



Everyday language input and production in 1,001 children from six continents

Elika Bergelson^{a,1} , Melanie Soderstrom^b , Iris-Corinna Schwarz^{c,d} , Caroline F. Rowland^{e,f,g} , Nairán Ramírez-Esparza^h, Lisa R. Hamrickⁱ , Ellen Marklund^c, Marina Kalashnikova^{j,k} , Ava Guez^l , Marisa Casillas^{e,g,m} , Lucia Benettiⁿ , Petra van Alphen^o, and Alejandrina Cristia^{1,1} 

Edited by Jenny Saffran, University of Wisconsin, Madison, WI; received January 20, 2023; accepted October 14, 2023, by Editorial Board Member Susan A. Gelman

Language is a universal human ability, acquired readily by young children, who otherwise struggle with many basics of survival. And yet, language ability is variable across individuals. Naturalistic and experimental observations suggest that children's linguistic skills vary with factors like socioeconomic status and children's gender. But which factors really influence children's day-to-day language use? Here, we leverage speech technology in a big-data approach to report on a unique cross-cultural and diverse data set: >2,500 d-long, child-centered audio-recordings of 1,001 2- to 48-mo-olds from 12 countries spanning six continents across urban, farmer-forager, and subsistence-farming contexts. As expected, age and language-relevant clinical risks and diagnoses predicted how much speech (and speech-like vocalization) children produced. Critically, so too did adult talk in children's environments: Children who heard more talk from adults produced more speech. In contrast to previous conclusions based on more limited sampling methods and a different set of language proxies, socioeconomic status (operationalized as maternal education) was not significantly associated with children's productions over the first 4 y of life, and neither were gender or multilingualism. These findings from large-scale naturalistic data advance our understanding of which factors are robust predictors of variability in the speech behaviors of young learners in a wide range of everyday contexts.

human diversity | language | socioeconomic status | speech | infancy

Typically developing children readily progress from coos to complex sentences within just a few years, leading some to hypothesize that the universal language abilities of humans develop uniformly, with only incidental effects of individual- or group-level variation (1). And yet, studies using a variety of proxies for language development find some evidence of such variation in early language skills, with differences reported between girls and boys (2) as well as those raised in socioeconomically privileged compared to disadvantaged households (3, 4).

However interesting, these studies tend to rely on Western-centric samples and methods and may not reflect everyday language use in children. Moreover, prior work often stops after only considering individual predictors in a binary way (i.e., do they significantly impact language development or not), while failing to ask the more informative question of how large their relative impact is (5), especially in freely occurring, everyday speech behavior.

Recent research on mice and whales shows the promise of machine learning for examining everyday animal behavior (6, 7). We leverage advances in wearables and machine-learning-based speech technology to catalyze a similar breakthrough in language development research. Our dataset is composed of >40,000 h of audio from >2,500 d in the lives of 1,001 2- to 48-mo-olds from six continents and diverse cultural contexts (Fig. 1). Within this dataset, we focused on the amount of speech or speech-like vocalization young children produce in their everyday life. Critically, these automatically extractable “quantity” measures correlate robustly with gold-standard “quality” measures of children's language skills and knowledge, like vocabulary estimates (*SI Appendix, section 1D* for relevant evidence) (4).

We query and compare the effects of two types of factors. First, there are factors with undeniable effects on early language production, namely, child age and language-relevant clinical risks and diagnoses. Second, there are individual- and family-level factors that are reported to correlate with variability in early language skills: socioeconomic status (SES; operationalized here as maternal education; *SI Appendix, section 2B*), gender, language input quantity, and multilingual background. Because small and homogeneous samples make universal claims more questionable,

Significance

Harnessing a global sample of >40,000 h of child-centered audio capturing young children's home environment, we measured contributors to how much speech 0- to 4-y-olds naturally produce. Amount of adult talk, age, and normative development were the sole significant predictors; child gender, socioeconomic status, and multilingualism did not explain how often children vocalized or how much adult talk they heard. These findings (strengthened by our validation of existing automated speech algorithms) open up interesting conversations regarding early language development to the broader public, including parents, clinicians, educators, and policymakers. The factors explaining variance also inform our understanding of humans' unique capacity for learning and potentially large-scale applications of machine technology to everyday human behavior.

The authors declare no competing interest.

This article is a PNAS Direct Submission. J.S. is a guest editor invited by the Editorial Board.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Although PNAS asks authors to adhere to United Nations naming conventions for maps (<https://www.un.org/geospatial/mapsgeo>), our policy is to publish maps as provided by the authors.

¹To whom correspondence may be addressed. Email: elika_bergelson@fas.harvard.edu or alecristia@gmail.com.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2300671120/-/DCSupplemental>.

Published December 12, 2023.

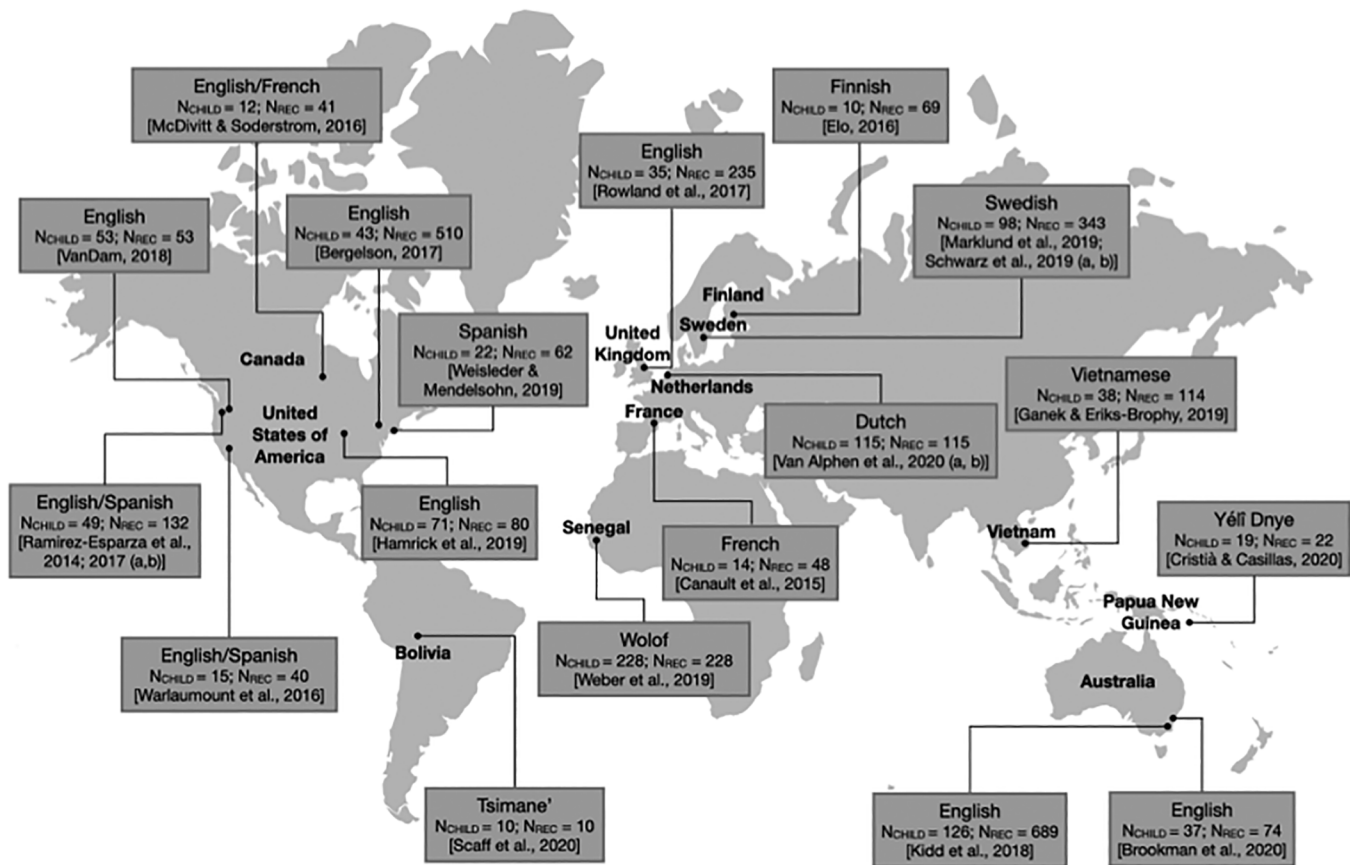


Fig. 1. Geographical location, primary language, number of children (N_{CHILD}), number of recordings (N_{REC}), and data citation for each corpus.

a key contribution of this work is its benchmarking of the level of stability and variability of everyday language use in a heterogeneous, richly diverse participant sample.*

Measuring Diverse, Real-life Language Use. Language skills and knowledge are not directly observable. As a result, all studies use a proxy when estimating them in individual children. These proxies have variable validity and predictive power relative to other measures, both concurrently and predictively, and likely vary in the extent to which they reflect children’s everyday language behavior. For instance, parental report measures are indirect and—especially for receptive knowledge—can be difficult for caretakers to estimate (9), even in relatively homogeneous Western-centric contexts.

Here, we adopt a very different approach. We employed the LENA™ system, which captures what children hear and say across an entire day through small wearable recorders (10); this ecologically valid sampling method reduces observer effects relative to, e.g., shorter video recordings (11). The LENA™ system uses standardized algorithms that estimate who is speaking when, alongside automated counts of adult and child linguistic vocalizations (4) (see definition and validation in *SI Appendix, sections 1C–1E*). The resulting LENA™ measures correlate with and predict other measures of language skills in children with and without clinical risks or diagnoses, as revealed by manual transcriptions, clinical instruments, and parent questionnaires

*While these data collectively span living circumstances, geography, and family structure, some data donors were concerned that highlighting differences when minoritized communities are involved poses ethical challenges, in terms of honorable representation and potential harm. Individual data stewards are actively engaging in richer descriptions of included samples (*SI Appendix, section 5*), which may enable future work on meaningful population-level differences (e.g., ref. 8).

(12, 13). We use LENA™’s validated, automated estimates to derive our measures of everyday language use: adult talk and child speech (see detailed motivation in *SI Appendix, section 3B*). We define child speech as the quantity of children’s speech-related vocalizations (e.g., protophones (14), babbles, syllables, words, or sentences, but not laughing or crying) per hour, and adult talk as the number of near and clear vocalizations per hour attributed to adults (both as detected by LENA™’s algorithm; see *Methods*). Assuaging concerns that these measures are merely capturing chattiness or repetition, both have a ≥ 0.7 correlation with measures of lexical diversity and language “quality”: Our child speech measure correlates with vocabulary in an independent sample, and the adult talk measure correlates with the number of word types from manual transcription in a subset of the data (*SI Appendix, section 1D*).

