1 Online Computerized Adaptive Tests (CAT) of Children's Vocabulary Development in
2 English and Mexican Spanish

3 George Kachergis¹, Virginia A. Marchman¹, Philip S. Dale², Jessica Mankewitz¹, &
4 Michael C. Frank¹

5 ¹ Stanford University

6 ² University of New Mexico

7 Author Note

8 Department of Psychology Department of Psychology

9 Correspondence concerning this article should be addressed to George Kachergis,
10 Building 420, 450 Jane Stanford Way, Stanford, CA, 94305. E-mail:
11 kachergis@stanford.edu

Abstract

Purpose: Measuring the growth of young children’s vocabulary is important for researchers seeking to understand language learning as well as for clinicians aiming to identify early deficits. The MacArthur-Bates Communicative Development Inventories (CDI) are parent-report instruments that offer a reliable and valid method for measuring early productive and receptive vocabulary across a number of languages. CDI forms typically include hundreds of words, however, and so the burden of completion is significant. We address this limitation by building on previous work using Item Response Theory (IRT) models to create Computer Adaptive Test (CAT) versions of the CDI. We created CDI-CATs for both comprehension and production vocabulary, for both American English and Mexican Spanish.

Method: Using a dataset of 7,633 English-speaking children ages 12-36 months and 1,692 Spanish-speaking children ages 12-30 months, across three CDI forms (Words & Gestures, Words & Sentences, and CDI-III), we found that a 2-parameter logistic IRT model fits well for a majority of the 680 pooled vocabulary items. We conducted CAT simulations on this dataset, assessing simulated tests of varying length (25-400 items).

Results: Even very short CATs recovered participant abilities very well with little bias across ages. An empirical validation study with N=204 children ages 15-36 months showed a correlation of $r=0.92$ between language ability estimated from full CDI vs. CDI-CAT forms.

Conclusions: We provide our item bank along with fitted parameters and other details, offer recommendations for how to construct CDI-CATs in new languages, and suggest when this type of assessment may or may not be appropriate.

Keywords: computerized adaptive test, vocabulary development, early language learning

Word count: 7759

Online Computerized Adaptive Tests (CAT) of Children's Vocabulary Development in English and Mexican Spanish

Introduction

Measuring children's early language is important for caregivers, clinicians and researchers. The MacArthur-Bates Communicative Development Inventories (CDIs; Fenson et al., 2007) are a set of parent report forms that offer a holistic assessment of children's productive and receptive language skills. CDIs are low-cost to administer and produce reliable and valid estimates of early vocabulary and other aspects of early language (Fenson et al., 1994). The CDIs offer more comprehensive data than a short interaction in the lab with a child (Fenson et al., 2000), because they ask caregivers to report on vocabulary comprehension as well as production, and other milestones, such as communicative gesture use and use of word combinations. In addition, the CDIs are less contextually influenced than language samples, as caregivers are integrating across the full range of their experiences with their child. Vocabulary size is assessed via a checklist format, which allows caregivers to quickly scan and recognize words their child produces or understands, rather than relying on recall alone. Because of these properties, CDI forms have been adapted to dozens of languages. Data from CDIs are archived in a central public repository (Wordbank; Frank, Braginsky, Yurovsky, & Marchman, 2017), and insights from these data have been used to inform theories of early language learning (Frank, Braginsky, Yurovsky, & Marchman, 2021).

Although CDIs measure a variety of other constructs related to early language, our focus here is on vocabulary assessment. Across languages, measures of vocabulary on the CDIs are very tightly correlated with other aspects of early language like gesture and grammatical competence (Bates et al., 1994; Frank et al., 2021). From the perspective of the CDI, it is justifiable to say that the language system is "tightly woven" (Frank et al., 2021), meaning that precise measures of early vocabulary provide a good proxy

measurement of the language system as a whole.

In American English and Mexican Spanish (two of the original languages in which the CDIs were developed), there are two long-form CDI instruments that focus on different ages. The 396-item vocabulary checklist on the CDI: Words & Gestures (CDI:WG) was designed for children ages 8–16 months and measures both comprehension and production. The 680-item vocabulary checklist on the CDI: Words & Sentences (CDI:WS) targets children ages 16–30 months, and includes nearly all of the items from the CDI:WG form, but only measures production. In addition, the English CDI-III contains a 100-item checklist meant for children 30–37 months of age, extending the age range for which CDI data are available to 3 years of age. All CDI forms include words from a range of semantic and syntactic categories. For example, the CDI:WS form comprises 22 semantic categories representing common early-learned words, including nouns (subdivided into e.g., Body Parts, Toys, and Clothing), action words (e.g. verbs), descriptive words (e.g. adjectives), and closed-class words such as pronouns.

Although the CDIs have many advantages, one clear drawback is the long length of the vocabulary checklists, which make them time-consuming for caregivers to complete. The length of the CDIs also makes it difficult to include them in studies requiring a battery of tasks, as is often the case in both clinical and research settings, including the case of assessing dual language learners (requiring administration of multiple CDIs). Due to these challenges, there have been a variety of efforts to create shortened versions of the CDI. The 100-item short-form CDIs (Fenson et al., 2007, 2000) are derived from the 680-item CDI:WS form, with items selected based on difficulty and item-to-full score correlations, while also attempting to represent the diversity of semantic and linguistic categories. While the scores of the short-form CDI:WS are highly correlated with scores on the full CDI:WS, there is evidence for a ceiling effect for children older than 27 or 28 months of age. Like the full forms, short-forms that are fixed (i.e., the same for all test takers) are inefficient, in that there are likely to be many words which are substantially below a given

child's level, and hence all will be checked, as well as many words which are substantially above the child's level, and hence none will be checked. The most informative items are ones around the child's expanding frontier, and that set will vary from child to child.

Computerized Adaptive Testing (CAT; Linden & Glas, 2010) versions of the CDI offer an alternative approach to creating short forms that might remove ceiling (and floor) effects and also hold the potential of further reducing the number of items caregivers must assess. CAT is a technique that allows test questions to be chosen adaptively based on the learner's responses, analogous to staircasing procedures to find stimuli (e.g., in audiological testing) that are just above/below a given participant's threshold. CATs are widely used in education as a way to assess individuals across a broad range of abilities in an efficient and precise manner (e.g., Weiss & Kingsbury, 1984; Weiss, 2004). The basis of CAT is item-response theory modeling (IRT; Embretson & Reise, 2013), a technique for the analysis of test data that allows the inference of both the ability of individual test takers and the difficulty (and other information) of individual test questions along shared and standardized dimensions. CAT models use this item information – typically extracted from a larger dataset collected via standard testing methods – to select questions of the appropriate difficulty for a particular test-taker. An individual CAT includes a number of components, including the bank of possible items and their difficulties, as well as an algorithm that uses the responses received thus far to choose the next item to give to a test taker, and a rule for when to stop (e.g., after a fixed number of items or after a desired precision has been reached). Given the established concurrent validity of the full CDIs (Fenson et al., 2007), correlations, especially empirical ones, between the CAT and the original CDIs provide strong evidence for the validity of the CAT.

Previous work has applied CAT and related techniques to CDI forms, leveraging the availability of large datasets from previous CDI studies where caregivers filled out the full forms (Chai, Lo, & Mayor, 2020; Cobo-Lewis et al., 2016; Makransky, Dale, Havmose, & Bleses, 2016; Mayor & Mani, 2019). For example, Makransky et al. (2016) used IRT

models fit to normative data from the CDI:WS to develop CAT versions. They conducted a simulation study comparing full scores on the CDI to different fixed-length CAT versions (with 5–400 items). Scores from a CAT of only 50 items had a correlation of $r = 0.95$ with the full CDI. However, for the youngest age group (16–18 month-olds), the correlation with the full CDI was somewhat lower ($r = 0.87$).

