-F. Figure 5.8.1 and Figure 5.8.2 graphically display the raw summed score distributions of the two measures. Figure 5.8.3 shows the distribution for the combined. Figure 5.8.4 is a scatter plot matrix showing the relationship of each pair of raw summed scores. Pearson correlations are shown above the diagonal. The correlation between PROMIS Fatigue and FACIT-F was -0.96. The disattenuated (corrected for unreliabilities) correlation between PROMIS Fatigue and FACIT-F was -0.99. The correlations between the combined score and the measures were 1.00 and 0.96 for PROMIS Fatigue and FACIT-F, respectively.

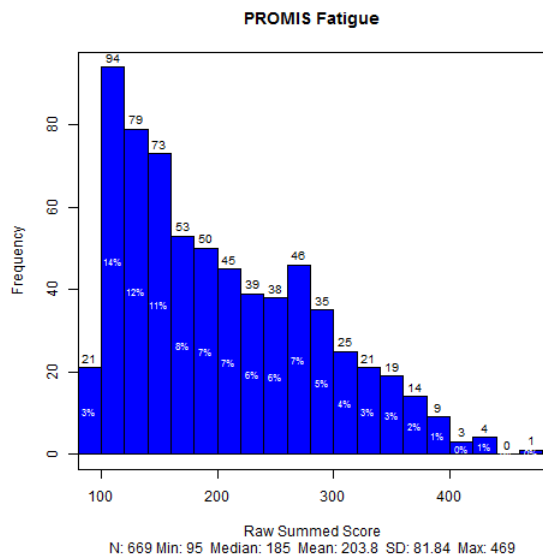


Figure 5.8.1: Raw Summed Score Distribution - PROMIS Instrument

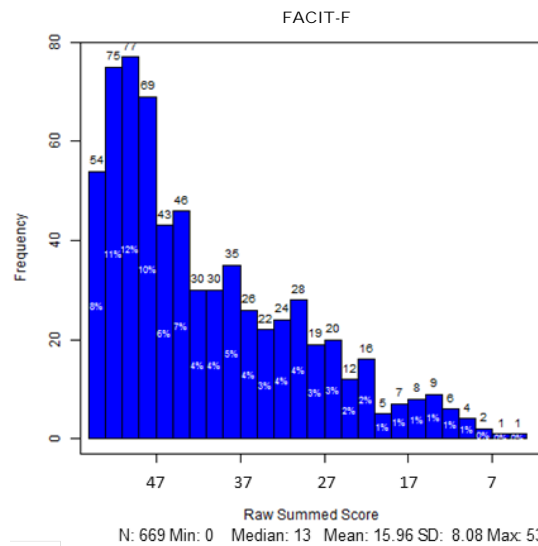


Figure 5.8.2: Raw Summed Score Distribution - Linking Instrument

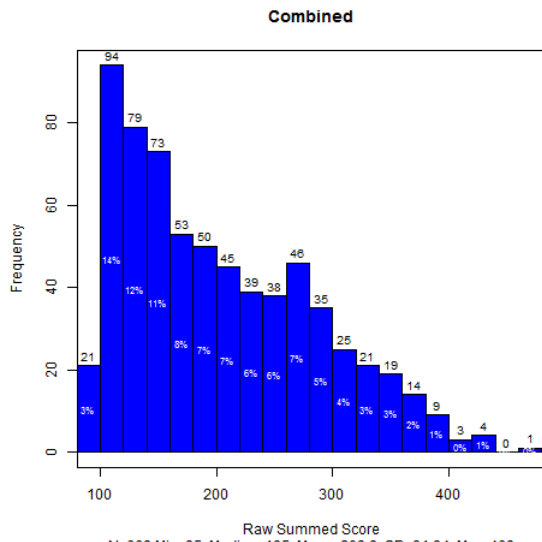


Figure 5.8.3: Raw Summed Score Distribution – Combined

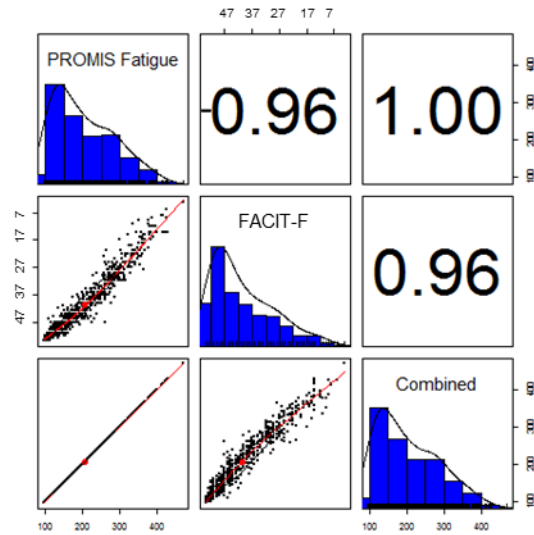


Figure 5.8.4: Scatter Plot Matrix of Raw Summed Scores

5.8.2. Classical Item Analysis

We conducted classical item analyses on the two measures separately and on the combined. Table 5.8.1 summarizes the results. For PROMIS Fatigue, Cronbach's alpha internal consistency reliability estimate was 0.994 and adjusted (corrected for overlap) item-total correlations ranged from 0.509 to 0.883. For FACIT-F, alpha was 0.958 and adjusted item-total correlations ranged from 0.610 to 0.877. For the 95 items, alpha was 0.994 and adjusted item-total correlations ranged from 0.509 to 0.883.

Table 5.8.1: Classical Item Analysis

	No. Items	Cronbach's Alpha Internal Consistency Reliability Estimate	Adjusted (corrected for overlap) Item-total Correlation		
			Minimum	Mean	Maximum
PROMIS Fatigue	95	0.994	0.509	0.805	0.883
FACIT-F	13	0.958	0.610	0.781	0.877
Combined	95	0.994	0.509	0.805	0.883

5.8.3. Confirmatory Factor Analysis (CFA)

To assess the dimensionality of the measures, a categorical confirmatory factor analysis (CFA) was carried out using the WLSMV estimator of Mplus on a subset of cases without missing responses. A single factor model (based on polychoric correlations) was run on each of the two measures separately and on the combined. Table 5.8.2 summarizes the model fit statistics. For PROMIS Fatigue, the fit statistics were as follows: CFI = 0.969, TLI = 0.968, and RMSEA = 0.062. For FACIT-F, CFI = 0.989, TLI = 0.986, and RMSEA = 0.102. For the 95 items, CFI = 0.969, TLI = 0.968, and RMSEA = 0.062. The main interest of the current analysis is whether the combined measure is essentially unidimensional.

