Environmental influences on infants' native vowel discrimination:
The case of talker number in daily life

Anonymized

Anonymized

Abstract

Both quality and quantity of speech from the primary caregiver have been found to impact language development. A third aspect of the input has been largely ignored: the number of speakers who provide input. Some infants spend most of their waking time with only one person, while others hear many different speakers, for example in daycare. Children receive very different input as the number of people they spend their time with varies. Even if the very same words are spoken the same number of times, the pronunciations can be more variable when several talkers pronounce them. The question then arises whether language acquisition is affected by the number of people who provide input. Laboratory studies testing how infants learn from multiple talkers have yielded all possible patterns: Some studies suggest no effect, others a negative one, and yet others a positive one. Extrapolating to language acquisition from such diverse findings is difficult because (a) they arise from various paradigms and tasks, and (b) they may not reflect acquisition as it happens naturally outside the lab. In the present work, we sought to shed light on the possible link between how many people provide input in daily life and infants' native vowel discrimination. We tested three age groups: 4-month-olds (before attunement to native vowels), 6-month-olds (at the cusp of native vowel attunement) and 12-month-olds (well attuned to the native vowel system). We find no relationship between input variability and native vowel discrimination skills in 4- and 6-month-olds, who are overall able to discriminate the vowel contrast. At 12 months, we observe a small positive relationship, but further analyses reveal that the data are also compatible with the null hypothesis of no relationship. We discuss implications of our findings in the context of infant language acquisition and cognitive development.

Word count: 8400

Forming discrete categories based on continuous input is a fundamental human skill, and it takes places across many domains. One instance is visual perception, where, for example, golden retrievers and Chihuahuas can be grouped together into a "dog" category.

Even before infants acquire object labels, they are remarkably capable of building such visual categories (Quinn, Eimas, & Rosenkrantz, 1993). Another instance of category formation takes place during early language acquisition, when infants tune into the sound system of their native language. There are many possibilities to divide up the acoustic space into units that are combined to form distinct word forms. Infants have to discover which sound contrasts (such as /r/ and /l/, which change word meaning in English, but not Japanese) to pay attention to. This process takes place early in infancy, as infants begin to show indicators of a native vowel system at the age of six months (see the meta-analysis by Tsuji and Cristia 2014).

Given the timeline, mechanisms that can account for infants' ability to build sound categories must work from the signal, involving limited or no top-down information. Consequently, all currently proposed mechanisms of sound category acquisition rely heavily on the computation of statistics over some input representations. This includes distributional learning (Maye, Werker, & Gerken, 2002); the perceptual magnet effect (Kuhl, 2004); and proposals that allow for interactions with other sources of information, such as the developing proto-lexicon (e.g., Elsner, Goldwater, Feldman, and Wood 2013; Feldman, Griffiths, Goldwater, and Morgan 2013; Swingley 2009; Yeung and Werker 2009). Infants must both draw the lines between native sound categories and discover how many categories actually exist; clearly a difficult joint learning problem.

It is possible that speaker characteristics are a hurdle during language processing, and specifically sound categorization, because talkers vary a great deal in their productions along the same acoustic dimensions that distinguish speech sounds. As a result, one person's /a/ (as in "cot") might be indistinguishable from another's /ɔ/ (as in "caught") on the acoustic level. Speaker differences can be traced to a number of factors ranging from physical differences to idiosyncratic ways of articulating sounds.

If infants were to track the statistics of raw acoustic realizations across salient dimensions, they might be misled as to their native category system. Speaker specific variation can mask linguistically relevant information, and the overlap between categories increases dramatically when introducing multiple speakers (Hillenbrand, Getty, Clark, and Wheeler 1995; specifically Fig. 4). Alternatively to computing possibly misleading statistics over all speakers in the environment, infants might be able to separate talkers (Kleinschmidt & Jaeger, 2015). How could they go about doing so? One option is to use contextual information, so that statistics within speakers are tracked independently. This should lead to less information being available for the formation of each speaker-dependent category when being exposed to multiple people (assuming that equal amounts of speech input is available across low- and high-talker-variability scenarios). The resulting prediction in both scenarios is that increased talker variation has *negative* consequences for sound category development.

Not all researchers who work on the topic of variability and early language acquisition agree with these negative predictions. A sizable community believes that speaker variability could play no role at all. This position rests on the assumption that, from birth, infants are able to ignore the difference between speakers, by either abstracting away features that are confusing, or computing speaker-invariant representations early on during processing (Dehaene-Lambertz & Peña, 2001; Kuhl, 1979, 2004). With this innate ability in place, talker variation should have *no impact* on sound category development.