Following this work but using a slightly different approach, Mayor and Mani (2019) assessed a range of short forms created via random sampling of words but then scored the form for each child using a procedure that used individual item responses to identify the closest age- and sex-matched participants from the Wordbank database. This method performed well in simulations using pre-existing data from English, German, and Norwegian WS forms, yielding correlations with full CDI scores in excess of .94 for even a 25-item test. In addition, this study empirically validated the approach with a small group of caregivers (19 and 25 for the 25- and 50-item tests, respectively) and found good correlations between the short forms and the full CDI scores for German ($r = .96$ and $r = .94$). However, it should be noted that these correlations are not directly comparable to those recovered by the earlier CAT study, because the Mayor and Mani (2019) method makes use of age and sex information, which is on its own highly diagnostic of vocabulary size ($r = 0.78$ for predicting English WS production data with just age and sex).¹

Most recently, Chai et al. (2020) built on the method of Mayor and Mani (2019) to

¹ The use of age and sex information to predict vocabulary is both a strength and a weakness. It is a strength because it generally allows precise estimation with very limited assessment, based on the strong prior expectations that we can have about the vocabulary of a child of a particular age and sex. This precision can be very helpful for research. On the other hand, the use of prior information of this type in assessment can also be a weakness for individual assessment (especially in the clinical context) because it biases our view of a particular individual based on their demographic characteristics. A child with an unusually low or high vocabulary for their age and sex would tend to be assessed as closer to the typical score using this instrument. Another way of expressing this issue is that, on average, the use of demographic information trades some bias in exchange for lower variance.

create an efficient and less data-hungry adaptive technique. They used IRT models and standard CAT testing to sample informative items (as opposed to sampling them randomly) but then extrapolate to full CDI scores by using the prior method of finding age- and sex-matched participants whose item responses most closely match those of the current test taker. The addition of the IRT-based item selection further boosted performance in simulated CATs, leading to correlations greater than .95 for tests with 25 items and sometimes even fewer across American English, Danish, Mandarin, and Italian (with the latter two having substantially smaller datasets).

In sum, previous work strongly suggests that short CAT versions of CDI WS forms achieve strong performance in real-data simulation, and Mayor and Mani (2019) have provided initial validation of these results with a small sample of German speakers. Our goal in the current work is to build on this foundation and to create and validate a maximally versatile set of CATs for English and Spanish. These are the two most widely spoken languages in the US, and have particularly extensive datasets available. In particular, our contributions are:

1. To assess the fit of a variety of IRT models, determining the best IRT models to serve as a foundation for CDI-CATs.
2. To create and assess CDI-CATs for productive vocabulary in both American English and Mexican Spanish that cover the full age span covered by CDIs – 12 to 36 months of age.
3. To create and assess CDI-CATs for receptive vocabulary in these two languages, which have not to our knowledge previously been available.
4. To provide evidence for the validity of the productive vocabulary assessment using a larger-scale, web-based concurrent validation study, and
5. To provide access to the CDI-CATs through Web-CDI, an online platform for data collection using CDI-type instruments (DeMayo et al., 2021).

We first introduce the candidate Item Response Theory (IRT) models, how they were fitted to the datasets, and present an assessment of which model is appropriate as the basis for the adaptive CDI. Next, various CAT simulations were conducted to assess the effects of different design choices on test performance, and a set of preferred CAT parameters were chosen. Finally, we report results from a validation study. We end by discussing the strengths and weaknesses of our approach, and its potential value to clinicians, caregivers, and researchers to reduce the amount of time needed to assess children’s early language ability, making it possible to include language assessment in a broader array of studies.

Methods

Item Response Theory Models

A range of IRT models for diverse types of testing scenarios and item types exist (see Baker, 2001 for an overview); we considered four standard models of increasing complexity that may be appropriate for the dichotomous responses that caregivers make for each item (word) regarding whether their child can produce that word. The four models we considered (1-parameter logistic (1PL), 2PL, 3PL, and 4PL) are all widely-used, and build on each other, each introducing an additional item-level parameter. These models and their parameters are introduced sequentially in detail below, but in general the 1PL and 2PL are commonly used for many test scenarios with binary responses, while the 3PL and 4PL models tend to be more suitable for tests with higher rates of guessing (e.g., multiple-choice tests) and lapses (e.g., incorrect responding due to complex questions).

In all four models, the goal is to jointly estimate for each child j a latent ability θ_j , and for each item i one or more parameters capturing characteristics of the item, e.g. difficulty b_i , contained in all four models. In the Rasch (aka 1-parameter logistic; 1PL) model, where only item difficulty is modeled, the probability of child j producing a given item i is

$$P_i(x_i = 1|b_i, \theta_j) = \frac{1}{1 + e^{-D(\theta_j - b_i)}}$$

188 where D is a scaling parameter (fixed at 1.702) to make the logistic more closely
 189 match the ogive function in traditional factor analysis (Chalmers, 2012; Reckase, 2009).
 190 Thus, children with high latent ability (θ) will be more likely to produce any given item
 191 than children with lower latent ability, and more difficult items will be produced by fewer
 192 children (at any given θ) than easier items. The 2-parameter logistic (2PL) model
 193 additionally estimates a second parameter per item i , discrimination a_i , modifying the
 194 slope of the logistic (which is 1 in the 1PL model):

$$P_i(x_i = 1|b_i, a_i, \theta_j) = \frac{1}{1 + e^{-Da_i(\theta_j - b_i)}}$$

195 Items with higher discrimination (i.e. slopes) better distinguish children above/below
 196 that item's difficulty level. The 3-parameter logistic (3PL) model adds a pseudo-guessing
 197 parameter c_i , an asymptotic minimum probability for producing item i :

$$P_i(x_i = 1|b_i, a_i, c_i, \theta_j) = c_i + \frac{1 + c_i}{1 + e^{-Da_i(\theta_j - b_i)}}$$

198 The 3PL model may be appropriate if caregivers “guess” with a non-zero probability
 199 that their child knows a given word. Finally, the 4PL model adds an asymptotic upper
 200 bound d_i for producing item i , providing an upper limit on the probability of correct
 201 response item to item i :

$$P_i(x_i = 1|b_i, a_i, c_i, d_i, \theta_j) = c_i + \frac{d_i - c_i}{1 + e^{-Da_i(\theta_j - b_i)}}$$

202 representing a probability of randomly responding incorrectly. We used model
 203 comparisons to determine which of these four IRT models (1PL, 2PL, 3PL, or 4PL) best
 204 explain the CDI data before conducting CAT simulations.

Datasets

We report analyses for four datasets from Wordbank (Frank et al., 2017): comprehension data from children learning American English or Mexican Spanish (CDI:WG vocabulary checklists), and production data from children learning American English or Mexican Spanish (combining CDI:WG and CDI:WS vocabulary checklists).

Participants. The English comprehension dataset consists of Wordbank data from 2394 children (12-18 months of age) from the American English CDI:WG. The Spanish comprehension dataset consists of Wordbank data from 759 children (12-18 months of age) from the Mexican Spanish CDI:WG. All participating children’s home language environments were essentially monolingual, with at least 80% of their exposed language being in the target language. Data from Spanish-speaking participants was from the CDI norming sample in Mexico. Figure 1A and 1B shows children’s CDI:WG scores vs. age for the English and Spanish comprehension datasets, respectively.

The English production dataset consists of the combined production data from Wordbank for 2394 children aged 12-18 months from the American English CDI:WG, data from 5573 children aged 16-30 months from the English CDI:WS, and data from 69 children aged 31-36 months from the English CDI-III, for a total of 7633 participants. Note that these data are composed of the norming dataset (Fenson et al., 2007), as well as other contributions². The Spanish production dataset consists of the combined production data for 759 children from the Mexican Spanish CDI:WG (aged 12 or more months) as well as data from 1092 children from the Spanish CDI:WS form (16-30 months), for a total of 1691 participants. Figure 1C and 1D shows children’s production scores vs. age for the English and Spanish production datasets, respectively. Note that the socioeconomic distributions of these datasets are not matched. (See Frank et al. (2021) for a discussion of possible effects of socioeconomic status on vocabulary development.)

² <http://wordbank.stanford.edu/contributors>

Figure 1. Children’s CDI vocabulary plotted by age and sex in each dataset: (A) English comprehension, (B) Spanish comprehension, (C) English production, and (D) Spanish production. Note that we plot fitted quadratics showing extrapolated vocabulary sizes beyond the maximum CDI score, rather than an asymptotic logistic.

Instruments. The 680-item vocabulary checklist of the English CDI:WS form is organized into 22 semantic categories (e.g., furniture, games and routines, people). The Spanish CDI:WS vocabulary checklist consists of 680 words, organized into 23 semantic categories. The English CDI:WG vocabulary checklist comprises 396 of the easier vocabulary items from the English CDI:WS, and the Spanish CDI:WG is a subset of 428 of the easier items from the Spanish CDI:WS.³ Of the 100 English CDI-III items, 45 are from the English CDI:WS, and the remaining 55 words were selected to be more appropriate for older children (ages 30-36 months; Fenson et al., 2007). The 100 items of the Spanish CDI-III, however, overlaps with only 12 items of the Spanish CDI:WS.

When either of the CDI:WG forms was administered, caregivers were asked to indicate for each vocabulary item whether their child 1) understands that word (“comprehends”) or 2) both understands and says (“produces”) that word. Leaving the item blank indicates that the child neither comprehends nor produces that word. When either of the CDI:WS forms was administered, caregivers were asked to indicate for each vocabulary item on the instrument whether or not their child can recognizably produce (say) the given word.