Table 5.8.2: CFA Fit Statistics

	No. Items	n	CFI	TLI	RMSEA
PROMIS Fatigue	95	738	0.969	0.968	0.062
FACIT-F	13	738	0.989	0.986	0.102
Combined	95	738	0.969	0.968	0.062

5.8.4. Item Response Theory (IRT) Linking

We conducted concurrent calibration on the combined set of 95 items according to the graded response model. The calibration was run using `MULTILOG` and two different approaches as described previously (i.e., IRT linking vs. fixed-parameter calibration). For IRT linking, all 95 items were calibrated freely on the conventional theta metric (mean=0, SD=1). Then the 95 PROMIS Fatigue items served as anchor items to transform the item parameter estimates for the FACIT-F items onto the PROMIS Fatigue metric. We used four IRT linking methods implemented in `plink` (Weeks, 2010): mean/mean, mean/sigma, Haebara, and Stocking-Lord. The first two methods are based on the mean and standard deviation of item parameter estimates, whereas the latter two are based on the item and test information curves. Table 5.8.3 shows the additive (A) and multiplicative (B) transformation constants derived from the four linking methods. For fixed-parameter calibration, the item parameters for the PROMIS items were constrained to their final bank values, while the FACIT-F items were calibrated under the constraints imposed by the anchor items.

Table 5.8.3: IRT Linking Constants

	A	B
Mean/Mean	1.012	0.588
Mean/Sigma	1.002	0.591
Haebara	1.004	0.595
Stocking-Lord	1.002	0.596

The item parameter estimates for the FACIT-F items were linked to the PROMIS Fatigue metric using the transformation constants shown in Table 5.8.3. The FACIT-F item parameter estimates from the fixed-parameter calibration are considered already on the PROMIS Fatigue metric. Based on the transformed and fixed-parameter estimates we derived test characteristic curves (TCC) for FACIT-F as shown in Figure 5.8.5. Using the fixed-parameter calibration as a basis we then examined the difference with each of the TCCs from the four linking methods. Figure 5.8.6 displays the differences on the vertical axis.