Yet others argue that talker variability might be *helpful* during sound category acquisition. This is the case if variability helps infants focus on linguistically informative dimensions of the speech signal (Rost & McMurray, 2009; Seidl, Onishi, & Cristia, 2014). The core assumption is that infants base their learning on a representation where speaker-dependent information is not (or only weakly) correlated with linguistically important information. As a consequence, infants who are exposed to more variable input should benefit, because they learn to ignore irrelevant information and start learning the abstract system governing their ambient language faster.

Empirical support exists for all three stances and we will give a brief overview here (for a more extensive review focusing on sound discrimination, see Bergmann, Cristia, and Dupoux 2016). For example, Jusczyk, Pisoni, and Mullennix (1992) showed a negative effect of multiple talkers for two-month-olds in a sound discrimination task, especially when introducing a short delay that likely increases task demands. In contrast, Kuhl (1979) found that six-month-olds succeeded in distinguishing two vowels, both in the presence and absence of multiple talkers. The third scenario was supported by a study showing that, when learning phonotactic rules, four- and eleven-month old infants benefit from hearing multiple talkers during the learning phase (Seidl et al., 2014). While the diversity in the findings just mentioned could be due to methodological and/or age differences, across the literature it is not possible to discern a clear pattern which is modulated by those factors (Bergmann et al., 2016). More importantly, it is impossible to know for a fact that these laboratory results represent precisely categorization and learning as they occur in daily life. Indeed, laboratory experiments are only *one* way to understand how early life experiences shape infants' language acquisition.

Another approach that does not rely solely on short-term learning and performance in the lab involves individual variation among infants. This line of research relies on studying differences between (groups of) infants by tracking both diverging features of infants' environments and their performance in fairly standard lab-based tasks. Correlations between the two measures are compatible with the interpretation that there may be an underlying association. This method has prominently been applied when investigating the possible link between language abilities and the quantity of infants' input (e.g., Huttenlocher, Haight, Bryk, Seltzer, and Lyons 1991) as well as the role of home environment quality (Melvin et al., 2016).

On a more subtle level, qualitative aspects of infants' input also shape their path into language. For example, mothers' vowel space size when talking to their infant aged 6 to 12 months correlated with the infants' performance in a speech perception task at the same age (Liu, Kuhl, & Tsao, 2003). In another study documenting a more specific association, differences in how distinct /s/-/ʃ/ (as in "sip" versus "ship") were in a group of caregivers predicted their respective infants' ability to discriminate the same sound pair beyond overall speech rate or pitch (Cristia, 2011). The latter study suggests that the more confusing infants' input, the less likely they are to perform well when having to distinguish two similar sounds in laboratory studies. In other words, the effects of characteristics of infants' everyday environment on language development may be measurable via laboratory tasks.

Our study uses the individual variation method to assess the key question of whether talker variation in infants' everyday life could potentially affect early language acquisition.

If all infants received input from just one main caregiver, it would not be possible to observe any differences. However, in many cultures, different daycare models and diverse household structures (e.g., with multiple siblings or including extended family) lead to natural variation in the number of people who habitually talk to an infant.

The present work is the first (to our knowledge) to systematically assess the role that talker variability plays in infant language acquisition. We sought to shed light on the possible link between how many people provide input and how infants zoom into their native vowel system. We chose a vowel contrast because speakers are expected to differ more in vowel than consonant articulation, as the former are greatly impacted by variation in vocal tract length and structural configuration (Hillenbrand et al., 1995). We specifically selected a vowel pair previously described as difficult to distinguish for young infants (Pons, Albareda-Castellot, & Sebastián-Gallés, 2012), which also varies greatly across French adult speakers (Gendrot and Adda-Decker 2005, Table 5), and used the implementation of this contrast in a single speaker, so as to focus native vowel discrimination (rather than other abilities potentially affected by experience with talker variability, such as cross-speaker normalization). To estimate talker variability in infants' input, we asked parents about their child's schedule during a typical week. We tested three age groups to track development, namely 4-month-olds (before attunement to native vowels), 6-month-olds (at the cusp of native vowel attunement) and 12-month-olds (well attuned to the native vowel system). As outlined above, there are three mutually exclusive predictions on the impact of speaker variability. If talker variation hinders acquisition of native vowel categories, we should observe a negative association between the number of speakers in the environment, and infant performance in a vowel discrimination task. A positive association would be consistent with views stating that infants use talker variation to determine what acoustic changes are linguistically relevant. Finally, no association would be most compatible with views that infants ignore, or automatically compensate for, talker variation.