For the production datasets, “produces” responses were coded as 1 and all other responses were coded as 0. For the comprehension datasets, both “comprehends” responses (from WG forms) and “produces” responses (from all forms) were coded as 1, and all other responses were coded as 0. Thus, our datasets consisted of four dichotomous-valued response matrices of size N subjects $\times$ W words ([English, Spanish] $\times$ [production,

³ When the CDI is adapted to a new language, it is up to the researchers creating the adaptation to selectively translate items that are culturally appropriate, and to introduce new ones, or even to reorganize the semantic categories as appropriate for the new language. Broadly speaking, the full English and Spanish CDI forms have considerable overlap, and thus comparable diversity in their content. It should be noted that the semantic categories play no role in the CAT, which is based entirely on the relationship between individual words and estimated language ability.

comprehension])).

Procedure. Before constructing a CAT, the IRT model with the appropriate structure for the CDI data had to be selected. Thus, we fitted the four IRT models described above (Rasch/1PL, 2PL, 3PL, 4PL) to each of the four datasets (English and Spanish, comprehension and production) and performed model comparisons.

For the favored model, we then identified ill-fitting items, pruned them, and then used the pruned item bank in a variety of computerized adaptive test (CAT) simulations on Wordbank data, investigating how many items needed to be given to achieve reliable estimates of children’s word learning ability across the range of intended ages, and benchmarking CAT performance against randomly-selected baseline tests of the same length. Finally, based on these simulations we chose a set of preferred CDI-CAT settings and fitted parameters that are recommended to be used by CAT developers for English and Spanish tests of comprehension and production. After reporting CAT simulations for each dataset with these preferred settings, we conducted an empirical experiment to test the validity of the English production CDI-CAT.

Results

All models, simulations, and other materials are available on OSF⁴.

IRT Model Selection

For each dataset, we fitted each of the four standard psychometric models (1PL-4PL) and compared the fits using Akaike and Bayesian Information Criteria (AIC and BIC) to select models with the appropriate level of complexity (i.e. number of parameters) justified by the data. Model fits were obtained in R (R Core Team, 2019) using the **mirt** package

⁴ OSF repository: <https://osf.io/xdp73/>.

(Chalmers, 2012). The results of the model comparisons are shown in the Appendix (Tables A1-A4).

For three of the four datasets, the 2-parameter logistic (2PL) model was preferred by both model selection techniques. For the large English production dataset (7633 children), the 4PL was preferred by AIC, but the more conservative BIC preferred the 2PL. In light of the 2PL being preferred in most model comparisons, and for the sake of interpretability, we selected the 2PL model as the basis for our further analyses for all four datasets. Next, we examined the fitted item difficulty and discrimination parameters of each of the 2PL models to determine if there were any ill-fitting words that should be pruned from the item banks.

Pruning the Item Banks

IRT models assume that items are locally independent—that is, items in a test should not be related (replying “yes” to one item should not lead to a “yes” on another item). Local dependence between one or more items (i.e., a significant correlation between items after the effect of the underlying trait is controlled for) may indicate that the latent ability is not unidimensional, and can lead to anomalous estimates for both the item parameters and children’s ability. Thus, we chose to prune any items from use in the CATs that showed both strong local dependence (LD) with one or more other items (Cramer’s $V \geq 0.5$; Chen & Thissen, 1997) as well as poor fit to the 2PL model ($\chi^2_{df}, p < .001$; Stone & Zhang, 2003). For example, items with non-monotonic item response functions would show poor fit, since with increasing ability, parents should show equal/higher rates of responding “yes” for their child on any given item.

English Comprehension. For the English comprehension model, 18 items in the full 2PL model had poor fit: “animal”, “arm”, “baa baa”, “back”, “bad”, “bottle”, “chair”, “close”, “cute”, “diaper”, “don’t”, “happy”, “in”, “love”, “out”, “telephone”, “tonight”, and

“toothbrush”. Only one item showed strong LD (“grandma”), but this item did not show poor fit; thus no items were pruned from the English comprehension model.

Spanish Comprehension. For the Spanish comprehension model, 14 items in the full 2PL model had poor fit: “abeja”, “acerrín”, “ardilla”, “cabra”, “cuna”, “galleta”, “la”, “leche”, “oso”, “pelota”, “señora”, “sol”, “tijeras”, and “tortillitas”. However, the only item that showed strong LD was “tía”, which did not show ill fit; thus, no items were pruned from this model.

English Production. For the English production model, 142 items in the full 2PL model had poor fit (χ^2_{df} , $p < .001$; see Appendix B for a full listing). However, only one item showed strong LD: “daddy”. This item also showed poor fit, and was thus pruned from the data for simulated CATs, and the 2PL model used in the CATs was re-fitted without that item.

Spanish Production. For the Spanish production model, 38 items in the full 2PL model had poor fit (χ^2_{df} , $p < .001$; see Appendix B). Six items showed strong LD: “abuela”, “agua”, “mamá”, “no”, “prima”, and “tía”. Only “no” also showed poor fit, and thus only “no” was pruned from the data. The 2PL model was re-fitted without this item.

Data Imputation

Using the per-child, per-item responses, we simulated a variety of CATs with different settings. However, a final preprocessing step was needed for the production datasets: for participants’ whose data comes from either the CDI:WG form or from the CDI-III, it was necessary to impute their responses for any of the 680 CDI:WS items they were not asked. We did this imputation via the marginal maximum likelihood method (Bock & Aitkin, 1981) using the participants’ estimated production ability and the appropriate 2PL model. Overall, 11.40% of the data was missing from the English production dataset, and was imputed. 12.60% of the data was missing from the Spanish production dataset, and was

imputed.

CAT Simulations

We used real data simulations to estimate how well the CDI would perform were it administered with varying CAT procedures. That is, we used the four Wordbank datasets of full CDI administrations to simulate CATs, selecting the succession of items based on caregiver’s actual responses from the full CDI, and assuming that their responses to the subset of the items on the CAT would be the same as they had been on the full CDI. Thus, for each simulated CAT procedure, we derived an estimate of each child i ’s language ability (θ_i), similar to a latent factor score. These CAT-estimated abilities were then compared (e.g., by correlation, reliability, and error) to children’s IRT-estimated ability from the full CDI.⁵ We used this general approach to conduct three sets of CAT simulations, for each of the four datasets. The three simulations tested CATs with 1) a fixed number of items, 2) a threshold for early stopping, and 3) our preferred CAT settings, on the basis of the prior simulations.

Fixed-length CATs. Following the approach taken by Makransky et al. (2016), we first ran a series of fixed-length CAT simulations and compared children’s CAT-estimated ability to their IRT-estimated ability from their full CDI data. Each successive item (including the first item) was selected to be maximally informative of ability, given by the maximum a posteriori (MAP) estimate. For the production datasets, we varied test lengths from 25 to 400 items, while for the comprehension datasets we simulated CATs only up to 300 items due to the shorter length of the WG forms. We also added a comparison to a baseline test that selected a random set of CDI items for each child of the same length as the CAT for that child. For each test length, comparisons of children’s CAT-estimated

⁵ Note that we use the term “IRT-estimated ability from full CDI” to mean a score based on the IRT parameters for all items, whereas “full CDI vocabulary” refers to the conventional score of the number of positively checked items, which considers every item to be of equal difficulty (1).

ability to their full CDI IRT-estimated ability are shown in the Appendix (Tables C1-C4). Overall, results were quite good even for 25-item tests, with $r = .99$ for both production datasets, $r = .98$ for English comprehension, and $r = .97$ for Spanish comprehension. 50-item tests brought the production correlations close to 1, and the comprehension correlations to .99. Compared to the production correlations reported by Makranksy et al. (2016) for the same test lengths, these values are somewhat higher, likely due to the larger Wordbank datasets used here. However, we also found that the baseline of randomly-selected items also showed a strong association with ability from the full CDI ($r \sim .94$ for 25 items; $r \sim .97$ for 50 items). Nonetheless, the mean standard error of measurement (SEM) of ability estimates from the short CATs are roughly half that of the random tests (e.g., English production: 50-item CAT mean SEM = .13; 50-item random test mean SEM = .24). Even the longest tests only brought the mean SEM down to ~ 0.1 —and the longer random tests all have similar values, showing a diminishing advantage of longer CATs.