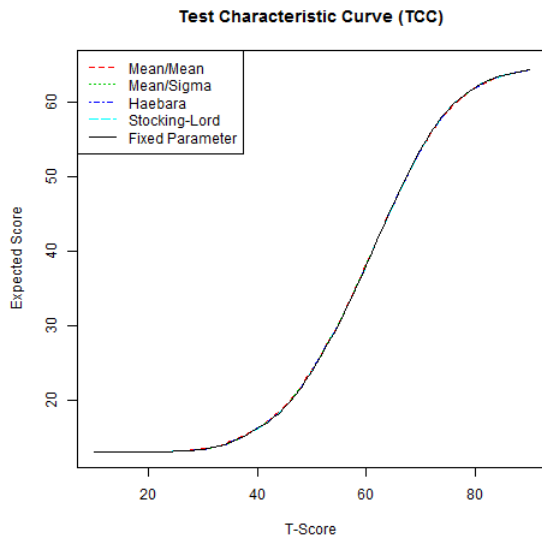


Figure 5.8.5: Test Characteristic Curves (TCC) from Different Linking Methods

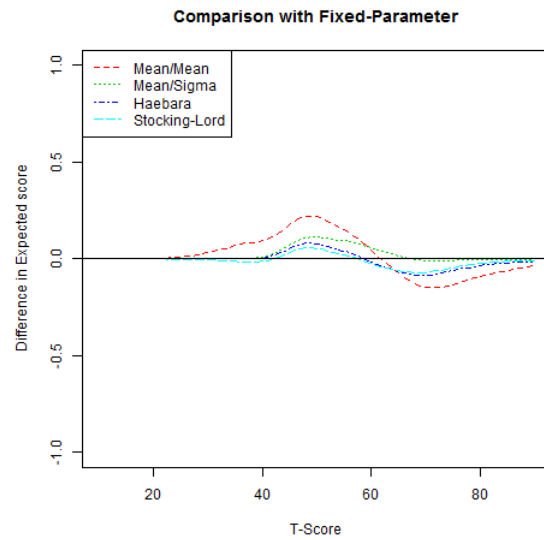


Figure 5.8.6: Difference in Test Characteristic Curves (TCC)

Table 5.8.4 shows the fixed-parameter calibration item parameter estimates for FACIT-F. The marginal reliability estimate for FACIT-F based on the item parameter estimates was 0.945. The marginal reliability estimates for PROMIS Fatigue and the combined set were 0.991 and 0.991, respectively. The slope parameter estimates for FACIT-F ranged from 1.64 to 4.35 with a mean of 3.1. The slope parameter estimates for PROMIS Fatigue ranged from 1.17 to 4.77 with a mean of 3.17. We also derived scale information functions based on the fixed-parameter calibration result. Figure 5.8.7 displays the scale information functions for PROMIS Fatigue, FACIT-F, and the combined set of 95. We then computed IRT scaled scores for the three measures based on the fixed-parameter calibration result. Figure 5.8.8 is a scatter plot matrix showing the relationships between the measures.

Table 5.8.4: Fixed-Parameter Calibration Item Parameter Estimates

Slope	Threshold 1	Threshold 2	Threshold 3	Threshold 4
4.320	-1.140	0.081	0.956	1.770
2.690	0.161	0.916	1.660	2.420
3.270	-0.203	0.744	1.370	2.270
3.300	-1.360	0.099	0.844	1.800
4.350	-0.311	0.504	1.180	1.960
3.400	-0.341	0.582	1.280	2.140
1.640	-0.433	0.874	1.560	2.520
2.310	0.948	1.720	2.470	3.540
2.310	0.753	1.520	2.200	3.020
3.900	-0.088	0.613	1.070	1.430
3.610	0.147	0.829	1.420	1.870
2.710	-1.390	0.003	1.010	2.170
2.550	-0.498	0.597	1.680	2.600