Capitalizing on this standardized and deidentified numeric output, we solicited LENA™ datasets that researchers had previously collected to study mono- and multilingual children (i.e. those learning >1 language) in urban, farmer-forager, and subsistence-farming contexts worldwide (Fig. 1). This resulted in a dataset reflecting the state of current knowledge in ecologically valid speech samples from children’s daily lives (*SI Appendix, section 3A*; see *Methods* for more sample details).

The dataset includes children from wide-ranging SES backgrounds, based on maternal education levels spanning from no formal education to advanced degrees (*SI Appendix, section 2B*). This SES proxy was selected not only because it was available in all 18 corpora (only 3 had alternative SES proxies), but most importantly because it is the most commonly employed SES proxy in language acquisition research, as established in meta-analyses (15, 16). This allows our findings to inform ongoing

discussions. Theories of how SES relates to children's language development have proposed a wide range of pathways in which maternal education is predictive of children's language experiences, including the connection between maternal education and the tendency to employ verbal over physical responsiveness (17), the diversity in mothers' vocabulary (18), and the frequency of verbally rich activities (19). Maternal education also correlates highly with other SES proxies [e.g., $r = 0.86$ in a study of children growing up in 10 European or North American countries, (20)], suggesting it may also indirectly pick up on other pathways linking SES to language development, through, e.g., differential access to resources and nutrition, or exposure to stress perinatally (21). At the same time, we recognize that comparing a variable like education across countries, although commonly done (22), is not straightforward. Therefore, we supplement our preregistered approach with numerous exploratory checks and analyses examining alternative implementations (*SI Appendix, sections 3G and 3H* described further below).

Crucially, by including children aged 2 to 48 mo, we span a wide range of linguistic skills, allowing us to better capture the effects of our variables over a broad span of development within our socioculturally and geographically broad-ranging participants. We also include children with a variety of diagnoses of language delays and disorders, as well as those at high risks for them (*Methods* and *SI Appendix, section 2A* for definitions and detailed justification). Such children's language development is by definition nonnormative. Thus, age and nonnormative status provide useful yardsticks for considering the significance and effect size of other child- and family-level factors (SES through maternal education, child gender, mono- vs. multilingual status, and how much adults talk to and around the child). That is, if a factor (e.g., gender) has an effect far smaller than that of age or nonnormative development, it would suggest that individual differences within it are relatively limited in their connection to everyday language use. If the effects are comparable in size, it would instead suggest that the amount of speech humans produce in everyday interactions is undergirded by substantial and structured individual differences, rather than striking uniformity. Given that effects could vary as a function of child age, we make sure to include key interaction terms. For instance, if older children are more sensitive to adults' talking to them than younger ones, then we can expect age to interact with adult talk.

Results

Predicting Children's Speech Production. We employed a hypothesis-testing approach: In a two-step preregistration, we first established exploration and confirmation data subsets (see *Methods* and *SI Appendix, section 3A* for detailed explanation, and *SI Appendix, sections 3D and 3E* for the procedure used to derive preregistered hypotheses and analyses). We then leveraged the held-out confirmation subset to answer our key question: How well do specific individual- and family-level factors predict variation in how much speech young children produce? At stake in these analyses is whether systematic differences in children's lives have measurable links to their language production, and if so, what the strength of these relationships is both overall, and in relation to one another (see Table 1 for results[†]).

[†]All β s in tables and text are based on treatment-coded models. See *SI Appendix, section 3H* for sum-coded models, which give the same pattern of results.

Table 1. Model results predicting child speech

	β	SE	q	
Intercept	0.109	0.128	0.681	
Child gender (Male)	0.026	0.051	0.852	
SES(<H.S.(1))	0.001	0.111	0.991	
SES(H.S.(2))	-0.033	0.115	0.932	
SES(B.A.(4))	-0.064	0.079	0.681	
SES(>B.A.(5))	-0.002	0.090	0.991	
Control	-0.085	0.029	0.035	*
Norm	-0.220	0.087	0.036	*
Adult talk	0.260	0.037	<0.001	*
Age	0.647	0.024	<0.001	*
Mono	0.045	0.095	0.852	
Norm \times Adult talk	-0.005	0.063	0.991	
Norm \times Age	-0.217	0.051	<0.001	*
Adult Talk \times Age	0.125	0.022	<0.001	*
Adult Talk \times Mono	0.092	0.072	0.45	
Mono \times Age	-0.048	0.056	0.681	
Norm \times Adult talk \times Age	0.019	0.043	0.852	
Mono \times Adult talk \times Age	0.137	0.065	0.094	

q -values show FDR-corrected P -values.

Note. Betas show deviation from the following baseline levels: Child gender: female; SES: some university(3); Norm: Norm(ative development); Mono: Mono(lingual). SES = child SES based on maternal education (<H.S.(1) = less than high school, H.S.(2) = high school, B.A.(4) = college degree, >B.A.(5) = advanced degree); Control = overlap rate control; Adult talk = adult vocalization count rate.

As expected, we found that older children produced more speech than younger ones ($\beta = 0.647$, $SE = 0.024$). Children with nonnormative development produced less speech than children with normative development ($\beta = -0.22$, $SE = 0.087$),[‡] an effect that strengthened with age ($\beta = -0.217$, $SE = 0.051$; see Fig. 2*B*). This is expected because for some groups in our nonnormative subset (e.g., those with familial risk of a speech impairment), older children are more likely to have an actual diagnosis (as opposed to risk factor) than younger ones (see *SI Appendix, section 2A* for details on nonnormative classification).

Our results further revealed that young children's speech production correlated with the amount of adult talk they heard ($\beta = 0.26$, $SE = 0.037$). This correlation strengthened with age ($\beta = 0.125$, $SE = 0.022$; see Fig. 2*A*), perhaps because variation in adult talk rate has less effect on infants [whose early babbles occur frequently even when infants are alone, (14)]. The effect of adult talk is a substantial one. Taking the effects of age and normativity as convenient (but unrelated) gauges for what counts as a considerable effect, we see that the effect size of adult talk is about a third of that for age and similar to that for normativity (adult talk: 0.26; interaction adult talk by age: 0.125; age: 0.647; nonnormative development: -0.22; interaction nonnormative by age: -0.217; all effect size betas expressed as SDs).

To provide these results in more intuitive units, we fit the same model centering variables without scaling. Children produced 66 more vocalizations per hour with each year of life. For every 100 adult vocalizations per hour, children produced 27 more vocalizations; this effect grew by 16 vocalizations per year. Relative to infants with typical development, those with nonnormative development produced 20 fewer vocalizations per hour; this difference grew by 8 vocalizations per year.

Surprisingly, and in contrast to previous results using smaller and less diverse datasets and/or other language proxies, we found

[‡]The normativity estimate is negative because normative development is the baseline.

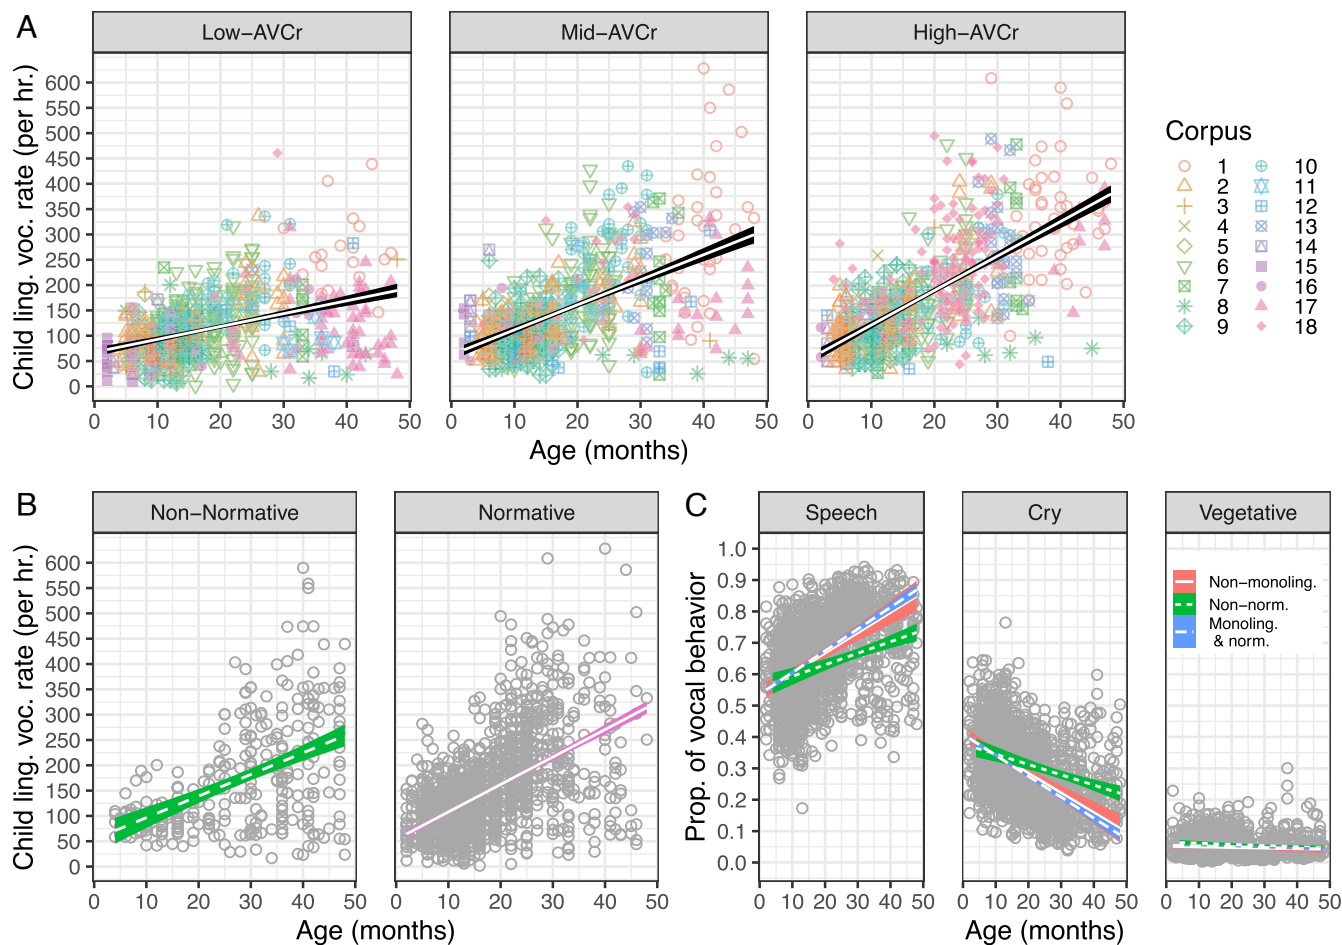


Fig. 2. Effects of adult talk, child age, and normative development on children's speech production. Points show each daylong recording; lines show linear regression with 95% CIs. Child speech is quantified as child linguistic vocalization rate; adult talk as adult vocalization count rate (AVCr). (A) Child speech by age, split by low/mid/high tertiles of adult talk. Lines depict significant adult talk \times age interaction. Color-shape combinations show each unique corpus, numbered to preserve anonymity. (B) Child speech by age and normative status. Lines depict significant age \times normative status interaction. (C) Proportion of vocal behavior classified as speech, cry, or vegetative, by age. The line type/color indicates monolingual and normative statuses. N.B. Monolingual normative CI (blue) falls fully within that for multilingual children (pink) for all three types of vocal behavior, highlighting these groups' equivalent patterns.

that child gender, SES (indexed here by maternal education), and monolingual status did not explain significant variation in child speech. As our raw data figures and model outcome results show, these null effects hold both when considering covariates (as in our model; Table 1) and when considering these variables individually (as in Fig. 3; *SI Appendix* sections 3F–3H). In our full model controlling for other variables (Table 1), the largest estimate for main effects or interactions involving child gender, SES, and monolingual status was about half of that for normativity, and one-sixth of that for age; none reached thresholds for statistical significance.