We also note that on the shorter CAT simulations, a large number of items were never selected. For example, on the 50-item production CATs, 347 of the 680 items were never selected for any of the 7633 English CATs, and a similar number of items were never selected for any of the 1692 simulated Spanish CATs. For comprehension, roughly 150 of the WG items were never selected for any of the simulated 50-item CATs. These items were not useful for estimating the ability of these samples of children either because they were too easy (known by nearly all children, e.g. “ball”), too difficult (known by very few children, e.g. “would”), or because an item with similar difficulty and higher discrimination was selected instead.

Early-stopping CAT Simulations. Having ascertained that fixed-length forms as short as 25 to 50 items can estimate children’s language ability with high reliability, we next evaluated how well early-terminating CATs operate. We chose to evaluate CATs that stop after reaching either a maximum item length, or an estimated standard error of

measurement (SEM) for a given child's ability reaches 0.10, a quite strict criterion. Note from the fixed-length CAT simulations that on average, SEM only approaches this criterion, but lower SEM for many children will allow for earlier termination for many and thus potentially a much shorter average administration. For each participant in each dataset, we simulated CATs using a maximum of 25, 50, 75, 100, 200, 300, or (for production datasets only) 400 items, with early termination if the estimated SEM reaches 0.10 for a given subject. The results of these CAT simulations are shown in the Appendix (Tables C5-C8).

The simulation results for English and Spanish comprehension (Tables C5 and C6) were broadly similar, with reliability for maximum 50-item tests at 0.981 (English) and 0.979 (Spanish), and improving by $< .007$ for even the longest (300-item) tests.

The results for English production (Table C7) show that reliability for the 25-item early stopping CAT is the same as for the fixed-length CAT (0.975), which is unsurprising since the CAT terminated early for very few participants (mean items = 24.7). For the 50-item CAT, the participants received a median of 45 items (mean = 39.5), and reliability was quite high (0.981). Allowing for longer CATs (maximum 75-400 items) showed minimal gains in reliability (0.983 to 0.985), and while the mean length of the CATs steadily increased (58.7 to 143.1), the median number of items administered remained constant at 45.

The simulation results for Spanish production (Table C8) were largely consistent with those for English: participants rarely received fewer than 25 items, and the median number of items administered on tests that allowed 50 or more items stayed constant at 36 items. Tests with a maximum of 50 items were already quite reliable (0.975), and longer tests did not greatly increase reliability (0.976 to 0.979).

On balance, for all datasets it seems that tests with a minimum of 25 items are needed to approach an estimated SEM of 0.10 for many participants, and that allowing up to 50 items in total can sufficiently improve the reliability of the CAT without much

increasing the mean/median test length. Tests with more than 50 items show only small improvements on reliability that seem unwarranted given the increasing demands on administration time. We use these insights to specify and test our preferred settings for English and Spanish production and comprehension CATs.

Preferred CAT Settings. Based on the fixed-length and early-terminating CAT simulations, it seems clear that even a minimum of 25 items can often lead to a stable estimate of children’s language ability, but that in some cases it is helpful to have the flexibility to have slightly longer tests of up to 50 items to further reduce estimation error. Thus, we chose to test – first via simulation, and then in an empirical validation of English production – CATs with a minimum of 25 items, a maximum of 50 items, and early termination when the estimated $SE \leq .15$.⁶

We further compare two ways of estimating subject’s ability from the set of actual CAT responses: *maximum likelihood* (ML) returns the ability value that is most likely based only on the observed data, whereas the *maximum a posteriori* (MAP) estimate differs slightly in that it incorporates a prior expectation that ability is drawn from a normal distribution. Since MAP estimates are regularized (i.e., extreme values are pulled in slightly towards the higher-probability center) in comparison to ML estimates, MAP estimates may be more stable and perhaps less extreme-valued in the tails.

Finally, we tested two ways of choosing the starting item: we either used the maximally informative item, as in all prior simulations, or we used age-based starting items that were chosen to be slightly easy relative to each subject’s age i (i.e., with difficulty $d_i \approx \theta_i - .5$, where θ_i is the average ability of children of that age). We evaluated using starting items that were somewhat easy for children of a given age so that most caregivers will be presented with at least one item that their child is likely to know. If the maximally

⁶ In the early-stopping CAT simulations, we found that the stricter $SE \leq .10$ criterion was never met on average, even for tests of 300 or more items.

informative start item across all ages were used, caregivers of younger and/or lower-than-average ability children may be presented predominantly with words that their child does not know, which may concern them. Thus, if an age-based starting item does not impact the CAT's performance, we will recommend their use to mitigate this concern and to grant most caregivers the opportunity to use all response alternatives.

Tables 1-4 show the results of the CAT simulations for each dataset with the preferred settings. MAP ability estimates resulted in lower mean SE, slightly higher correlations with ability estimates from the full CDI, and also resulted in an average of 1 fewer item being asked per CAT. MAP ability estimates were strongly correlated with ability estimates from the full CDI, with $r = .99$ for both English and Spanish production, and $r = .98$ or higher for comprehension. On average, 32-34 items were asked on the simulated CATs (median = 25 for all four datasets), showing that for most children the CATs terminated much before the 50-item maximum. Finally, the use of easy age-based starting items had no discernible impact on the reliability of the CAT-estimated ability in comparison to using the maximally informative starting item. Tables C9-C12 in the Appendix show the correlation between ability estimates for these preferred CAT settings and ability estimates from the full CDI broken down by age (3-month age bins for production; 2-month age bins for comprehension), which revealed consistently high correlations ($r_s = [.96, .99]$) across English and Spanish for both production and comprehension.

In summary, these simulation results suggest that a CAT with 25-50 items, terminating when the estimated SEM $\leq .15$, using MAP-based ability estimation and age-based starting items will allow for efficient and reliable estimation of young children's ability for both language comprehension and production, in both English and Spanish.

Table 1

English production CAT simulations with preferred settings.

Scoring / Start Item	Mean Items	r with full CDI	Mean SE	Reliability
ML / MI	32.3	0.990	0.157	0.975
MAP / MI	31.9	0.992	0.146	0.979
ML / age-based	32.2	0.990	0.157	0.975
MAP / age-based	31.9	0.992	0.146	0.979

Table 2

Spanish production CAT simulations with preferred settings.

Scoring / Start Item	Mean Items	r with full CDI	Mean SE	Reliability
ML / MI	34.6	0.982	0.193	0.963
MAP / MI	34.3	0.990	0.168	0.972
ML / age-based	34.6	0.982	0.192	0.963
MAP / age-based	34.4	0.990	0.168	0.972

Table 3

English comprehension CAT simulations with preferred settings.

Scoring / Start Item	Mean Items	r with full CDI	Mean SE	Reliability
ML / MI	33.9	0.983	0.165	0.973
MAP / MI	33.0	0.985	0.160	0.974
ML / age-based	34.0	0.983	0.165	0.973
MAP / age-based	33.1	0.985	0.160	0.974

Table 4

Spanish comprehension CAT simulations with preferred settings.

Scoring / Start Item	Mean Items	r with full CDI	Mean SE	Reliability
ML / MI	34.7	0.974	0.187	0.965
MAP / MI	34.1	0.977	0.175	0.969
ML / age-based	34.8	0.974	0.187	0.965
MAP / age-based	34.3	0.977	0.175	0.969

Validation Study

To empirically test the validity of the model as a short version of the CDI vocabulary checklist, we administered both the full MB-CDI:WS form and the English production CDI-CAT to the caregivers of English-speaking children.

Methods

Participants. We recruited caregivers of 251 children between 15 and 36 months of age from the online survey vendor Prolific (child demographics: 22 Latino/Hispanic (48 not reported), 203 white, 29 black, 13 Asian, 4 Native American, and 4 other race/ethnicity). The average number of years of education attained by the primary caregiver was 15.85 (sd=2.23). 47 participants were excluded according to our preregistered exclusion criteria.⁷ The exclusion criteria included discrepancies between questions about the child's age (n = 4), child's sex (n = 4), child's birth weight (n = 14), caregiver's zipcode (n = 1), caregiver's year of birth (n = 13), amount of exposure to languages other than English in the home (n = 2), or multiple violations of any of the previous exclusion criteria (n = 9). Data from the remaining 204 (103 female) children were analyzed.