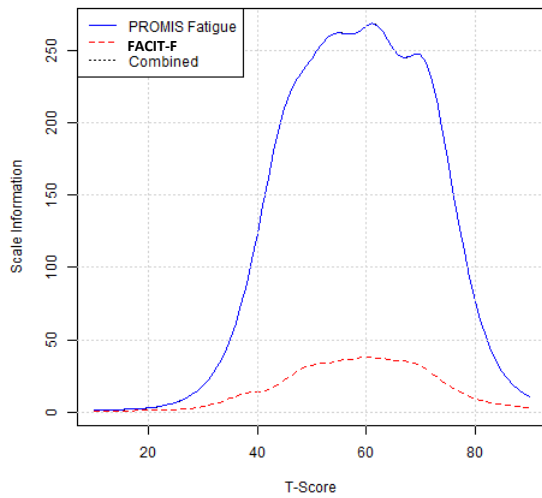


Figure 5.8.7: Comparison of Scale Information Functions

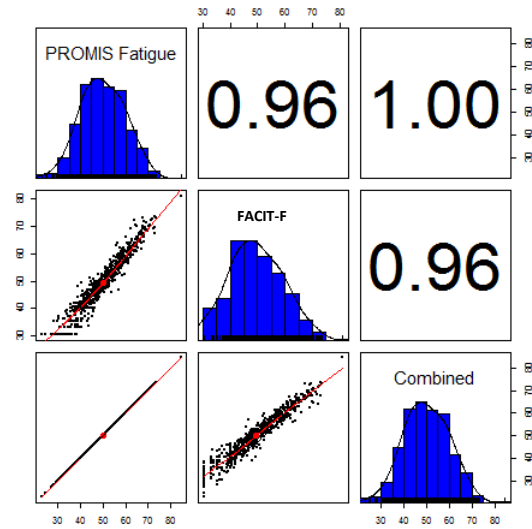


Figure 5.8.8: Comparison of IRT Scaled Scores

5.8.5. Raw Score to T-Score Conversion using Linked IRT Parameters

The IRT model implemented in PROMIS (i.e., the graded response model) uses the pattern of item responses for scoring, not just the sum of individual item scores. However, a crosswalk table mapping each raw summed score point on FACIT-F to a scaled score on PROMIS Fatigue can be useful. Based on the FACIT-F item parameters derived from the fixed-parameter calibration, we constructed a score conversion table. The conversion table displayed in Appendix Table 22 can be used to map simple raw summed scores from FACIT-F to T-score values linked to the PROMIS Fatigue metric. Each raw summed score point and corresponding PROMIS scaled score are presented along with the standard error associated with the scaled score. The raw summed score is constructed such that for each item, consecutive integers in base 1 are assigned to the ordered response categories.

5.8.6. Equipercentile Linking

We mapped each raw summed score point on FACIT-F to a corresponding scaled score on PROMIS Fatigue by identifying scores on PROMIS Fatigue that have the same percentile ranks as scores on FACIT-F. Theoretically, the equipercentile linking function is symmetrical for continuous random variables (X and Y). Therefore, the linking function for the values in X to those in Y is the same as that for the values in Y to those in X. However, for discrete variables like raw summed scores the equipercentile linking functions can be slightly different (due to rounding errors and differences in score ranges) and hence may need to be obtained separately. Figure 5.8.9 displays the cumulative distribution functions of the measures. Figure 5.8.10 shows the equipercentile linking functions based on raw summed scores, from FACIT-F to PROMIS Fatigue. When the number of raw summed score points differs substantially, the equipercentile linking functions could deviate from each other noticeably. The problem can be exacerbated when the sample size is small. Appendix Table 23 and Appendix Table 24 show the equipercentile crosswalk tables. The result shown in Appendix Table 23 is based on the direct (raw summed score to scaled score) approach, whereas Appendix Table 24 shows the result based on the

indirect (raw summed score to raw summed score equivalent to scaled score equivalent) approach (Refer to Section 4.2 for details). Three separate equipercentile equivalents are presented: one is equipercentile without post smoothing (“Equipercentile Scale Score Equivalents”) and two with different levels of postsmoothing, i.e., “Equipercentile Equivalents with Postsmoothing (Less Smoothing)” and “Equipercentile Equivalents with Postsmoothing (More Smoothing).” Postsmoothing values of 0.3 and 1.0 were used for “Less” and “More,” respectively (Refer to Brennan, 2004 for details).