While our models are well powered to estimate associations of child speech with age, normativity, adult talk, gender, SES (as measured by maternal education), and monolingual status, this is predicated upon pooling the data and accounting statistically for corpus- and child-level variance via random effects, as described in *Methods*. This makes it beyond this paper's scope to analyze language or population/cultural differences in detail, i.e., in a way that might allow the consideration of additional, culture-specific variables (hence their omission in Figs. 2 and 3); see *SI Appendix*, section 5 for citations to research on individual datasets, some of which tackle such differences directly.

Noting that the results above have the strongest inferential value thanks to being preregistered, we also addressed certain

alternative hypotheses and interpretations that could have rendered our conclusions unjustified through a series of follow-up analyses. These checked for robustness of our key results with different operationalizations and statistical implementations of SES, when considering only children under or over 18 mo, when considering causal paths, and when incorporating speech from other children as a predictor; our key results held in all cases (*SI Appendix*, section 3H).

We highlight here the results that may run most counter to many readers' assumptions, namely, that in this large sample, SES (indexed by maternal education) does not come out as a significant predictor of child speech. This conclusion held when declaring SES as an ordinal and as a continuous variable based on levels or years of maternal education, when binarizing SES levels based on individual countries' average education completion rate and when declaring a random slope for SES within corpus (which allows SES effects to vary across corpora).

Some readers may wonder whether there were some corpora for which SES did matter. If so, the analysis with random SES slopes by corpus would have indicated this, but it did not (*SI Appendix*, section 3H). The relationship between SES and child speech was weak and inconsistent across corpora (as evident in Fig. 4).

Perhaps most convincingly, results also held when constraining our analysis to our largest homogeneous subset, the North

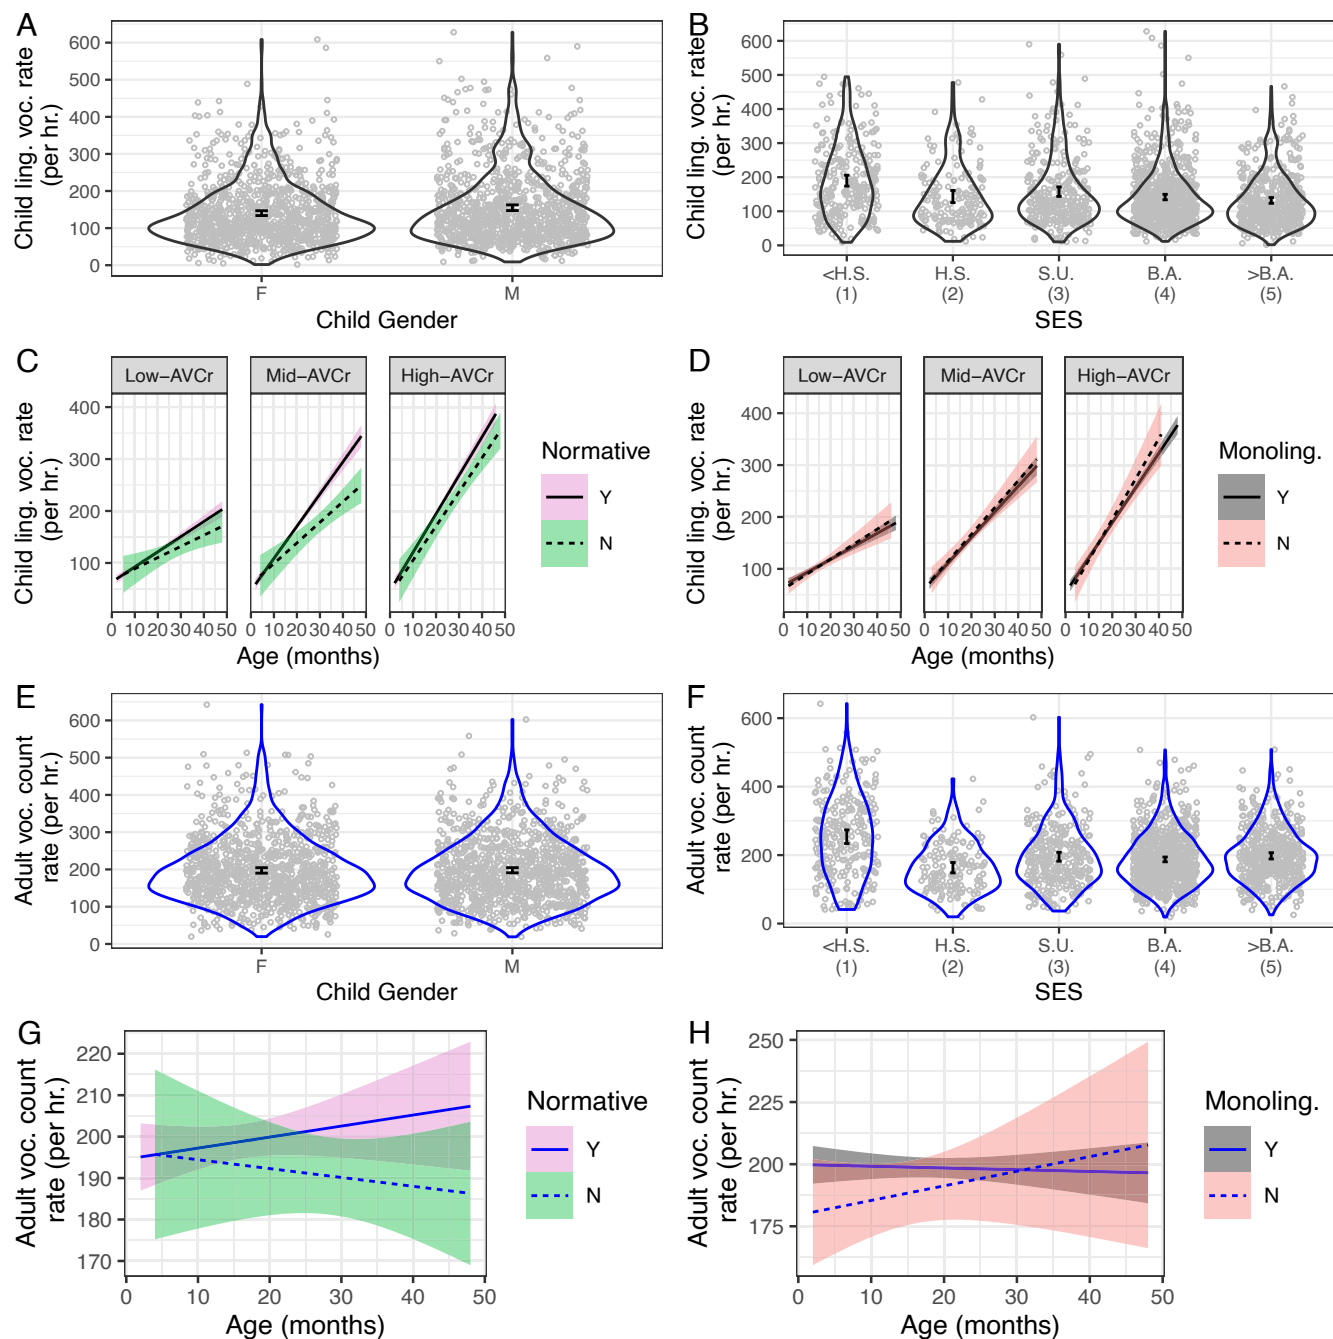


Fig. 3. Factors that do not predict child speech or adult talk. Points = individual recordings, jittered horizontally. Lines = linear fit with 95% confidence intervals. Error bars = 99% bootstrapped CIs of sample means. Child speech is quantified as child linguistic vocalization rate; adult talk as adult vocalization count rate (AVCr). **A & B:** null effects of child gender (A) and socioeconomic status (SES) (B) on child speech. (C) Null three-way effect of normative development \times adult talk \times age (N.B.: normative \times age and adult talk \times age are significant; see Fig. 2). (D) null three-way effect of age \times adult talk \times monolingual status. (E and F) null effects of child gender (E) and SES (F) on adult talk. (G and H) null effect of normative development (G) and monolingual status (H) on adult talk.

American subsample (642 daylong recordings from 206 infants in 7 corpora; *SI Appendix, section 3G*). We essentially replicated the full-sample results in this subsample: Adult talk and age were significant predictors, whereas gender and SES (based on maternal education) were not. The significant adult talk \times age interaction also replicated. The main effect of normativity did not, likely because normativity's interaction with age was larger than in the full-sample analysis. Finally, we also tested whether removing the adult talk variable would result in an SES effect, i.e., testing whether adult talk was absorbing variance that would otherwise be accounted for by SES. This was not the case:

Removing the adult talk predictor, SES still does not account for significant variance in child speech in our analysis. A central contribution of this work is thus the clear lack of evidence we find for effects of SES (under several operationalizations focused on maternal education), on how much speech young children produce in day-to-day life.