⁷ Preregistration: <https://aspredicted.org/h2kz4.pdf>

Procedure. Caregivers were asked to complete both the full CDI:WS and a production CDI-CAT alongside a comprehensive demographics form. Both surveys and the demographics form were administered within the Web-CDI interface (DeMayo et al., 2021). In Web-CDI, caregivers receive a link that takes them to each form to complete. Instructions are provided in written and pictorial form. The user interface of the CDI-CAT differs from that of the full Web-CDI. In the full CDI, each page presents all of the items in each of the semantic categories, ranging from 6 items (connecting words) to 103 items (action words). In the CDI-CAT, caregivers are presented with one word at a time and asked to select “Yes” or “No” to indicate whether their child can say that word (see Figure 2). For items that are ambiguous when presented in isolation, information about the type of word was included in parentheses to help caregivers identify the target word (e.g., “drink (verb)” vs. “drink (object)”). The order of survey administration was counterbalanced so $n = 106$ caregivers received the full CDI followed by the CDI-CAT and demographics while $n = 98$ participants received the CDI-CAT and demographics followed by the full CDI:WS. This was done to enable analysis of possible carryover effects from one version of the assessment to the other. Caregivers were asked to complete both surveys in one sitting, but 2.45% of caregivers ($n = 5$) returned their second survey within 24 hours and 3.43% ($n = 7$) returned their second survey within 5 days.

Figure 2. Screenshots of the English production CDI-CAT user interface.

Results and Discussion

Median time for caregivers to complete the full CDI was 19.8 minutes, and median time to complete the CAT was 5.1 minutes. We carried out our preregistered⁸ comparisons of children’s estimated latent language ability (θ) from the CDI-CAT to their IRT-estimated ability from the full CDI:WS, as well as to their total vocabulary sumscore from the full CDI:WS. Figure 3 shows participating children’s IRT-estimated ability scores from the full CDI:WS vs. their ability scores from the CDI-CAT, which were strongly correlated ($r = 0.92$), for both females and males. The correlation between participants’ total vocabulary sumscore from the full CDI:WS vs. their CAT-estimated ability was $r = 0.86$. For reference, participants’ total vocabulary from the full CDI:WS was also strongly correlated with their IRT-estimated ability from the full CDI:WS ($r = 0.95$). To test whether there was an effect of test administration order, we conducted a t-test on the squared differences between children’s CAT-based ability (θ) and their IRT-estimated ability from the full CDI:WS, and found no significant difference in the amount of squared error as a function of test order ($t(185.8) = 0.46$, $p = 0.65$).

The mean squared error between CDI-CAT ability and full CDI:WS ability was 0.55 (median = 0.17, sd = 1). Of the 18 participants who showed extreme discrepancy ($> \text{mean} + 1.5 \times \text{sd} = 2.05$; 11 of whom did the full CDI first), all yielded a much higher CDI-CAT ability as compared to their ability on the full CDI (mean CAT ability - full CDI ability = 1.75). Assuming that their responses on the full CDI were veridical, these participants generally responded “knows” to many more items on the CDI-CAT than expected. It is not clear if these caregivers were too quick to endorse their child’s knowledge of items on the CDI-CAT, or perhaps neglected to attend sufficiently to each of the items on the full CDI, where a default response indicates lack of knowledge. To attempt to calibrate caregivers, it may be advisable to stress in the instructions for a CDI-CAT that the test is intended to

⁸ Preregistration: <https://aspredicted.org/h2kz4.pdf>

be difficult, and that they should expect to be asked about many words (at least 50%) that their child is unlikely to know yet.

Nonetheless, the results largely validate that the English production CDI-CAT works well for measuring children’s early language ability across a broad age range. Table 5 shows the correlations between children’s ability from the CDI-CAT vs. the full CDI by age group, indicating generally strong associations across the entire age range. The 27-30 month age bin shows a slightly weaker association, but with small sample sizes (<30) for some age bins this may be due to sampling variability. Figure D1 in the Appendix shows that the relationships between children’s age and 1) their total vocabulary from the full CDI:WS (left), 2) their full CDI ability (middle), and 3) their ability measured from the CDI-CAT (right) are all of approximately the same strength, and show the modest female advantage commonly found in assessments of young children’s language ability (Eriksson et al., 2012; see Frank et al., 2021, Ch. 6 for an overview).

Figure 3. Children’s estimated ability from the full CDI:WS vs. estimated ability from the CDI-CAT, by sex of child.

Table 5

Validation study ability correlations (CDI-CAT vs. full CDI) by age group.

	[15,18)	[18,21)	[21,24)	[24,27)	[27,30)	[30,33)	[33,36]
r CAT vs. full CDI	0.96	0.88	0.9	0.87	0.68	0.86	0.89
N	26	22	26	30	28	24	48

General Discussion

We set out to build a computerized adaptive version of the CDI (CDI-CAT) to assess language ability in young children by asking caregivers 25-50 vocabulary questions – much shorter than the 680 vocabulary items on the full CDI:WS. After finding that our preferred CAT settings work well in simulation for estimating the language ability of thousands of children from Wordbank (Frank et al., 2017) in American English and Mexican Spanish, for both production and comprehension, we ran a validation study to establish the empirical reliability of the CDI-CAT for measuring English production ability. The validation study found a strong correlation between ability estimated by the CDI-CAT and ability estimated from the full CDI:WS, and this correlation was strong across age ranges and sex.

Building on earlier work using IRT-based models to construct CAT measures of vocabulary (Chai et al., 2020; Makransky et al., 2016; Mayor & Mani, 2019), we have improved on a few aspects. First, we have leveraged larger datasets than prior work by stitching together production Wordbank data (Frank et al., 2017, 2021) from multiple CDIs, allowing for more precise, stable item parameter estimates that we expect will generalize well to future samples across a broad age range. The use of such large datasets also enabled us to conduct model comparisons of several standard IRT models (1PL-4PL), which established that the 2-parameter logistic (2PL) IRT model achieves the right balance of complexity and fit for all four of our datasets, fitting per-item difficulty and

discrimination (slope) parameters. As the 2PL model works best for both languages investigated here, we expect that the 2PL model will also operate best when extending this approach to build CDI-CATs for other languages. Another extension made in our work is the application of IRT models and CAT simulations to data regarding children's language comprehension. The simulation results were promising, suggesting that these comprehension CATs will likely function as well as those for production. Moreover, the proposed production CDI-CATs operate over a wider age range (12-36 months) than any of the individual CDI forms upon which they are based (CDI:WG, CDI:WS, or CDI-III). These production CDI-CATs should serve to robustly estimate language ability for outlier children at both extremes (i.e., young children with high ability, and older children with lower ability), and can largely mitigate ceiling and floor effects produced by static short-form CDIs (see Appendix F). The CDI-CAT may also be an excellent choice for longitudinal studies: due to the small number of items given per administration and its adaptive nature, as children age they will be unlikely to see the same item more than once, although longitudinal use warrants additional testing. Finally, we have conducted an empirical study to test the validity of the preferred CAT English production settings which is larger than has been run in prior studies, with encouraging results. In summary, we have shown how to build CDI-CATs of only 25-50 items using 2PL IRT models for both comprehension and production in two languages (American English and Mexican Spanish), and offer these item banks and parameters⁹ for other researchers. We hope that a similar approach will be taken in the future to build other CAT-based vocabulary measures.

Clinical applications of the CDI-CAT

The brevity and high validity of the CDI-CAT make it highly appropriate for some clinical purposes. In the present validity study, the median completion time for the CDI-CAT (5.1 minutes) was 74% shorter than that for the full CDI:Words & Sentences

⁹ CAT item bank and parameters: <https://osf.io/xdp73/>

(19.8 minutes). The correlation between the two assessments (.86) was not far below the internal consistency (.96) and test-retest reliability coefficients (.95), as reported in Fenson et al. (2007), indicating that the concurrent validity is near the theoretical maximum possible.

Beyond simply reducing the need for caregivers' time, the brevity offers several additional advantages. Although the CDI-CAT has simple and user-friendly instructions, in those cases where assistance may be needed due to limited parental literacy or for other reasons, it is practical to provide one-to-one assistance, such as oral presentation. (As noted by Makransky et al. (2016), it may be useful to inform caregivers that the words have been chosen to be somewhat difficult for children at their child's age, and that it is expected that they will be making several "no" responses.) Although the CDI-CAT requires online access, it can easily be completed on a tablet with internet connection; the brevity suggests that it would be practical in the office of a pediatrician, educational psychologist, or speech-language pathologist. The brevity also enables completion of the CDI-CAT in two languages at little cost in time if one caregiver can do both (in other cases, different caregivers may be most appropriate for assessment in each language). At present, only American English and Mexican Spanish are available, but the range of instruments is likely to increase in the future, as the modeling is easily extendable to norming datasets in other languages.