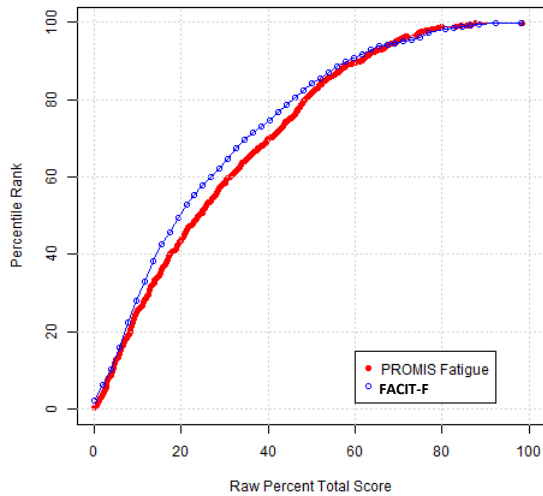


Figure 5.8.9: Comparison of Cumulative Distribution Functions based on Raw Summed Scores

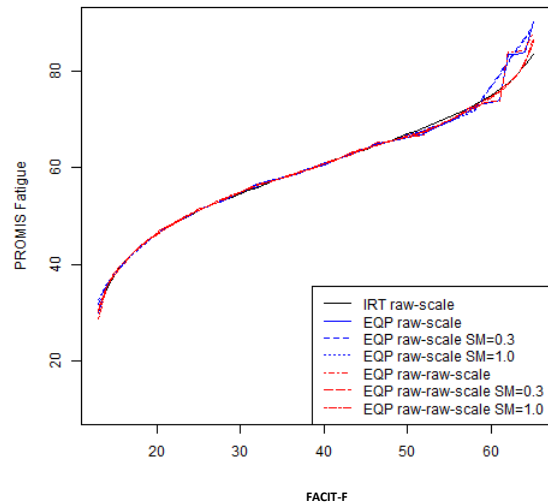


Figure 5.8.10: Equipercentile Linking Functions

5.8.7. Summary and Discussion

The purpose of linking is to establish the relationship between scores on two measures of closely related traits. The relationship can vary across linking methods and samples employed. In equipercentile linking, the relationship is determined based on the distributions of scores in a given sample. Although IRT-based linking can potentially offer sample-invariant results, they are based on estimates of item parameters, and hence subject to sampling errors. A potential issue with IRT-based linking methods is, however, the violation of model assumptions as a result of combining items from two measures (e.g., unidimensionality and local independence). As displayed in Figure 5.8.10, the relationships derived from various linking methods are consistent, which suggests that a robust linking relationship can be determined based on the given sample.

To further facilitate the comparison of the linking methods, Table 5.8.5 reports four statistics summarizing the current sample in terms of the differences between the PROMIS Fatigue T-scores and FACIT-F scores linked to the T-score metric through different methods. In addition to the seven linking methods previously discussed (see Figure 5.8.10), the method labeled “IRT pattern scoring” refers to IRT scoring based on the pattern of item responses instead of raw summed scores. With respect to the correlation between observed and linked T-scores, IRT pattern scoring produced the best result (0.958), followed by EQP raw-scale SM=0.0 (0.955). Similar results were found in terms of the standard deviation of differences and root mean squared difference (RMSD). IRT pattern scoring yielded smallest RMSD (2.867), followed by EQP raw-scale SM=0.3 (2.942).

Table 5.8.5: Observed vs. Linked T-scores

Methods	Correlation	Mean Difference	SD Difference	RMSD
IRT pattern scoring	0.958	0.472	2.830	2.867
IRT raw-scale	0.955	-0.026	2.952	2.950
EQP raw-scale SM=0.0	0.955	-0.029	2.948	2.946
EQP raw-scale SM=0.3	0.955	-0.127	2.941	2.942
EQP raw-scale SM=1.0	0.954	-0.158	2.944	2.946
EQP raw-raw-scale SM=0.0	0.955	-0.056	2.944	2.943
EQP raw-raw-scale SM=0.3	0.955	-0.054	2.951	2.950
EQP raw-raw-scale SM=1.0	0.954	-0.050	2.978	2.976

One approach to evaluating the robustness of a linking relationship is comparing the observed and linked scores in a new sample independent of the sample from which the linking relationship was obtained. Such a sample can be used to examine empirically the bias and standard error of different linking results. Because of the small sample size (N=669), however, subsetting out a sample was not feasible. Instead, a resampling study was used where small subsets of cases (e.g., 25, 50, and 75) were drawn with replacement from the study sample (N=669) over a large number of replications (i.e., 10,000).