Another potential concern is that our conclusions hinge on the use of LENA's particular algorithm; they do not. The findings above successfully replicate in the subset of data for which data stewards were able to share raw audio (11/18 corpora), which was analyzed with a wholly different algorithmic approach,

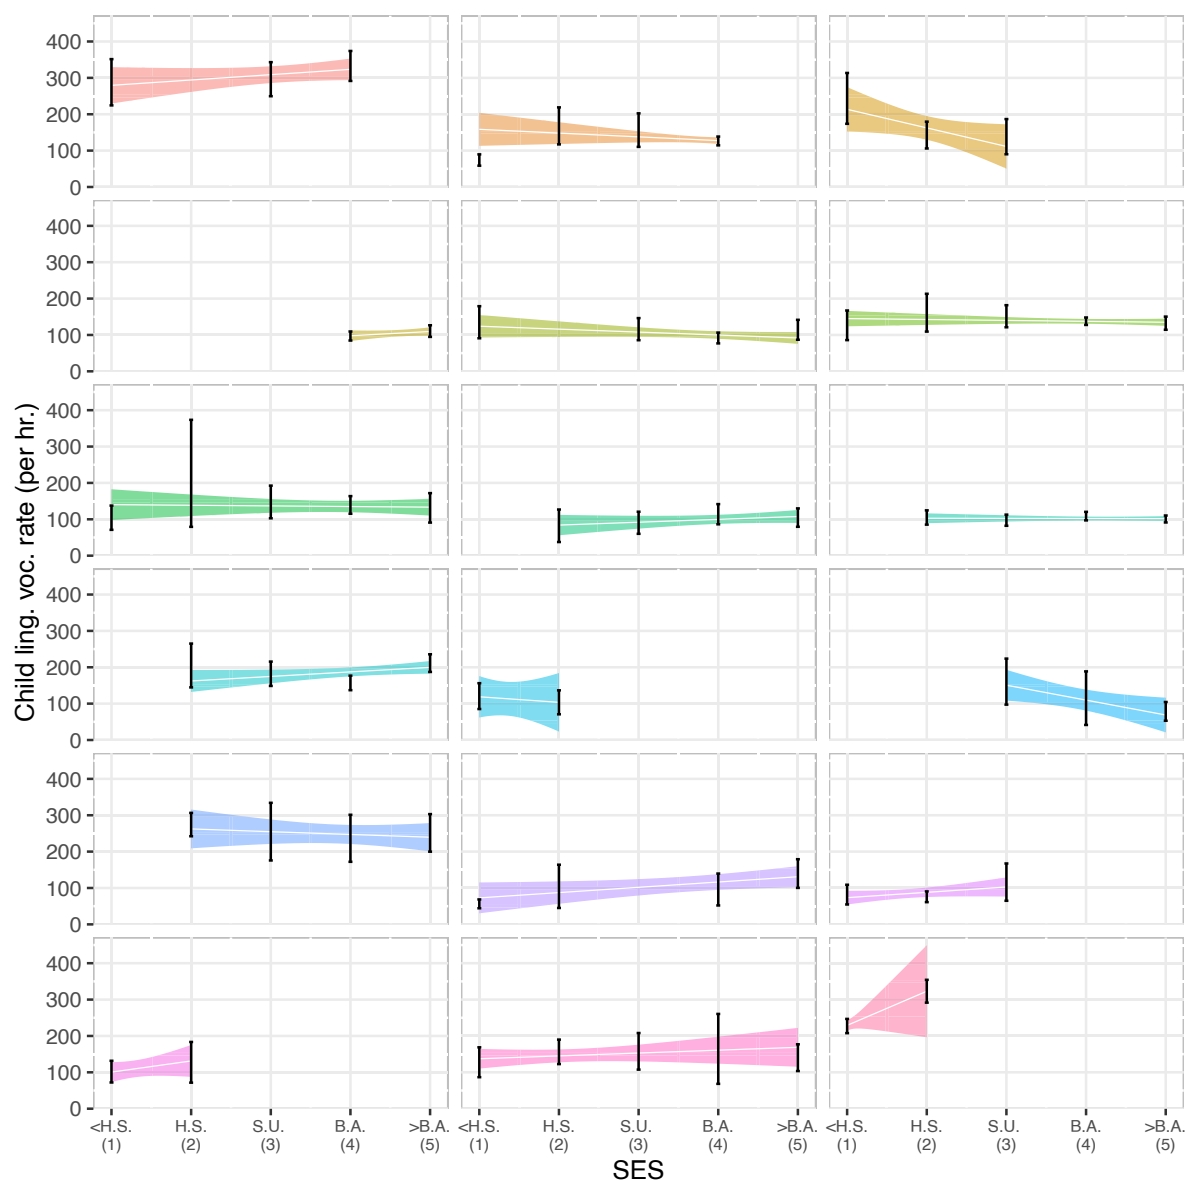


Fig. 4. Child speech as a function of SES within individual corpora. SES = maternal education levels as in Table 1. White lines = linear fit with 95% CIs in color, color = corpus. Black lines = 99% CIs of sample means bootstrapped separately from linear fit for each level of SES. These data (as well as our main models and further analyses in SI 3H/G) do not reveal an SES effect on child speech.

the Voice Type Classifier or VTC (*Methods* and *SI Appendix, section 3F*).[§] Yet another worry is that our focus on adult talk may mask other important contributions to children's language experiences, for instance, speech from other children. Testing this in a supplemental analysis, we confirm that the level of association found between adult talk and children's speech was unaffected by including other children's talk measured by LENA as a predictor variable (*SI Appendix, section 3H*), confirming that our key conclusions hold when factoring this other source of input in.

Finally, we also ran a model predicting adult talk (rather than child speech). The amount of adult talk was not significantly predicted by SES, child age, gender, and monolingual or normative status (Table 2, Fig. 3 *E–H*, and *SI Appendix, sections 3G and 3H*). Importantly, these null results replicated in the North American subset (*SI Appendix, section 3G*) as well as in every

other alternative analysis we attempted (*SI Appendix, section 3H*). Together, these analyses suggest that the relationship we find between adult talk and child speech in the child speech models is not attributable to child- or family-level factors affecting adult talk.

Speech and Other Early Vocal Behavior. While our central query concerned variability within early speech production, we conducted a further descriptive analysis examining how much of children's vocalizations were speech or speech-like, as opposed to the two other classes of LENA™-identified vocalizations: crying and vegetative sounds (e.g., burps, hiccups). We examined these vocalization types as a function of age, monolingual status, and normative status. As Fig. 2*C* shows, for children with normative development, the proportion of vocalizations that were speech increased from just over half to the vast majority over 2–48 mo. In contrast, the crying proportion fell steeply over the same period, from nearly half of vocalizations to a small fraction of them; the proportion of vegetative sounds was low and

[§]VTC too has been robustly validated relative to various gold standard manual measures (*SI Appendix, section 1E*)

Table 2. Model results predicting adult talk (i.e., adult vocalization count rate)

	β	SE	q
Intercept	-0.100	0.160	0.778
Child gender (Male)	0.174	0.148	0.547
SES(<H.S.(1))	0.239	0.173	0.547
SES(H.S.(2))	-0.015	0.194	0.939
SES(B.A.(4))	0.148	0.131	0.547
SES(>B.A.(5))	0.098	0.150	0.778
Control	0.084	0.055	0.547
Norm	0.013	0.103	0.939
Age	-0.030	0.029	0.547
Mono	-0.028	0.112	0.939
Gender (Male) \times SES(<H.S.(1))	-0.375	0.196	0.547
Gender (Male) \times SES(H.S.(2))	-0.263	0.252	0.547
Gender (Male) \times SES(B.A.(4))	-0.220	0.176	0.547
Gender (Male) \times SES(>B.A.(5))	0.016	0.201	0.939
Norm \times Age	-0.076	0.060	0.547
Mono \times Age	0.035	0.068	0.804

q -values show FDR-corrected P -values.

Note. None of the variables in our model predicted adult talk. All abbreviations and baselines are as in Table 1.

constant. Convergent with our speech analyses, monolingual status did not alter these patterns but normative status did: While the same overall patterns held for children with nonnormative development, their decrease in crying and increase in speech production with age was less steep (Fig. 2C).

As with more narrowly defined nonnormative populations [e.g., children with autism spectrum disorder (23)], we find clear divergences in language trajectories in our normative vs. nonnormative samples. This is notable because a) our nonnormative sample is heterogeneous (*SI Appendix, section 2A*) and b) as 2- to 48-mo-olds, many children with nonnormative classifications here were at risk of (but not yet diagnosed with) language delays or disorders. Automated speech analyses in naturalistic recordings thus hold promise for future research into early diagnostics (24, 25).

Discussion

Adult Talk and Child Speech. Children who heard more adult talk produced dramatically higher rates of speech, and this effect increased with age. This result feeds into ongoing theoretical debates regarding the relevance of individual differences (26). Although we cannot infer causality from our correlational data, it is useful to consider possible causal paths that could in principle have led to our results. A correlation between child speech and adult talk is compatible with at least three explanations: 1) Children who produce more speech elicit more talk from adults; 2) language-dense environments lead children to produce more speech; or 3) a third variable causes increases in both adult talk and child speech.⁴

Our model predicting adult talk (Table 2) can be brought to bear on Explanation 1. If children talking more elicited more talk from adults, then we would have expected to see that age and normative status were significant predictors of adult talk. Instead, we find that neither these (nor any other variables in our model) predicted the quantity of adult talk (Fig. 3G). Nonetheless, the precise statistical analyses we carried out do

⁴Our analyses suggest that one such potential third variable, differences in activities across recordings, is not a likely candidate for the correlation between child speech and adult talk (*SI Appendix, section 4*).

not allow us to directly rule out any of the explanations, a combination of which may be jointly true. Establishing a precise causal chain will require careful consideration of a variety of proximal and ultimate pathways through which child and adult behaviors are shaped. As one example, given that most children here are genetically related to their adult caregivers, we may be observing covariance in amount of talk and its linguistic correlates (Explanation 3). Evaluating these alternatives requires evidence from children raised by unrelated caregivers or from genome-wide association studies, as genetic and environmental factors remain challenging to disentangle (27). In this vein, recent work with adopted 15- to 73-mo-olds provides evidence for input effects (maternal utterance length and/or lexical diversity) on adopted children's vocabulary size (measured via a caretaker checklist) (28). This study suggests that shared genetics is not the sole contributor to links between (at least these proxies for) caretaker input and child language outcomes. Moreover, shared genetics is just one of the ways in which adult and child behavior may be independently shaped by an unmeasured confounded variable (as per Explanation 3). For instance, other third variables related to dimensions like personality, neighborhood, and childcare context too may be contributors (29, 30). These explanations can only be definitively teased apart by future work.

Insight on Child and Family Factors. Our main models, figures showing the raw data, and additional analyses (in the North American subset of the data, as well as using an alternative algorithm, see *SI Appendix, section 3F*) reveal effects of normativity, age, and adult talk but not SES (measured here through maternal education), child gender, or monolingualism. To illustrate the complexities involved in determining causal links between child and family factors and child language skills, we again consider how causal links might manifest, using SES as a central example.

Our findings bear on debates regarding SES-associated academic achievement differences in Western industrialized societies (31, 32). Slower language development has often been attributed to parents from lower-SES backgrounds providing less input to their children [viewed from a middle-class Western-centric perspective (32)], leading to calls for behavioral interventions aiming to increase it. Proponents of such interventions might highlight our correlation between adult talk and child speech; critics might instead underscore our finding that SES was not significant in our main analyses nor in every other reanalysis we attempted (*SI Appendix, sections 3E–3G*).