Experiments

The present study was preregistered in two steps on Open Science Framework prior to data collection. The pre-registrations are available on the project website, along with all stimuli, analysis scripts, anonymized data, and supplementary materials: `https://osf.io/q9cpa/?view_only=230c57313df04d0fb70e3907eaf94c07`. Details on the pre-registration, all exploratory analyses conducted, and documentation of all changes in data processing with respect to pre-registration can be found on the project website. The experiments were approved by local ethical authorities.

Participants

Birth records and contact information are provided by official sources to the babylab; about three months after their child's birth parents receive a letter describing the babylab. Parents who respond to this letter are added to the lab database, and those whose child matches the pre-set age criteria were then contacted for this experiment. All children were, according to parental reports, monolingual speakers of French (at least 90% exposure) and born full-term.

135 In each age group, we included 46 participants in our final sample (girls per age group: 27 4-month-olds; 21 6-month-olds, 19 12-month-olds). The mean ages (along with minimum and maximum age) in days were 138 (121 to 151) for the 4-month-olds, 180 (154 to 211) for the 6-month-olds, and 358 (339 to 392) for the 12-month-olds.

  To arrive at the final sample we had to test 75 additional infants, who were excluded
140 for the following reasons: audible crying (17 4-month-olds; 10 6-month-olds; 8 12-month-olds), fussiness (14 4-month-olds; 4 6-month-olds; 1 12-month-olds), less than 1.5s looking time in any of the test trials (4 4-month-olds; 3 6-month-olds), and parental interference (4 4-month-olds; 8 6-month-olds; 2 12-month-olds).

  The children in our final sample came from a homogeneous sample, growing up in an
145 urban setting and with mothers who had all completed high school and, with two exceptions, attended university for several years (number of years in higher education: M = 5.2, SD = 1.99).

Stimuli

  The acoustic stimuli were two native vowels /i/ and /e/, presented as nuclei in syl-
150 lables with /g/ as the onset. The resulting monosyllabic minimal pair is probably low frequency in infants' input (/gi/ is equivalent to the name "Guy", none of our participants in the final sample was called Guy; /ge/ means "gay"). The resulting CV syllables /gi/ and /ge/ were recorded by a native female speaker of French in a lively manner, as if speaking to a young infant, and thus contained variable intonation patterns. We created three syllable
155 lists, one for habituation and two for test, each containing seven unique tokens separated by a 1s inter-stimulus interval for a total duration of 20s (vowel duration was matched across the three lists). The test syllables were matched across lists for their intonational pattern so that there was equal intonational variation in the /i/ and /e/ lists.

  The average formant values in Bark for the habituation and test lists are: Habituation
160 list /gi/ F1 = 3.37, F2 = 14.70; Test list /gi/ F1 = 3.36, F2 = 14.70; Test list /ge/ F1 = 5.13, F2 = 14.29. T-tests confirmed that the habituation and test items of /gi/ did not differ significantly ($p = .95$). We further confirmed that voice onset time and prevoicing duration in [g] across both test lists were not significantly different with two t-tests (both $p > .7$).

165 ### Procedure

  Infants sat in a sound-proof booth on their parent's lap approximately 60 cm away from a wall-mounted 27" monitor (type Iiyama PROLITE E2773HS-GB1). Sound was presented from two speakers next to the monitor (model JBL Control 1Pro). Parents were instructed not to talk or point at the screen and listened, via noise-canceling headphones,
170 to masking music which was overlaid with the experimental stimuli in random order and varied intensity. Via a camera above the screen, infants' reactions were observed by the experimenter, who controlled the experiment outside the sound-proof booth without being aware of the trial type currently being presented. The experimenter wore noise-canceling headphones to further ensure blinding.

175 We presented our experiment using the Lincoln Lab LOOK software (Meints & Woodford, 2008). The experimenter recorded on-line whether the infant looked to the screen or

not via a button press. Trials ended early when infants looked away from the screen for more than two consecutive seconds, and a new trial only started when infants looked back towards the screen for at least one second. On the screen, infants saw a bull's eye on a gray background during the trials, and moving colorful shapes on a black background between trials to redirect their attention to the screen. All infants were tested by the same experimenter (the first author).

The habituation criterion was set to 50% of the average of the three longest looking times. The maximal trial duration was the duration of the files, 20s. After completing the habituation phase with the background syllable /gi/, which could last up to 24 trials, infants heard one trial with the novel yllable /ge/, followed by two trials with new tokens of the habituated syllable /gi/ and a final trial with the novel syllable /ge/. No trials were repeated.