The factors just discussed suggest that the CDI-CAT is very well-suited to serve as a screening instrument for late talkers. In addition to the standardized ability score analyzed here, it is also possible to generate percentile scores that provide information about a child's relative status in relation to a normative dataset, thereby, aiding in the interpretation of a child's score for clinical screening purposes. Another advantage for that purpose is that it is not necessary to choose among the various levels of CDI instruments (CDI:WG, CDI:WS, CDI-III, short forms Level 1 and Level 2), as the single CDI-CAT covers the entire range from 12-36 months. The CDI-CAT also has considerable potential

for longitudinal administration, as there is likely to be little repetition of specific items from one administration. Consequently it is appropriate for “watch and see” situations, where the rate of change over time is a relevant clinical measure. Finally, the variation across administrations minimizes the possibility of “teaching to the test” by parents/caregivers, teachers, or therapists.

The clinical limitations of the CDI-CAT should also be kept in mind. It assesses only vocabulary, not morphology, syntax, phonology or pragmatics. Given the substantial correlation between vocabulary and these other dimensions, it is a reasonable first choice for assessment, but it is far from a broad assessment and may miss children whose difficulties lie elsewhere. Furthermore, it assesses only vocabulary size, not composition. Thus it is a tool for identification, not diagnosis. Another limitation is that due to the variable set of words queried, without any alignment across languages, it is not possible to compute a total conceptual vocabulary score for bilingual children; consequently it can be difficult to interpret the performance of children who score somewhat low in both languages. Different items will be queried pre- and post-test, and thus there will be variability in overlap between intervention target and CAT items both between participants and between testing times, and this would be a confounder for intervention effect (Makransky et al., 2016). Finally, as noted earlier, it may not be appropriate as a measure of intervention effectiveness, particularly when vocabulary is a core focus of the intervention. This concern may be less relevant for interventions with other foci, such as grammar or pragmatics. In the future, it may be possible for intervention designers to review the full CDI vocabulary checklist to identify intervention-relevant words, and mark them as not to be used in the CDI-CAT.

The clinical advantages and limitations of the CDI-CAT listed above generally apply to research applications as well. Of particular relevance for research is the suitability of the CDI-CAT for repeated administration in longitudinal research, although this must be qualified by the lack of information obtained about the composition of vocabulary. Second,

unlike paper-based forms which can simultaneously ask caregivers about children's comprehension and production of each item, the current CDI-CATs assess comprehension and production separately, requiring two sessions if both are needed. Future work may consider fitting a multidimensional IRT model to estimate production and comprehension difficulties for each item, and testing a multidimensional CAT to jointly assess production and comprehension abilities simultaneously (e.g., Groen, Eggen, & Veldkamp, 2016).

The IRT parameters in these CDI-CATs have been fitted on monolingual children, and it is an open question whether the items function the same way for bilingual or multilingual children. Pandemic- and funding-related constraints meant it was only possible to empirically validate the English production CAT. But the similarity of the development approach for the CAT in the two languages, and the excellent results of the simulations suggest that the CATs provide a valid approach for obtaining a vocabulary measure in each language separately when that is the goal. However, further investigation is needed to explore the clinical utility of using the CATs to help screen for language delay in dual-language learners and other situations in which it is desired to combine the measures across languages for a given child. It is likely that an age-typical ability estimate from the CAT in either of the languages indicates lack of impairment, and clinically low ability in both languages, such as a score below the 15th percentile, suggests language impairment. However, further research is needed on the interpretation of low, but not clinical range scores in both languages as it is not possible to compute total conceptual vocabulary. Despite these caveats, we are encouraged by the results of both the simulations and the empirical validation study, and endorse the application of the approach taken here to create CATs for other languages.

In summary, assessment of early vocabulary skill is an important part of many studies of early development. Here, we have developed four computerized adaptive tests to quickly measure young children's language comprehension and production ability in American English or Mexican Spanish, and empirically validated this approach for English

production. We believe that these tests – and brethren in other languages created following the same procedure – will greatly ease the time and financial burdens of including language assessments in batteries of tasks investigating children’s early development.

Conflict of Interest Statement

No potential conflict of interest was reported by the authors.

Funding Statement

This research was funded in part by the MacArthur-Bates CDI Advisory Board.

References

- Baker, F. B. (2001). *The basics of item response theory*. ERIC.
- Bates, E., Marchman, V., Thal, D., Fenson, L., Dale, P., Reznick, J. S., . . . Hartung, J. (1994). Developmental and stylistic variation in the composition of early vocabulary. *Journal of Child Language*, 21(1), 85–123.
- Bock, R. D., & Aitkin, M. (1981). Marginal maximum likelihood estimation of item parameters: Application of an em algorithm. *Psychometrika*, 46(4), 443–459.
- Chai, J. H., Lo, C. H., & Mayor, J. (2020). A Bayesian-inspired item response theory-based framework to produce very short versions of MacArthur-Bates Communicative Development Inventories. *Journal of Speech, Language, and Hearing Research*, 63(10), 3488–3500.
- Chalmers, R. P. (2012). mirt: A multidimensional item response theory package for the R environment. *Journal of Statistical Software*, 48(6), 1–29.
<https://doi.org/10.18637/jss.v048.i06>
- Chen, W.-H., & Thissen, D. (1997). Local dependence indexes for item pairs using item response theory. *Journal of Educational and Behavioral Statistics*, 22(3), 265–289.
- Cobo-Lewis, A. B., Meadow, C., Markowsky, G., Pearson, B. Z., Collier, S. A., & Eilers, R. E. (2016). Computerized adaptive assessment of infant-toddler language development: Demonstration and validation of an app for screening.
- DeMayo, B., Kellier, D., Braginsky, M., Bergmann, C., Hendriks, C., Rowland, C. F., . . . Marchman, V. (2021). Web-CDI: A system for online administration of the MacArthur-Bates Communicative Development Inventories.
- Embretson, S. E., & Reise, S. P. (2013). *Item response theory*. Psychology Press.
- Eriksson, M., Marschik, P. B., Tulviste, T., Almgren, M., Pérez Pereira, M., Wehberg, S., . . . Gallego, C. (2012). Differences between girls and boys in emerging language

skills: Evidence from 10 language communities. *British Journal of Developmental Psychology*, 30(2), 326–343.

Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D. J., Pethick, S. J., . . . Stiles, J. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development*, i–185.

Fenson, L., Marchman, V. A., Thal, D. J., Dale, P. S., Reznick, J. S., & Bates, E. (2007). *MacArthur-Bates Communicative Development Inventories: User's guide and technical manual (2nd ed.)*. Baltimore, MD: Brookes.

Fenson, L., Pethick, S., Renda, C., Cox, J. L., Dale, P. S., & Reznick, J. S. (2000). Short-form versions of the MacArthur Communicative Development Inventories, 21, 95–116.

Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2017). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, 44(3), 677.

Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2021). *Variability and consistency in early language learning: The wordbank project*. MIT Press.

Groen, M. M. van, Eggen, T. J., & Veldkamp, B. P. (2016). Multidimensional computerized adaptive testing for classifying examinees with within-dimensionality. *Applied Psychological Measurement*, 40(6), 387–404.

Linden, W. J. van der, & Glas, C. A. (2010). *Elements of adaptive testing*. Springer.

Makransky, G., Dale, P. S., Havmose, P., & Bleses, D. (2016). An item response theory-based, computerized adaptive testing version of the MacArthur-Bates Communicative Development Inventory: Words & Sentences (CDI:WS). *Journal of Speech, Language, and Hearing Research*, 59(2), 281–289.

Mayor, J., & Mani, N. (2019). A short version of the MacArthur-Bates Communicative

699 Development Inventories with high validity. *Behavior Research Methods*, 51(5),
700 2248–2255.

701 R Core Team. (2019). *R: A language and environment for statistical computing*. Vienna,
702 Austria: R Foundation for Statistical Computing. Retrieved from
703 <https://www.R-project.org/>

704 Reckase, M. D. (2009). Multidimensional item response theory models. In
705 *Multidimensional item response theory* (pp. 79–112). Springer.

706 Stone, C. A., & Zhang, B. (2003). Assessing goodness of fit of item response theory
707 models: A comparison of traditional and alternative procedures. *Journal of*
708 *Educational Measurement*, 40(4), 331–352.

709 Weiss, D. J. (2004). Computerized adaptive testing for effective and efficient measurement
710 in counseling and education. *Measurement and Evaluation in Counseling and*
711 *Development*, 37(2), 70–84.

712 Weiss, D. J., & Kingsbury, G. G. (1984). Application of computerized adaptive testing to
713 educational problems. *Journal of Educational Measurement*, 21(4), 361–375.

Appendix A

IRT Model Comparisons

For each dataset, we fitted each of the four standard psychometric models (1PL-4PL) and compared the fits using Akaike and Bayesian Information Criteria (AIC and BIC) to select models with the appropriate level of complexity (i.e. number of parameters) justified by the data. Tables A1-A4 show the results of these model comparisons.