A full understanding of how SES may relate to children's language input is complicated for empirical and conceptual reasons, leaving strong conclusions premature. On the empirical side, two recent meta-analyses have investigated SES–input correlations, one focused on LENA™ measures (15) and the other based on human-annotated measures (mostly from short lab recordings) (16). The former finds evidence consistent with a publication bias; correcting this bias statistically nearly halves the association between SES and LENA™'s adult talk measure ($r = 0.19$ vs. 0.12). The latter finds a sizeable SES effect when inspecting infant-directed speech ($r = 0.34$) and a much smaller one when analyzing overall input quantities ($r = 0.09$). Together, these studies suggest that our best estimate of the association between overall input quantities and SES is small ($r = 0.1$) and may not be detectable even with a sample as large as ours (where the effect was estimated at $|d| = 0.06$, or $|r| = 0.03$, which did not reach the threshold for significance). Similarly, descriptive plots of the potential correlation between our SES proxy and children's speech (Fig. 4) did not suggest a

strong or stable relationship across the 18 corpora, leading to our conclusion that, in the sample as a whole, on average, maternal education does not predict how much adults and children talk.

On the conceptual side, SES differences in input and language skills may depend on how language is measured (33). For instance, we speculate that SES effects may be magnified by measures like prevalence of low-frequency words and complex sentence structures common in written text. Such words and structures may occur more in the input to Western, higher-SES children because of parenting practices stereotypical in these groups (34). Moreover, such measures may predict academic achievement better than others because of the importance literacy has in Western schooling today. In contrast, SES differences in input may be minimized by holistic measures of speech quantities. Indeed, a strength of daylong recordings is that they provide a relatively neutral (rather than Western, high SES-centric) measure, as they tap into how much children are contributing (via speech) to their community's conversational interactions instead of how many rare words or complex constructions they have been taught.

An exclusive focus on word counts or speech quantities likely misses certain behaviors. As machine learning advances (35), it may soon be possible to automatically transcribe conversations happening in daylong recordings (at least in monolingual high-resource language contexts). We suspect that analysis of conversational content may reveal SES differences in, e.g., rare word use or family practices around book-reading even in naturalistic samples (36). Future work with a high-density longitudinal lens is also needed to assess the predictive value of global quantitative measures of speech (like those we employ) relative to more specialized measures (e.g., book-reading practices) with respect to culturally relevant outcomes (e.g., academic achievement, pragmatic competence in multiparty conversation, etc.).

In our view, causal links between parental behavior and children's outcomes can best be illuminated by randomized control trials. Discovering and leveraging such links to change long-term language outcomes depends on community partnership-based approaches that are informed by the role that structural inequalities play in these outcomes and engage with culturally informed perspectives (37). The present results should not be used to deny families access to resources that are linked with better outcomes for children and their families.

Complicated causal effects are integral to all developmental processes. While we illustrated this with our SES null results, we also found no differences in child speech or adult talk as a function of child gender or multilingual status. Regarding multilingualism, we could not examine relative input in each language the child was exposed to. Future machine learning advances will permit the separate quantification of different languages in daylong recordings, but this must happen alongside reflection on how to fairly measure input and outcomes in such heterogeneous populations (38–40).

Automated Tools and What They Count. A key benefit of our approach is that we were able to pool and identically process 40,933 h of independently collected data (*SI Appendix, section 3A*). Moreover, unlike parental surveys, clinical assessments, lab instruments, or hand-annotated data, current published evidence suggests that the LENA™ algorithm's results do not vary systematically by language [though they do vary somewhat across samples, (12)]. More relevant here, in analyzing the algorithm's accuracy as a function of samples grouped by language and cultural features, we found no significant differences (*Methods and SI Appendix, section 1E*).

While children's language skills grow dramatically over 2–48 mo, our measure is not an index of comprehension [which can show quite a different trajectory, 41] but rather of observable linguistic behavior, focusing exclusively on children's rate of linguistic vocalizations (*SI Appendix, section 3B*). These results certainly do not deny effects found on proxies of more narrow-scope linguistic developments (e.g., vocabulary, processing efficiency, or syntactic complexity), given that some predictors that fail to explain variance here may nonetheless be significant there (3, 42).

The same holds for our measure of adult talk, which is quantitative and holistic; additional research is needed to distinguish child-directed from child-available speech, with the latter including all speech the child hears. Although some research suggests child-directed speech shows tighter correlations with children's vocabulary than child-available speech does (43, 44), the importance of the latter has not been as fully studied for other types of language knowledge (45). Notably, this paper specifically documents a significant link between adult-produced child-available speech and everyday child speech behavior. Therefore, it would be relevant to further investigate the strength of the predictive value of overall adult talk (which was a significant predictor here) versus child-directed talk, in a similarly large and diverse sample as the present one. Unfortunately, automated tools for separating child-directed from overheard speech are not yet sufficiently accurate to make this possible (46). Future work could also develop promising approaches for considering other sources of speech (e.g., other children) given their relevance as a function of family structure (47). These approaches were not possible here due to both technical algorithmic constraints and family structure information not being available in our data subsets. Another fruitful future direction could consider conversational dynamics, studying both children's tendency to vocalize around adults and the complexity of such vocalizations. Recent work (that is critically reliant on human annotation of social intent) raises particularly interesting ideas in this domain (14, 48). Relatedly, novel exploratory analyses describing the acoustics of children's vocalizations (49) hold promise for driving future hypothesis-testing work building on the present results.

Whatever measures are employed in the future as proxies of child language production and input, we strongly encourage researchers to consider psychometric properties and ecological validity. The current approach demonstrates measure validity that is comparable to that of other standard infant instruments (*SI Appendix, sections 1D and 1E*). As context, measures used as proxies for infant language and cognitive knowledge are inherently noisier than the best batteries used to assess highly educated adults in Western-centric settings. Notably, even there, reliabilities can fall well below $r = 1$.[#]

Moreover, standardized tests face ecological validity threats, particularly when applied cross-culturally. If our goal is to measure and understand the human mind, we need implementable, culturally sensitive, and appropriate ways of measuring human behavior on a large scale. To our knowledge, there are no such measures whose reliability has been examined, driving us to conduct extensive quantification of the reliability of the metrics we employed here (*SI Appendix, sections 1D and 1E*). We found that our measures show levels of reliability that are consistent with those already in use for research

[#] For instance, prior work finds test-retest reliabilities as low as $r = 0.6$ for certain sections of the widely used Wechsler Adult Intelligence Scale among North American English-speaking adults (50).

and clinical purposes in infant populations. For example, the MacArthur-Bates Communicative Development Inventory (a parental report instrument used largely as a proxy for vocabulary size) has been the basis for cross-linguistic, demographic, and clinical research (9, 51–53) and reports a median correlation between itself and laboratory measures of 0.61 (54). Our median accuracy comparing automated and manual annotation for each of our algorithms (LENA™ and VTC) is 0.74, squarely in line with field standards (*SI Appendix, section 1E*). Indeed, converging evidence across these two wholly separate algorithms regarding overall accuracy of our measure serves to increase confidence in the validity of our results.

In sum, rather than eliciting knowledge or caregiver-child interaction in a constrained lab setting, or using checklists in contexts where they make little sense socioculturally, we measure everyday language use *en masse*. Our measure of early speech production is global, since we simply measure more versus less speech or speech-like production on the part of adults and children as they go about their daily life. And yet, these measures have important advantages, which led us to select them as proxies here, including comparable reliability to other measures of language development commonly used in both research and applied settings (*Methods* and *SI Appendix, sections 1D and 1E*); reported correlations between them and finer-grained, “qualitative” measures of language development (*SI Appendix, section 1D*), and convergent validity with respect to standardized language tests (13). Most importantly, our speech measure merits consideration as one of many possible proxies of language development thanks to its cross-cultural adaptability, observer-free sampling volume, and sheer ecological validity. Indeed, our results raise the possibility that more ecologically valid lexical, phonetic, or grammatical measures will also reveal stability across factors like SES (55), gender, and multilingualism. Exploring these factors, however, awaits machine-learning developments that can extract such fine-grained linguistic measures from the raw audio collected with child-worn devices.

Conclusion

Our analysis of speech behavior in daily life around the world evinces scientific progress on two fronts. First, by revealing substantial variation in young children’s speech, we provide evidence against a monolithic picture of language development. Instead, this work reveals individual variation as *fundamental* to our understanding of this species-wide ability. Second, by tapping into natural speech interactions at unprecedented scale and diversity, we are able to move beyond prior work by simultaneously considering the interlocking factors that affect speech production over early development. Our results reveal not only the expected correlations with age and clinical factors but also substantial associations with adult talk. All other factors paled in comparison with these three, the null effect of our SES proxy being of particular noteworthiness. These findings open exciting avenues for both theoretical research and potential applications, including the prospect of behavioral interventions to harness adult talk in the context of speech and language diagnoses. Small-scale experimental and observational research has been fundamental to our understanding of language, development, and the human mind. Machine learning (like that in speech technology) promises to extend our scientific reach by exploding the range of everyday interactions we are able to capture and analyze. Just as recent technological innovations have opened new vistas in understanding the vocalizations of mice and whales (6, 7), so too does speech technology have the potential to

reveal how everyday human communication gives rise to language learning in children around the world.

Methods

All code used to generate our analysis and the manuscript is available at https://osf.io/9v2m5/?view_only=50df17fc0844145ae692c35b78c6b08.

Data Discovery and Integration. We took steps to counter a prevalent bias for normative North American data (see *SI Appendix, section 3A* for corpus constitution procedure). Included data were independently collected by 18 stewards (56–77); see *SI Appendix, section 5* for the list of publications based on individual datasets. We note that while our corpora covered a much greater variety of participants than prior work, it would not be appropriate to interpret our samples as comprehensively representative of the country or language community from which they are drawn.

Socioeconomic status and normative development were streamlined for cross-corpus consistency (*SI Appendix, sections 2A, 2B, and 3A, and Fig. S3A.1*). For socioeconomic status, we use maternal education, a reliable proxy for SES in previous research on language development (18, 78). Maternal education was available across all datasets and could be converted into a 5-point maternal education scale with levels corresponding to less than high school degree, high school degree or equivalent, some college/vocational/associate degree-level training, university/college degree, and advanced degree (*SI Appendix, section 2B and Table S2B.1*).

For nonnormative development, data stewards had tagged a wide variety of infant or familial characteristics as potentially nonnormative. We confirmed that the classification was backed up by extant literature (*SI Appendix, section 2A*). Infants ultimately classified as having nonnormative development in the present sample include those who met one or more of the following criteria: preterm birth (<37 wk); diagnosed speech or language delay; global developmental delay; low birth weight (<2,500 g when specified); hearing loss, hearing aids, or cochlear implants; familial risk of autism spectrum disorder, specific language impairment, and/or dyslexia; and other relevant genetics syndromes. Notably, our child vocalization rate measure is not a standardized normed clinical evaluation, and thus nonnormative status may not necessarily translate to behavior that falls >1 SDs below the norm in these naturalistic recordings.

Analysis Details. We first randomly partitioned the data within each corpus such that 35% of monolingual, normative children were placed in an exploration set (N children = 264; N recordings = 850), and all others in a confirmation set (N children = 737; N recordings = 2,025) (*SI Appendix, section 3A*). The exploration set was used to study the psychometric properties of potential language input and output variables (*SI Appendix, section 3B*), resulting in the selection of the output variable referred to as child speech above, and CVCr (Child Vocalization Count rate) in analysis and supplementary files (*SI Appendix, section 3B and Table S3B.1*); and the input variable referred to as adult talk above, and AVCr (Adult Vocalization Count rate) in analysis and supplementary files (*SI Appendix, section 3B and Table S3B.2*). Note that this includes both child-directed and child-available speech.