Table 5.8.6 summarizes the mean and standard deviation of differences between the observed and linked T-scores by linking method and sample size. For each replication, the mean difference between the observed and equated PROMIS Fatigue T-scores was computed. Then the mean and the standard deviation of the means were computed over replications as bias and empirical standard error, respectively. As the sample size increased (from 25 to 75), the empirical standard error decreased steadily. At a sample size of 75, IRT pattern scoring produced the smallest standard error, 0.31. That is,

the difference between the mean PROMIS Fatigue T-score and the mean equated FACIT-F T-score based on a similar sample of 75 cases is expected to be around ± 0.62 (i.e., 2×0.31).

Table 5.8.6: Comparison of Resampling Results

Methods	Mean (N=25)	SD (N=25)	Mean (N=50)	SD (N=50)	Mean (N=75)	SD (N=75)
IRT pattern scoring	0.481	0.556	0.471	0.383	0.470	0.310
IRT raw-scale	-0.030	0.577	-0.023	0.402	-0.024	0.323
EQP raw-scale SM=0.0	-0.018	0.584	-0.035	0.399	-0.031	0.322
EQP raw-scale SM=0.3	-0.127	0.576	-0.130	0.401	-0.126	0.319
EQP raw-scale SM=1.0	-0.165	0.581	-0.162	0.400	-0.163	0.315
EQP raw-raw-scale SM=0.0	-0.067	0.585	-0.075	0.399	-0.062	0.319
EQP raw-raw-scale SM=0.3	-0.053	0.580	-0.057	0.403	-0.052	0.322
EQP raw-raw-scale SM=1.0	-0.049	0.578	-0.049	0.407	-0.047	0.325

Examining a number of linking studies in the current project revealed that the two linking methods (IRT and equipercentile) in general produced highly comparable results. Some noticeable discrepancies were observed (albeit rarely) in some extreme score levels where data were sparse. Model-based approaches can provide more robust results than those relying solely on data when data is sparse. The caveat is that the model should fit the data reasonably well. One of the potential advantages of IRT-based linking is that the item parameters on the linking instrument can be expressed on the metric of the reference instrument, and therefore can be combined without significantly altering the underlying trait being measured. As a result, a larger item pool might be available for computerized adaptive testing or various subsets of items can be used in static short forms. Therefore, IRT-based linking (Appendix Table 22) might be preferred when the results are comparable and no apparent violations of assumptions are evident.

6. Appendix Table 22: Raw Score to T-Score Conversion Table (IRT Fixed Parameter Calibration Linking) for FACIT-Fatigue to PROMIS Fatigue (PROMIS Study) RECOMMENDED

FACIT-F Score	PROMIS T-score	SE	FACIT-F Score	PROMIS T-score	SE
52	30.3	4.8	12	68.9	2.0
51	35.0	3.5	11	69.6	2.0
50	38.0	3.0	10	70.4	2.0
49	40.3	2.8	9	71.2	2.1
48	42.1	2.6	8	72.0	2.2
47	43.7	2.5	7	72.9	2.3
46	45.0	2.3	6	73.9	2.4
45	46.3	2.2	5	75.0	2.5
44	47.3	2.1	4	76.2	2.7
43	48.3	2.0	3	77.5	2.9
42	49.3	2.0	2	79.1	3.1
41	50.1	1.9	1	81.2	3.3
40	51.0	1.9	0	83.5	3.4
39	51.7	1.9			
38	52.5	1.9			
37	53.2	1.9			
36	53.9	1.8			
35	54.6	1.8			
34	55.3	1.8			
33	55.9	1.8			
32	56.6	1.8			
31	57.2	1.8			
30	57.8	1.8			
29	58.4	1.8			
28	59.0	1.8			
27	59.6	1.8			
26	60.2	1.8			
25	60.8	1.8			
24	61.4	1.8			
23	62.0	1.8			
22	62.6	1.8			
21	63.2	1.8			
20	63.8	1.8			
19	64.4	1.8			
18	65.0	1.8			
17	65.6	1.8			
16	66.2	1.9			
15	66.9	1.9			
14	67.5	1.9			
13	68.2	1.9			

7. Appendix Table 23: Direct (Raw to Scale) Equipercentile Crosswalk Table – From FACIT-Fatigue to PROMIS Fatigue (PROMIS Study). Note: Table 22 is recommended.