English Comprehension. Table A1 shows the model comparison for the 2394 children in the English comprehension data. Both AIC and BIC prefer the 2PL model over either the simpler Rasch (1PL) model or the more complex 3PL or 4PL models. Note that the 4PL model converged very quickly, yet with a worse overall fit than the other models: this remained the case even after setting a much lower tolerance for convergence. Phil Chalmers, the author of the mirt R package, indicates that the 4PL model has too many parameters (and thus flexibility) unless the dataset is in excess of 5000 participants.

Table A1

Comparison of 1PL-4PL models for English comprehension.

Model	AIC	BIC	logLik	df
Rasch	718,890.44	721,185.38	-359,048.22	NA
2PL	708,467.70	713,046.03	-353,441.85	395.00
3PL	708,909.35	715,776.85	-353,266.68	396.00
4PL	709,382.34	718,539.00	-353,107.17	396.00

Spanish Comprehension. Table A2 shows the model comparison for the 759 children in the Spanish comprehension data. The 2PL model is preferred over both the Rasch (1PL) model and the 3PL model. As for the English comprehension dataset, the

728 4PL again converged quickly, but with worse fit than the 3PL, likely due to the small size
 729 of the dataset relative to the number of parameters.

Table A2

*Comparison of 1PL-4PL models for Spanish
 comprehension.*

Model	AIC	BIC	logLik	df
Rasch	238,610.98	240,598.10	-118,876.49	NA
2PL	236,007.27	239,972.27	-117,147.64	427.00
3PL	236,590.80	242,538.29	-117,011.40	428.00
4PL	254,849.39	262,779.38	-125,712.70	428.00

730 **English Production.** Table A3 shows the model comparison for the 7633 children
 731 in the English production data. For both model selection criteria, smaller values are better.
 732 The 2PL model is preferred over the Rasch (1PL) model by both AIC and BIC, indicating
 733 that the 2PL's addition of a discrimination parameter (slope) for each item is justified.
 734 The 3PL model is preferred over the 2PL by AIC, but not by the more conservative BIC
 735 criterion. AIC again prefers the more complex 4PL model over the simpler 3PL model, but
 736 according to BIC the 2PL model has the appropriate level of complexity for the dataset.

Table A3

Comparison of 1PL-4PL models for English production.

Model	AIC	BIC	logLik	df
Rasch	2,542,533.33	2,547,259.63	-1,270,585.66	NA
2PL	2,475,247.30	2,484,686.02	-1,236,263.65	679.00
3PL	2,472,824.63	2,486,961.89	-1,234,375.31	677.00
4PL	2,466,697.98	2,485,547.67	-1,230,632.99	679.00

Spanish Production. Table A4 shows the model comparison for the 1610 children

in the Spanish production data. The 2PL model has the lowest AIC value, as well as the

lowest BIC value, indicating that it is of the appropriate complexity for this dataset.

Table A4

Comparison of 1PL-4PL models for Spanish production.

Model	AIC	BIC	logLik	df
Rasch	576,073.90	579,740.40	-287,355.95	NA
2PL	559,215.47	566,537.69	-278,247.73	679.00
3PL	560,716.88	571,700.22	-278,318.44	680.00
4PL	612,960.38	627,604.84	-303,760.19	680.00

Appendix B

Ill-fitting Items

Below we show the full list of items with ill fit (Stone's χ^2_{df} , $p < .001$) in the 2PL model for Spanish production (38 items) and for English production (142 items).

Table B1

Ill-fitting Spanish production items

1	2	3	4	5
¡am!	éste	mesa	pies	televisión
a	gritar	mío	quién	tengo manita
adíos/byebye	guaguá	mirar	recámara	tortilla
árbol	huevo	nana	saber	tú
atole	jugar	no	shhh	y
boca	las	o	suya	ya
caja	manos arriba	oír	taco	
caliente	me	pelo	te (pronouns)	

Table B2

Ill-fitting English production items

1	2	3	4	5	6
all	corn	hand	muffin	school	this
all gone	couch	hear	my	see	throw
baby	cow	helicopter	nice	sheep	tickle
babysitter's name	daddy*	hello	noodles	shoe	to
ball	did/did ya	hi	nose	sing	tooth
bathroom	dirty	home	not	sister	towel
bathtub	don't	hungry	orange (description)	sled	trash
beans	donkey	I	ouch	sleep	tree
bee	donut	is	out	sleepy	tummy
blanket	dress (object)	it	owie/boo boo	snow	uh oh
boots	ear	jacket	pancake	sock	vroom
bottle	eye	jar	peas	spoon	water (not beverage)
brother	fish (food)	juice	pen	stick	what
brown	french fries	keys	penis*	stop	which
bump	friend	look	pickle	story	white
by	garage	love	pizza	stroller	who
bye	get	man	popcorn	stuck	woof woof
child	girl	me	potty	sweater	yes
choo choo	good	meat	purse	swing (action)	yogurt
church*	goose	mine	raisin	table	you
coat	grrr	mommy*	read	teddybear	yucky
coffee	gum	money	red	thank you	yum yum
coke	hair	motorcycle	ride	that	
cold	hamburger	mouth	rooster	there	

Appendix C

Simulated CATs

The following sections show the results of the simulated CATs using American English and Mexican Spanish data, conducted for both production (combined WG and WS data) and comprehension (WG data).

Fixed-length CAT Simulations

Following Makransky et al. (2016), we ran a series of fixed-length CAT simulations and compared children's estimated ability (θ) from these CATs to children's estimated ability from the full CDI. The results are shown in Tables C1-C4 for English comprehension (C1), Spanish comprehension (C2), English production (C3), and Spanish production (C4). The results were quite good even for 25- and 50-item tests, but note that we added a comparison to tests of randomly-selected questions (per subject), and found that ability estimates from these tests were also strongly correlated with thetas from the full CDI. The mean standard error of the random tests showed more of a difference, especially for shorter tests (e.g., <100 items).

Table C1

Fixed-length CAT simulations compared to full CDI for English comprehension

Test Length	r vs. full CDI	Mean SE	Reliability	r Random vs. full CDI	Random SEM
25	0.980	0.175	0.969	0.949	0.296
50	0.990	0.137	0.981	0.975	0.220
75	0.994	0.120	0.986	0.982	0.184
100	0.996	0.110	0.988	0.989	0.161
200	0.999	0.093	0.991	0.996	0.118
300	1.000	0.087	0.992	0.999	0.097

Table C2

Fixed-length CAT simulations compared to full CDI for Spanish comprehension

Test Length	r vs. full CDI	Mean SE	Reliability	r Random vs. full CDI	Random SEM
25	0.970	0.192	0.963	0.935	0.317
50	0.983	0.152	0.977	0.962	0.238
75	0.989	0.135	0.982	0.975	0.202
100	0.992	0.124	0.985	0.982	0.180
200	0.998	0.105	0.989	0.994	0.133
300	1.000	0.098	0.990	0.998	0.112

Table C3

Fixed-length CAT simulations compared to full CDI for English production.

Test Length	r vs. full CDI	Mean SE	Reliability	r Random vs. full CDI	Random SEM
25	0.990	0.157	0.975	0.953	0.308
50	0.995	0.126	0.984	0.971	0.242
75	0.996	0.113	0.987	0.980	0.209
100	0.997	0.106	0.989	0.985	0.188
200	0.999	0.095	0.991	0.993	0.143
300	0.999	0.091	0.992	0.996	0.122
400	1.000	0.088	0.992	0.998	0.108

Table C4

Fixed-length CAT simulations compared to full CDI for Spanish production.

Test Length	r vs. full CDI	Mean SE	Reliability	r Random vs. full CDI	Random SEM
25	0.987	0.182	0.967	0.934	0.329
50	0.992	0.151	0.977	0.958	0.265
75	0.994	0.138	0.981	0.972	0.233
100	0.996	0.130	0.983	0.978	0.213
200	0.998	0.117	0.986	0.991	0.166
300	0.999	0.111	0.988	0.994	0.144
400	0.999	0.109	0.988	0.997	0.129

Early-stopping CAT Simulations

For each participant in each dataset, we simulated CATs using a maximum of 25, 50, 75, 100, 200, 300, or (for production datasets only) 400 items, with early termination if the estimated standard error of measurement (SEM) reached 0.10 for a given subject. For each of these simulations, we report 1) the median and mean number of items used, 2) the correlation of ability scores (θ) estimated from the CAT with IRT-estimated ability scores applied to the full CDI (*r with full CDI*), and 4) the mean estimated SEM achieved by the CATs, 5) the reliability of the ability estimates from the CAT compared to the ability estimated by the full CDI, and 6) how many items were never selected (*Unused Items*). The results are shown in Tables C5-C9.