In addition, we used the exploration set to check the robustness of results to variation in random effect structure and explored diverse model structures using mixed models in R’s lme4 package (79), checking whether the addition of effects or interactions explained additional variance (*SI Appendix, section 3C*). This led us to a) include overlap rate as a covariate (see *SI Appendix, Fig. S3C.1*) to control for the fact that in noisy environments, more child speech and adult talk within the same recordings may be labeled as “overlap” by LENA (and thus not attributed to either speaker type) and b) to not include random slopes for any of the predictors. Regarding the latter choice, our exploration of random effect structure revealed that models including random slopes for any of the predictors (notably including gender and SES) as a function of corpus led to nonconvergent models. While such nonconvergence could be due to various reasons, the most likely explanation is that the model is overparametrized (80), i.e., variance cannot be reliably attributed to predictors *within* each corpus (see *SI Appendix, section 3H* for additional checks, including one including random slopes for SES, and *SI Appendix, section 2B* for discussion of alternatives to our SES implementation).

Evaluation against human annotations. To assess the validity of our child speech and adult talk measures, we evaluated them against human annotations (see *SI Appendix, sections 1D and 1E* for further information). The median correlation of human to algorithm performance for the algorithms is >0.7 , i.e., comparable reliability to established developmental clinical and research instruments (81–83). As far as we know, the present multicultural validation exceeds those from prior research instruments. For example, the Ages and Stages Questionnaire (84) is a standard instrument used at well-child visits in the United States. It is also recommended by the World Bank as one of the most popular tools to measure child development, used in at least 20 countries (85). And yet, a recent systematic review (83) reports only six reliability analyses (averaging, e.g., 0.7 for internal consistency at 24 mo). Relative to this, our validation effort containing estimates for 14/18 corpora and finding strong validity is notable. Finally, one may wonder whether the LENA™ algorithm performs less well for languages and cultures that diverge from its training set, which was English-learning children growing up in an urban/suburban US setting. Although we observe considerable corpus variation, this variation is not attributable to whether children were learning English or growing up in an urban setting, as assessed by Welch's *t*-tests, for either our child speech measure (CVCr; English versus non-English medians 0.785 vs. 0.71, $t(6.04) = -0.5$, $P = 0.637$; urban versus rural medians 0.77 vs. 0.71, $t(8.11) = -0.46$, $P = 0.661$), or for our adult talk measure (AVCr; English versus non-English medians 0.75 vs. 0.74, $t(7.91) = 0.42$, $P = 0.686$; urban versus rural medians 0.75 vs. 0.74, $t(3.07) = -0.23$, $P = 0.835$). Instead, our results suggest that corpus variation more likely reflects how the human annotation was done rather than how well the algorithm worked, since the corpora with lower reliabilities were also those in which the human annotation was more coarse-grained (*SI Appendix, section 1E*).

Additional algorithm. To make sure that key conclusions were robust to methodological details, we reanalyzed the subset of the data for which data stewards shared audio with a newer, open-source alternative to LENA™: the Voice Type Classifier (VTC) (86). Like the LENA™ algorithm, VTC returns an estimation of child and adult vocalization counts. A total of 1,065 audio files from 11 corpora were available for this reanalysis (*SI Appendix, section 3F*).

The VTC algorithm employs a completely different approach than the proprietary algorithm developed by LENA™, including the use of neural networks running directly from the audio (rather than from MFCC features). VTC allows multiple talker classes to be activated at the same time, whereas in the LENA™ algorithm, overlap between talkers (or between a talker and noise) is tagged as "Overlap," which is not counted toward children's input or output. VTC also differs from LENA™ in its training set. While LENA™ was trained entirely on data from North American, monolingual English-learning, urban children, VTC was developed using the combination of various corpora of children residing in urban or rural settings and learning one or more of several languages (including the tonal language Minn, French, Ju|'hoan, Tsimane, English, and several others,

in rough order of quantity of data). Further information on accuracy is provided in *SI Appendix, section 1E*; both algorithms render similar accuracy when compared to human annotation as noted above.

Models. We used linear mixed regressions (Gaussian family) and established model structure from the exploration data (*SI Appendix, section 3C*). Hypotheses were derived from exploratory models and systematic reviews of the literature on monolingualism and normativity (*SI Appendix, section 3D*). The model predicting the rate of children's linguistic vocalizations (i.e., child speech) was the following: $child_gender + SES + child_normative * AVCr * age + child_monolingual * AVCr * age + overlap + (1 + overlap + AVCr|corpus) + (1|corpus : child_id)$. The model predicting the rate of adult linguistic vocalizations (i.e., adult talk) was the following: $child_gender + SES + child_normative * age + child_monolingual * age + overlap + (1 + overlap|corpus) + (1|corpus : child_id)$. Full model details and a link to model diagnostics are provided in *SI Appendix, section 3E*. We report estimates (standardized, which serve as effect sizes), standard errors of the estimates, and *q*-values (FDR-corrected *P*-values); see Tables 1 and 2.

Participants. Table 3 lists participant characteristics noting both 1) the exploration/confirmation split (*SI Appendix, section 3A*) and 2) that some children provided multiple recordings. We excluded 2/850 recordings from 1/264 children from the exploration set and 8/2,025 recordings from 5/737 children in the confirmation set from our models because data regarding their maternal education was missing. For child gender, there were slightly more boys than girls. This was in part because corpora with children with nonnormative development also include children with normative development matched in gender, leading to an overrepresentation of boys since more boys than girls have nonnormative development. See Table 3 and Fig. 5 for specific numbers and visualized distributions.

Language background. The languages represented in these data covered many languages and language families. Using classifications from Glottolog (87), we report that our 18 corpora feature 10 primary languages (Dutch, English, Finnish, French, Spanish, Swedish, Tsimane, Vietnamese, Wolof, and Yéfi Dnye) from 5 distinct language families and one isolate (Atlantic-Congo, Austroasiatic, Indo-European, Mosetén-Chimané, Uralic, and Yéfi isolate); see Fig. 1. Based on corpus metadata provided by each data steward, the recorded children were also exposed to an additional 33 languages (Arabic, ASL, Berber, Cantonese, Croatian, Danish, Farsi, Frisian, German, Greek, Hindi, Hungarian, Indonesian, Italian, Japanese, Khmer, Korean, Macedonian, Malay, Malayalam, Mandarin, Norwegian, Papiamentu, Polish, Portuguese, Romanian, Russian, Sahaptin, Slovenian, Solomon-Islands Pidgin, Thai, Turkish, and Yoruba), which add 11 further language families (Afro-Asiatic, Austroasiatic, Austronesian, Deaf Sign Languages–LSFic, Dravidian, Japonic, Koreanic, Sahaptian, Sino-Tibetan,

Table 3. Number of children and recordings by demographic variables, split by exploration and confirmation subsets

Variables	Levels	Exploration subset		Confirmation subset	
		Children	Recs.	Children	Recs.
Gender	Boys	156	516	398	1,016
	Girls	107	332	334	1,001
Normativity	Normative	263	848	550	1,731
	Nonnormative	0	0	182	286
Lingualism	Monolingual	263	848	662	1,847
	Multilingual	0	0	70	170
SES	<H.S.(1)	94	120	202	265
	H.S.(2)	10	26	60	159
	S.U.(3)	27	116	115	309
	B.A.(4)	86	355	241	786
	>B.A.(5)	46	231	114	498
Total N		263	848	732	2,017

Note. Children = # of children; Recs. = # of daylong recordings. In SES, <H.S. = children whose mothers have (the equivalent of) less than a high school degree; H.S. = high school degree; S.U. = some university; B.A. = bachelor's degree; >B.A. = more than a bachelor's degree. Multilingual children, children with nonnormative development, and 65% of all other children were reserved for the confirmation subset. N.B. the six children with missing data for maternal education are omitted from this table.

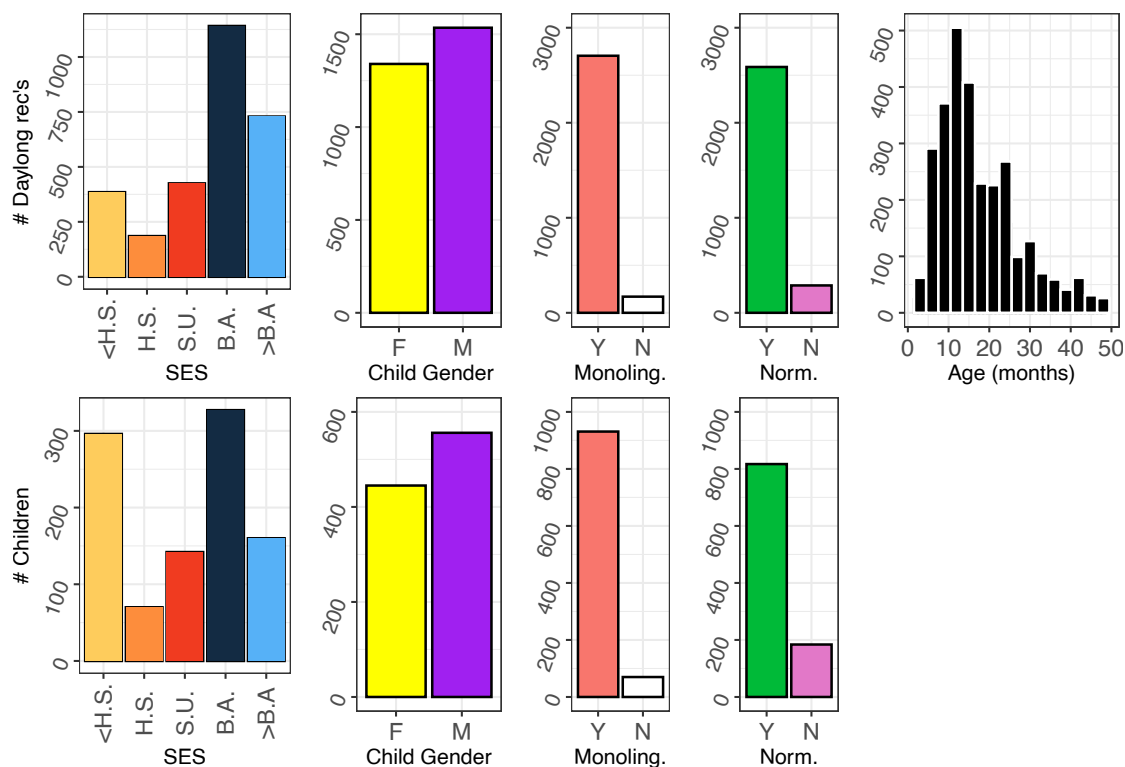


Fig. 5. Sample demographics. Number of daylong recordings (*Top row*) and children (*Bottom row*) in the full dataset across demographic variables. For socioeconomic status (SES), <H.S. = less than high school degree, H.S. = high school degree, S.U. = some university, B.A. = bachelor's degree, >B.A. = advanced degree. For child gender, F = female, M = male. For monolingual status (monoling.), Y = monolingual, N = not monolingual. For normative development (norm.), Y = normative, N = nonnormative.