FACIT-F Score	Equipercentile PROMIS Scaled Score Equivalents (No Smoothing)	Equipercentile Equivalents with Postsmoothing (Less Smoothing)	Equipercentile Equivalents with Postsmoothing (More Smoothing)	Standard Error of Equating (SEE)
52	30	32	33	1.28
51	36	36	36	0.60
50	38	38	38	0.32
49	40	40	40	0.27
48	42	42	42	0.21
47	44	44	43	0.24
46	45	45	45	0.53
45	46	46	46	0.19
44	48	47	47	0.24
43	48	49	48	0.21
42	50	49	49	0.33
41	50	50	50	0.24
40	51	51	51	0.26
39	52	52	52	0.29
38	52	53	53	0.26
37	53	53	53	0.28
36	54	54	54	0.30
35	55	55	55	0.41
34	56	56	55	0.33
33	56	56	56	0.26
32	57	57	57	0.21
31	57	57	57	0.20
30	58	58	58	0.18
29	58	58	59	0.20
28	59	59	59	0.30
27	59	60	60	0.26
26	60	60	60	0.29
25	60	61	61	0.45
24	61	61	62	0.39
23	62	62	62	0.39
22	63	63	63	0.70
21	64	63	63	0.33
20	64	64	64	0.30
19	65	64	64	0.56
18	65	65	65	0.45
17	66	65	66	0.36
16	66	66	66	0.32
15	66	66	67	0.30
14	66	67	67	0.31
13	67	67	68	0.53

FACIT-F Score	Equipercntile PROMIS Scaled Score Equivalents (No Smoothing)	Equipercntile Equivalents with Postsmoothing (Less Smoothing)	Equipercntile Equivalents with Postsmoothing (More Smoothing)	Standard Error of Equating (SEE)
12	68	68	68	0.35
11	69	69	69	1.09
10	70	70	70	1.09
9	70	70	70	0.75
8	72	71	71	0.82
7	72	72	72	0.67
6	73	74	74	0.67
5	74	77	77	0.47
4	74	79	79	1.41
3	84	82	82	1.41
2	84	84	84	1.41
1	84	87	87	1.41
0	90	89	89	1.41

8. Appendix Table 24: Indirect (Raw to Raw to Scale) Equipercentile Crosswalk Table - From FACIT-Fatigue to PROMIS Fatigue (PROMIS Study). Note: Table 22 is recommended.

FACIT-F Score	Equipercentile PROMIS Scaled Score Equivalents (No Smoothing)	Equipercentile Equivalents with Postsmoothing (Less Smoothing)	Equipercentile Equivalents with Postsmoothing (More Smoothing)
52	30	30	29
51	36	35	35
50	38	38	38
49	40	40	40
48	42	42	42
47	44	44	44
46	45	45	45
45	46	46	46
44	48	48	47
43	48	48	48
42	50	50	49
41	50	50	50
40	51	51	51
39	52	52	52
38	52	52	53
37	53	53	54
36	54	54	54
35	55	55	55
34	56	56	56
33	56	56	56
32	57	57	57
31	57	57	57
30	58	58	58
29	58	58	58
28	59	59	59
27	59	59	60
26	60	60	60
25	61	61	61
24	61	61	61
23	62	62	62
22	63	63	63
21	64	63	63
20	64	64	64
19	65	64	64
18	65	65	65
17	66	66	66
16	66	66	66
15	66	66	66
14	66	67	67
13	67	67	68

FACIT-F Score	Equipercntile PROMIS Scaled Score Equivalentents (No Smoothing)	Equipercntile Equivalentents with Postsmoothing (Less Smoothing)	Equipercntile Equivalentents with Postsmoothing (More Smoothing)
12	68	68	68
11	69	69	69
10	70	70	70
9	70	71	71
8	72	72	72
7	73	73	72
6	73	74	73
5	74	75	74
4	74	76	76
3	84	77	77
2	84	79	79
1	84	82	82
0	88	86	86