Table C5
Early-stopping CAT simulations for English comprehension.

Max. Items	Median Items	Mean Items	r with full CDI	Mean SE	Reliability	Unused Items
25	25	25.0	0.980	0.175	0.969	235
50	50	50.0	0.990	0.137	0.981	145
75	75	73.5	0.994	0.121	0.985	76
100	82	84.5	0.995	0.116	0.986	52
200	82	107.2	0.996	0.113	0.987	0
300	82	125.0	0.996	0.112	0.987	0

Table C6

Early-stopping CAT simulations for Spanish comprehension.

Max. Items	Median Items	Mean Items	r with full CDI	Mean SE	Reliability	Unused Items
25	25	25.0	0.970	0.192	0.963	276
50	50	50.0	0.983	0.152	0.977	172
75	75	70.4	0.988	0.137	0.981	114
100	78	81.6	0.990	0.132	0.983	83
200	78	113.1	0.993	0.125	0.984	2
300	78	138.3	0.994	0.123	0.985	0

Table C7

Early-stopping CAT simulations for English production.

Max. Items	Median Items	Mean Items	r with full CDI	Mean SE	Reliability	Unused Items
25	25	24.7	0.990	0.158	0.975	468
50	45	39.5	0.993	0.136	0.981	405
75	45	49.9	0.994	0.131	0.983	375
100	45	58.7	0.995	0.128	0.983	344
200	45	88.9	0.995	0.126	0.984	183
300	45	116.4	0.995	0.125	0.984	87
400	45	143.1	0.995	0.124	0.985	15

Table C8

Early-stopping CAT simulations for Spanish production.

Max. Items	Median Items	Mean Items	r with full CDI	Mean SE	Reliability	Unused Items
25	25	25.0	0.987	0.182	0.967	496
50	45	40.3	0.991	0.160	0.974	403
75	45	51.6	0.992	0.154	0.976	352
100	45	61.9	0.993	0.150	0.977	308
200	45	98.6	0.994	0.145	0.979	176
300	45	133.0	0.994	0.143	0.980	80
400	45	166.3	0.994	0.142	0.980	12

CAT Performance Across Age Groups

For the preferred settings, it is desirable to know if the CAT can be expected to perform well across the intended age range. Thus, for each dataset we examined correlations between the simulated preferred CAT's ability estimates to the IRT-estimated ability from the full CDI.

English Comprehension. Table C9 shows correlations between ability estimates from the full CDI compared to the estimated ability from the preferred CAT by age group (126 8-9 month-olds, 233 10-11 mos, 909 12-13 mos, 233 14-15 mos, and 893 16-18 mos.

Spanish Comprehension. Table C10 shows correlations between ability estimates from the full CDI compared to ability from the preferred CAT by age group (104 8-9 month-olds, 138 10-11 mos, 172 12-13 mos, 137 14-15 mos, and 208 16-18 mos.

English Production. Table C11 shows correlations between ability from the full CDI compared to ability from the preferred CAT split by age (1035 12-15 month-olds, 2156

Table C9

Correlation between the preferred CAT’s ability estimates and the full CDI for English comprehension by age group.

Scoring / Start Item	[8,10)	[10,12)	[12,14)	[14,16)	[16,18]
ML / MI	0.963	0.976	0.974	0.975	0.975
MAP / MI	0.983	0.977	0.974	0.976	0.977
ML / age-based	0.965	0.975	0.974	0.974	0.975
MAP / age-based	0.982	0.976	0.973	0.975	0.977

778 15-18 mos, 1128 18-21 mos, 665 21-24 mos, 1095 24-27 mos, 1190 27-30 mos, 322 30-33
779 mos, and 42 33-36 mos).

780 **Spanish Production.** Table C12 shows correlations between ability from the full
781 CDI compared to ability from the preferred CAT split by age (248 12-15 month-olds, 318
782 15-18 mos, 307 18-21 mos, 227 21-24 mos, 219 24-27 mos, and 291 27-30 mos).

Table C10

Correlation between the preferred CAT's ability estimates and the full CDI for Spanish comprehension by age group.

Scoring / Start Item	[8,10)	[10,12)	[12,14)	[14,16)	[16,18]
ML / MI	0.958	0.97	0.962	0.978	0.961
MAP / MI	0.964	0.978	0.963	0.978	0.961
ML / age-based	0.957	0.971	0.963	0.979	0.961
MAP / age-based	0.965	0.978	0.963	0.979	0.961

Table C11

Correlation between the preferred CAT's ability estimates and the full CDI for English production by age group.

Scoring / Start Item	[12,15)	[15,18)	[18,21)	[21,24)	[24,27)	[27,30)	[30,33)	[33,36]
ML / MI	0.966	0.975	0.978	0.984	0.977	0.971	0.943	0.934
MAP / MI	0.985	0.979	0.98	0.985	0.977	0.971	0.944	0.939
ML / age-based	0.966	0.976	0.978	0.985	0.977	0.973	0.955	0.919
MAP / age-based	0.985	0.98	0.979	0.985	0.977	0.973	0.957	0.94

Table C12

Correlation between the preferred CAT's ability estimates and the full CDI for Spanish production by age group.

Scoring / Start Item	[12,15)	[15,18)	[18,21)	[21,24)	[24,27)	[27,30]
ML / MI	0.912	0.971	0.975	0.974	0.976	0.969
MAP / MI	0.972	0.981	0.978	0.977	0.978	0.97
ML / age-based	0.918	0.971	0.974	0.974	0.977	0.97
MAP / age-based	0.973	0.981	0.977	0.977	0.979	0.972

Appendix D

Ability vs. Age

For participants in the validation study, Figure D1 shows the relationships between children's age and 1) their total vocabulary from the full CDI:WS (left), 2) their full CDI ability (middle), and 3) their ability measured from the CDI-CAT (right). These relationships are all of approximately the same strength, and show the expected female advantage.

Figure D1. Children's age vs. their full CDI vocabulary size, full CDI ability, and CAT-estimated ability, by sex of child.

Appendix E

Standard Error of Measurement by Ability

Finally, we examine the standard error of measurement (SEM) for children of varying ability using the preferred CAT settings with age-based starting items. The SEM theoretically varies with ability, as the CDI is most informative of children's ability for children who are of closer to average ability. This is because there are more items of intermediate difficulty than very difficult or very easy items, and so the test is better calibrated to distinguish children of intermediate ability than children at high/low ranges of ability. Figure E1 shows SEM vs. estimated ability (θ) for each child in the Wordbank data using the preferred CDI-CAT settings. The CDI-CATs show fairly minimal SEM for children with ability $-1 \leq \theta \leq 2$, meaning that the CDI-CATs will give a quite precise estimate of ability for $\sim 82\%$ of children. For children with estimated $\theta < -2.5$ or $\theta > 3$ (i.e. $< 2\%$ of the population) the SEM starts to rise, and it may be advisable to administer a full CDI if it is important to get a precise estimate of the child's language ability.

Figure E1. Standard Error of Measurement (SEM) vs. children's language ability (theta) using the preferred CDI-CAT settings, with ML-based estimation and age-based starting items).

Figure F1. Children’s CDI-CAT estimated language ability (theta) vs. IRT-estimated ability from the full CDI. The only children shown are the 308 who would have gotten a floor or ceiling score on one or more 100-item short-form CDIs. On the CDI-CAT, most of these children no longer show a ceiling or floor effect, and their CAT-estimated ability is strongly related to the IRT-estimated ability from the full CDI.

Finally, we investigated whether the CDI-CAT would mitigate floor and ceiling effects that would be present if participants in the Wordbank English production dataset (N=7633) were given existing 100-item short-form CDIs from Fenson et al. (2000). In our dataset, 213 children would receive a floor score on one or both of the short-form production CDI:WS forms (i.e., be reported by their caregiver to produce 0 of the 100 items), and 95 children would receive a ceiling score on one or both of the short-form production CDIs (i.e., be reported to produce all of the 100 items). Figure F1 shows the estimated ability for these 308 children as estimated by the production CDI-CAT with preferred settings (MAP-based estimation with age-based start items) vs. their

IRT-estimated ability from the full CDI. For these children who would have scored at floor or ceiling on the short-form CDIs, only 18 scored at floor on the CDI-CAT, and 1 scored at ceiling on the CDI-CAT. Overall, the CDI-CAT thetas vs. the full CDI thetas for these children were strongly related ($r = .999$, $t(306) = 337.9$, $p < .001$). Thus, we can conclude that the CDI-CAT will mitigate ceiling and floor effects for many children (289 out of 308 in this sample) that would be found if the 100-item short-form CDIs were used – and with at most half the number of tested items.