Tai-Kadai, and Turkic) and bolster data from three language families already represented by the primary languages (Atlantic-Congo, Indo-European, and Uralic).

Data, Materials, and Software Availability. Anonymized (tabular) data and all relevant code have been deposited with the Open Science Foundation (https://osf.io/9v2m5/?view_only=50df17fc0844145ae692c35b78c6b08) (88). The raw audio recordings are not able to be shared given the consent process participants underwent, but all derived tabular data are fully shared.

ACKNOWLEDGMENTS. We thank Adriana Weisleder, Ann Weber, Camila Scaff, Karmen McDivitt, Evan Kidd, Bridgette Kelleher, Hillary Ganek, Anne Fernald, Hanna Elo, Samantha Durrant, Yatma Diop, John Bunce, and Sarp Uner for organizing and/or sharing their data with us. We acknowledge the following funding sources: ANR-16-DATA-0004 ACLEW, ANR-14-CE30-0003 MechELex; J. S. McDonnell Foundation; ERC H2020 (ExELang, 101001095) (A.C.); NEH HJ-253479-17 (E.B.); NIH DP5-OD019812 (E.B.); NSF BCS-1844710 (E.B.); NSF SBE-0354453 (N.R.-E.); ESRC ES/L008955/1 (C.F.R.); SSHRC 435-2015-0628, 869-2016-0003 (M.S.); NSERC 501769-2016-RGPDD (M.S.); Netherlands Organisation for Scientific Research 275-89-033 (M.C.); NIMH K23MH111955; NIDCD F31DC018219 (L.R.H.); MAW 2011.0070 (I.-C.S. and E.M.); MAW 2013.0056 (I.-C.S. and E.M.); Basque Government through the BERC 2022-2025 program and the Spanish State Research Agency through BCBL Severo Ochoa excellence accreditation CEX2020-001010/AEI/10.13039/501100011033, and

the Ramon y Cajal Fellowship, RYC2018-024284-I (M.K.); ARC CE140100041 (Evan Kidd).

Author affiliations: ^aDepartment of Psychology, Harvard University, Cambridge, MA 02138, United Kingdom; ^bDepartment of Psychology, University of Manitoba, Winnipeg, CA R3T 2N2; ^cDepartment of Linguistics, Stockholm University, Stockholm SE-106 91, Sweden; ^dDepartment of Special Education, Stockholm University, Stockholm SE-106 91, Sweden; ^eLanguage Development Department, Max Planck Institute for Psycholinguistics, Nijmegen 6525 XD, Netherlands; ^fDonders Centre for Brain, Cognition and Behaviour, Radboud University, Nijmegen 6525 XZ, Netherlands; ^gAustralian Research Council Centre of Excellence for the Dynamics of Language, Australian National University, ACT 2601, Australia; ^hPsychological Sciences, University of Connecticut, Storrs, CT 06268; ⁱDepartment of Psychological Sciences, Purdue University, West Lafayette, IN 47907; ^jBasque Center on Cognition Brain and Language, Donostia-San Sebastian 20009, Spain; ^kIkerbasque - Basque Foundation of Science, Bilbao 48009, Spain; ^lDépartement d'études Cognitives, École normale supérieure, École des hautes études en sciences sociales, Centre National de la Recherche Scientifique, PSL University, Laboratoire de Sciences Cognitives et Psycholinguistique, Paris 75005, France; ^mComparative Human Development Department, University of Chicago, Chicago, IL 60637; ⁿSchool of Music, Ohio State University, Columbus, OH 43210; and ^oRoyal Dutch Kentalis Utrecht, Utrecht 3527 JP, Netherlands

Author contributions: E.B., M.C., and A.C. designed research; E.B., M.C., and A.C. performed research; E.B., M.S., I.-C.S., C.F.R., N.R.-E., L.R.H., E.M., M.K., A.G., M.C., P.v.A. and A.C. contributed new reagents/analytic tools; E.B., M.C., A.G., and A.C. analyzed data; E.B. and A.C. prepared materials for and/or led group decision-making; E.B., M.S., C.F.R., N.R.-E., A.G., M.C., L.B., and A.C. contributed to the decision-making on the analytic approach, including selection of exploratory and confirmatory sets, selection of variables, identification of hypotheses and/or specification of models; E.B., M.S., I.-C.S., C.F.R., N.R.-E., L.R.H., E.M., M.K., L.B., P.v.A., and A.C. contributed to supplementary materials, Open Science Framework project page and/or other documentation; E.B., M.K., M.C., and A.C. contributed to visualizations; and E.B., M.S., C.F.R., M.C., and A.C. wrote the paper.

1. S. Pinker, *The Language Instinct* (Morrow, New York, NY, 1994).
2. D. K. Oller *et al.*, Infant boys are more vocal than infant girls. *Curr. Biol.* **30**, R426–R427 (2020).
3. A. Fernald, V. A. Marchman, A. Weisleder, SES differences in language processing skill and vocabulary are evident at 18 months. *Dev. Sci.* **16**, 234–248 (2013).
4. J. Gilkerson *et al.*, Mapping the early language environment using all-day recordings and automated analysis. *Am. J. Speech Lang. Pathol.* **26**, 248–265 (2017).
5. R. Coe, It's the effect size, stupid (2002). <https://f.hubspotusercontent30.net/hubfs/5191137/attachments/ebe/ESguide.pdf>. Accessed 28 July 2021.
6. K. R. Coffey, R. G. Marx, J. F. Neumaier, DeepSqueak. *Neuropsychopharmacology* **44**, 859–868 (2019).
7. Y. Shiu *et al.*, Deep neural networks for automated detection of marine mammal species. *Sci. Rep.* **10**, 607 (2020).
8. T. Broesch *et al.*, Navigating cross-cultural research: Methodological and ethical considerations. *Proc. R. Soc. B.* **287**, 20201245 (2020).
9. M. C. Frank, M. Braginsky, V. A. Marchman, D. Yurovsky, Variability and Consistency in Early Language Learning (MIT Press, Cambridge, MA, 2021). <https://langcog.github.io/wordbank-book/index.html#>.
10. F. J. Zimmerman *et al.*, Teaching by listening. *Pediatrics* **124**, 342–349 (2009).
11. E. Bergelson, A. Amatuni, S. Dailey, S. Koorathota, S. Tor, Day by day, hour by hour. *Dev. Sci.* **22**, e12715 (2019).

12. A. Cristia, F. Bulgarelli, E. Bergelson, Accuracy of the language environment analysis system segmentation and metrics. *J. Speech Lang. Hear. Res.* **63**, 1093–1105 (2020).
13. Y. Wang, R. Williams, L. Dilley, D. M. Houston, A meta-analysis of the predictability of LENA™ automated measures for child language development. *Dev. Rev.* **57**, 100921 (2020).
14. D. K. Oller *et al.*, Preterm and full term infant vocalization and the origin of language. *Sci. Rep.* **9**, 14734 (2019).
15. L. Piot, N. Havron, A. Cristia, Socioeconomic status correlates with measures of Language Environment Analysis (LENA) system. *J. Child Lang.* **49**, 1037–1051 (2022).
16. S. Dailey, E. Bergelson, Language input to infants of different socioeconomic statuses. *Dev. Sci.* **25**, e13192 (2022).
17. A. L. Richman, P. M. Miller, R. A. LeVine, Cultural and educational variations in maternal responsiveness. *Dev. Psychol.* **28**, 614–621 (1992).
18. E. Hoff, The specificity of environmental influence. *Child Dev.* **74**, 1368–1378 (2003).
19. D. Hartas, Families' social backgrounds matter. *Br. Edu. Res. J.* **37**, 893–914 (2011).
20. C. F. Rowland, K. Alcock, K. Meints, The (null) effect of socio-economic status on the language and gestures of young infants. OSF. <https://osf.io/hwg4c>. Accessed 21 April 2023.
21. D. A. Hackman, M. J. Farah, Socioeconomic status and the developing brain. *Trends Cogn. Sci.* **13**, 65–73 (2009).
22. UNESCO Institute for Statistics, ISCED 2011 (UNESCO Institute for Statistics, 2012), <https://dx.doi.org/10.15220/978-92-9189-123-8-en>. Accessed 22 November 2023.
23. D. K. Oller *et al.*, Automated vocal analysis of naturalistic recordings from children with autism, language delay, and typical development. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 13354–13359 (2010).
24. J. Rankine *et al.*, Language ENvironment Analysis (LENA) in Phelan-McDermid Syndrome. *J. Autism Dev. Disord.* **47**, 1605–1617 (2017).
25. J. McDaniel *et al.*, Effects of pivotal response treatment on reciprocal vocal contingency in a randomized controlled trial of children with autism spectrum disorder. *Autism* **24**, 1566–1571 (2020).
26. E. Kidd, S. Donnelly, Individual differences in first language acquisition. *Annu. Rev. Linguist.* **6**, 319–340 (2020).
27. D. V. M. Bishop, Ten questions about terminology for children with unexplained language problems. *Int. J. Lang. Commun. Disord.* **49**, 381–415 (2014).
28. J. R. Coffey, C. L. Shafto, J. C. Geren, J. Snedeker, The effects of maternal input on language in the absence of genetic confounds. *Child Dev.* **93**, 237–253 (2022).
29. M. Hilton, K. E. Twomey, G. Westermarck, Taking their eye off the ball. *J. Exp. Child Psychol.* **183**, 134–145 (2019).
30. A. De Marco, L. Vernon-Feagans, Rural neighborhood context, child care quality, and relationship to early language development. *Early Educ. Dev.* **24**, 792–812 (2013).
31. R. M. Golinkoff, E. Hoff, M. L. Rowe, C. S. Tamis-LeMonda, K. Hirsh-Pasek, Language matters. *Child Dev.* **90**, 985–992 (2019).
32. D. E. Sperry, L. L. Sperry, P. J. Miller, Reexamining the verbal environments of children from different socioeconomic backgrounds. *Child Dev.* **90**, 1303–1318 (2019).
33. E. Ochs, T. Kremer-Sadl, Ethical blind spots in ethnographic and developmental approaches to the language gap debate. *Lang. Soc.* **170**, 39–67 (2020).
34. D. K. Dickinson, J. A. Griffith, R. M. Golinkoff, K. Hirsh-Pasek, How reading books fosters language development around the world. *Child Dev. Res.* **2012**, e602807 (2012).
35. M. Lavechin *et al.*, Brouhaha. arXiv [Preprint] (2022). <http://arxiv.org/abs/2210.13248>. Accessed 13 November 2023.
36. C. Nutbrown *et al.*, Families' roles in children's literacy in the UK throughout the 20th century. *J. Early Child. Lit.* **17**, 551–569 (2016).
37. A. Weber, A. Fernald, Y. Diop, When cultural norms discourage talking to babies. *Child Dev.* **88**, 1513–1526 (2017).
38. E. Bialystok, J. F. Werker, Special issue: Systematic effects of bilingualism on children's development. *Dev. Sci.* **20**, e12535 (2017).
39. D. K. Oller, B. Z. Pearson, A. B. Cobo-Lewis, Profile effects in early bilingual language and literacy. *Appl. Psycholinguist.* **28**, 191–230 (2007).
40. T. Grüter, N. Hurtado, V. A. Marchman, A. Fernald, "Language exposure and online processing efficiency in bilingual development" in *Input and Experience in Bilingual Development*, T. Grüter, J. Paradis, Eds. (John Benjamins Publishing Company, 2014), pp. 15–36.
41. E. V. Clark, B. F. Hecht, Comprehension, production, and language acquisition. *Annu. Rev. Psychol.* **34**, 325–349 (1983).
42. M. Eriksson *et al.*, Differences between girls and boys in emerging language skills. *Br. J. Dev. Psychol.* **30**, 326–343 (2012).
43. L. A. Shneidman, M. E. Arroyo, S. C. Levine, S. Goldin-Meadow, What counts as effective input for word learning? *J. Child Lang.* **40**, 672–686 (2013).
44. A. Weisleder, A. Fernald, Talking to children matters: Early language experience strengthens processing and builds vocabulary. *Psychol. Sci.* **24**, 2143–2152 (2013).
45. A. Cristia, Language input and outcome variation as a test of theory plausibility. *Dev. Rev.* **57**, 100914 (2020).
46. B. Schuller *et al.*, The INTERSPEECH 2017 Computational Paralinguistics Challenge (2017), pp. 3442–3446.
47. A. Cristia, L. Gautheron, H. Collieran, Vocal input and output among infants in a multilingual context. *Dev. Sci.* **26**, e13375 (2023).
48. G. M. Pretzer, L. D. Lopez, E. A. Walle, A. S. Warlaumont, Infant-adult vocal interaction dynamics depend on infant vocal type, child-directedness of adult speech, and timeframe. *Infant Behav. Dev.* **57**, 101325 (2019).
49. V. P. S. Ritwika *et al.*, Exploratory dynamics of vocal foraging during infant-caregiver communication. *Sci. Rep.* **10**, 10469 (2020).
50. E. Strauss *et al.*, *A Compendium of Neuropsychological Tests* (Oxford University Press, 2006).
51. D. J. Thal, E. Bates, J. Goodman, J. Jahn-Samilo, Continuity of language abilities. *Dev. Neuropsychol.* **13**, 239–273 (1997).
52. D. J. Thal, L. O'Hanlon, M. Clemmons, L. Fralin, Validity of a parent report measure of vocabulary and syntax for preschool children with language impairment. *J. Speech Lang. Hear. Res.* **42**, 482–496 (1999).
53. D. Thal, J. L. DesJardin, L. S. Eisenberg, Validity of the MacArthur-Bates communicative development inventories for measuring language abilities in children with cochlear implants. *Am. J. Speech Lang. Pathol.* **16**, 54–64 (2007).
54. L. Fenson *et al.*, Variability in early communicative development. *Monogr. Soc. Res. Child Dev.* **59**, 1–185 (1994).
55. J. Villar *et al.*, Neurodevelopmental milestones and associated behaviours are similar among healthy children across diverse geographical locations. *Nat. Commun.* **10**, 1–10 (2019).
56. E. Bergelson, Bergelson Seedlings HomeBank Corpus. HomeBank. <https://doi.org/10.21415/T5PK6D>. Accessed 22 November 2023.
57. R. Brookman *et al.*, Depression and anxiety in the postnatal period: An examination of infants' home language environment, vocalizations, and expressive language abilities. *Child Dev.* **91**, e1211–e1230 (2020).
58. M. Canault, M.-T. Le Normand, S. Foudil, N. Loundon, H. Thai-Van, Reliability of the Language Environment Analysis system (LENA™) in European French. *Behav. Res. Methods* **48**, 1109–1124 (2016).
59. A. Cristia, M. Casillas, LENA recordings gathered from children growing up in Rossel Island. OSF. <https://osf.io/juys6/>. Accessed 22 November 2023.
60. H. Elo, "Acquiring language as a twin: Twin children's early health, social environment and emerging language skills," PhD thesis, Tampere University (2016). <https://urn.fi/URN:ISBN:978-952-03-0296-2>.
61. H. Ganek, A. Eriks-Brophy, LENA its data from daylong recordings collected in Vietnam. OSF. <https://osf.io/d9453>. Accessed 13 November 2023.
62. L. Hamrick, A. Seidl, B. L. Tonnsen, LENA its data from daylong recordings gathered from children with typical and atypical development. OSF. <https://osf.io/n9pvq/>. Accessed 13 November 2023.
63. E. Kidd, C. Junge, T. Spokes, L. Morrison, A. Cutler, Individual differences in infant speech segmentation. *Infancy* **23**, 770–794 (2018).
64. E. Marklund, I.-C. Schwarz, F. Lacerda, LENA its data from daylong recordings in Swedish-speaking families with 3- to 10-month-olds (recorded 2016). OSF. <https://osf.io/wh9dt>. Accessed 13 November 2023.
65. K. McDivitt, M. Soderstrom, McDivitt HomeBank Corpus. HomeBank. <https://doi.org/10.21415/T5KK6G>. Accessed 13 November 2023.
66. N. Ramirez-Esparza, A. García-Sierra, P. K. Kuhl, Look who's talking: Speech style and social context in language input to infants are linked to concurrent and future speech development. *Dev. Sci.* **17**, 880–891 (2014).
67. N. Ramirez-Esparza, A. Garcia-Sierra, P. K. Kuhl, The impact of early social interactions on later language development in Spanish-English bilingual infants. *Child Dev.* **88**, 1216–1234 (2017).
68. C. F. Rowland, A. Bidgood, S. Durrant, M. Peter, J. M. Pine, The Language of 0-5 Project. OSF. <https://osf.io/kau5f/>. Accessed 13 November 2023.
69. C. Scaff, J. Stieglitz, A. Cristia, Tsimane' daylong recordings collected with LENA in 2017–2018. OSF. <https://doi.org/10.17605/OSF.IO/6NEZA>. Accessed 13 November 2023.
70. I.-C. Schwarz, E. Marklund, T. Gerholm, LENA its data from daylong recordings in Swedish-speaking families with 30-month-olds. OSF. <https://osf.io/yzp4b>. Accessed 13 November 2023.
71. I.-C. Schwarz, E. Marklund, C. Lam-Cassetari, U. Marklund, Longitudinal LENA its data from daylong recordings in Swedish-speaking families with infants at 6, 12, 16 and 24 months. OSF. <https://osf.io/38arg>. Accessed 15 November 2023.
72. P. Van Alphen, M. Meester, E. Dirks, LENA onder de loep; ITS files and metadata of daylong LENA recordings at the homes of preschoolers with DLD and TD peers (collected by the Royal Dutch Kennalis and the NDSK). OSF. <https://osf.io/2zyub>. Accessed 13 November 2023.
73. M. VanDam, VanDam Cougar HomeBank Corpus. HomeBank. <https://doi.org/10.21415/T5WT25>. Accessed 13 November 2023.
74. A. S. Warlaumont, G. M. Pretzer, S. Mendoza, E. A. Walle, Warlaumont HomeBank Corpus. HomeBank. <https://doi.org/10.21415/T54S3C>. Accessed 13 November 2023.
75. A. Weber, V. A. Marchman, A. Fernald, LENA its data collected in Kaolack Senegal in 2013. OSF. <https://doi.org/10.17605/OSF.IO/EMFB5>. Accessed 13 November 2023.
76. A. Weisleder, A. Mendelsohn, Daylong recordings of 2–12 month-old infants from Spanish-speaking homes in the US. OSF. <https://doi.org/10.17605/OSF.IO/JBINC>. Accessed 13 November 2023.
77. P. Van Alphen, N. Davids, E. Dijkstra, P. Fikkert, TiBENA: ITS files and metadata of daylong LENA recordings at the homes of preschoolers with DLD and TD peers (collected by the Royal Dutch Kennalis and the Radboud University). OSF. <https://osf.io/yymv7b>. Accessed 13 November 2023.
78. M. H. Bornstein, C.-S. Hahn, J. T. D. Suwalsky, O. M. Haynes, "Socioeconomic status, parenting, and child development" in *Socioeconomic Status, Parenting, and Child Development, Monographs in Parenting Series*, M. H. Bornstein, R. H. Bradley, Eds. (Lawrence Erlbaum Associates Publishers, Mahwah, NJ, 2003), pp. 29–82.
79. D. Bates, M. Mächler, B. Bolker, S. Walker, Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
80. D. Bates, R. Kliegl, S. Vasishth, H. Baayen, Parsimonious mixed models (2018). arXiv [Preprint] (2022). <https://arxiv.org/abs/1506.04967> (Accessed 29 April 2022).
81. P. S. Dale, The validity of a parent report measure of vocabulary and syntax at 24 months. *J. Speech Lang. Hear. Res.* **34**, 565–571 (1991).
82. H. M. Feldman *et al.*, Concurrent and predictive validity of parent reports of child language at ages 2 and 3 years. *Child Dev.* **76**, 856–868 (2005).
83. T. Velikonja *et al.*, The psychometric properties of the Ages & Stages Questionnaires for ages 2–2.5: A systematic review. *Child: Care Health Dev.* **43**, 1–17 (2017).
84. D. Bricker *et al.*, *Ages and Stages Questionnaire* (Paul H Brookes, Baltimore, MD, 1999).
85. L. C. H. Fernald, E. Prado, P. Kariger, A. Raikes, A toolkit for measuring early childhood development in low and middle-income countries (Ministerio de Educacion, 2017). Repositorio Institucional del Ministerio de Educación. <https://repositorio.minedu.gob.pe/handle/20.500.12799/5723>. Accessed 11 May 2022.
86. M. Lavechin, R. Bousbib, H. Bredin, E. Dupoux, A. Cristia, An open-source voice type classifier for child-centered daylong recordings. *Interspeech*. arXiv [Preprint] (2020). <http://arxiv.org/abs/2005.12656> (Accessed 11 September 2020).
87. H. Hammarström, R. Forkel, M. Haspelmath, S. Bank, Glottolog 4.2.1 (MPI for the Science of Human History, Jena) (2020). <https://glottolog.org/>. Accessed 4 June 2020.
88. E. Bergelson *et al.*, LENA Broad Strokes. Open Science Foundation (OSF). https://osf.io/9v2m5/?view_only=50df17fc0844145ae692c3578c6b08. Deposited 28 August 2023.