Before the first habituation trial and after the last test trial, we presented an attention trial to establish an independent criterion for infants' fussiness. To that end, infants saw a smiling baby's face on a grey background accompanied by "coucou" (*hi*) for maximally 20s. The same criteria (trial onset after looking to the screen for one second; trial end criterion after looking away for two consecutive seconds) were applied.

Questionnaire

To assess infants' input variability in daily life we used a questionnaire which parents filled out in the lab (see the project website). On this questionnaire parents were asked to complete a schedule of a "typical week" and fill in for each day's morning, afternoon, and evening who talked to the child for more than 20 consecutive minutes. This criterion aimed to exclude brief encounters to simplify the parents' task. If there were regular visitors or caretakers, but they could vary across weeks, parents marked this as well. A second sheet asked for details on each speaker or group of changing speakers, including sex, age (range), highest educational degree (following French conventions often reported in or converted to years post high school), and if they spoke with a non-local or non-native accent. If a change in the infant's routine had taken place in the last two months, parents filled out a second schedule and ensured that all speakers were mentioned in the sheet noting their details. The main analyses focus on the number of speakers older than two years the infant is currently exposed to. Interested readers can find additional analyses and the raw data on the project website.

Results

Preliminary analyses of the 6-month-olds' data suggested that we had not considered a number of factors in the pre-registration, most saliently the shape of the distribution of number of speakers. For the results of the pre-registered analyses and detailed discussion, see the supplementary materials on the project website. We report here analyses that deviate from the pre-registered ones but were more appropriate.

The number of non-infant speakers that parents reported as talking to the child in a typical week ranged from two to 27, with a median of seven for the 4-month-olds and eight for the other two age groups, as depicted in Figure 1. Since the distributions were not normal, this predictor was log-transformed. Note that log transformation coincides with

the intuition that differences in lower numbers (e.g. between two and four input speakers)
matter more than differences in larger numbers (e.g. between 22 and 24 input speakers).
We further opted to not split the data by their median, as some other individual differences
studies do, since this analysis would have yielded very uneven group sizes, as displayed in
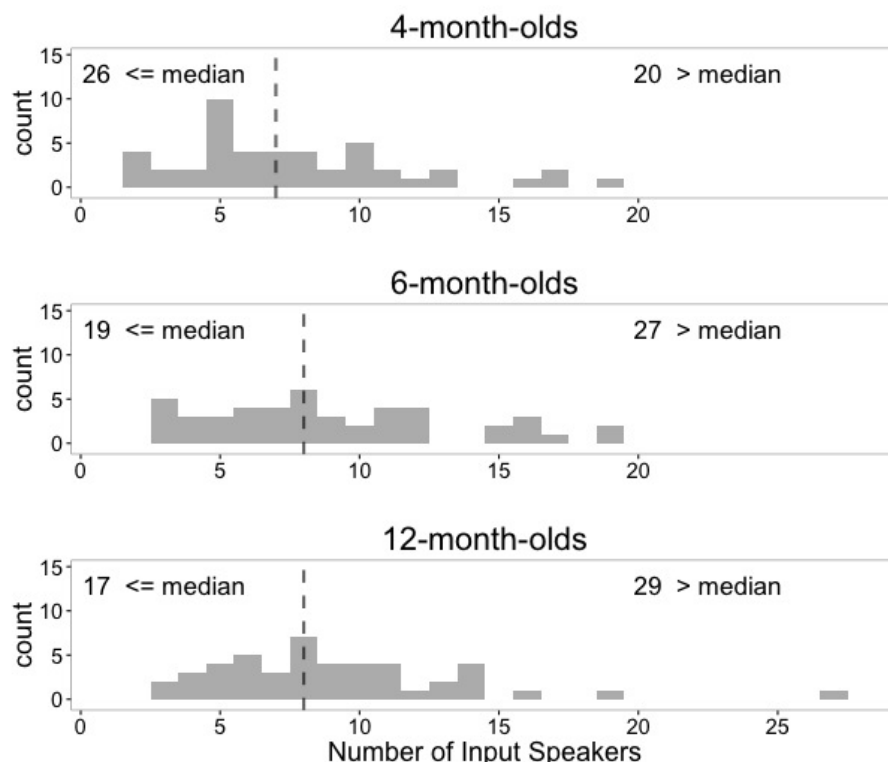Figure 1.

[Insert Figure 1 about here.]

Figure 1. **Histogram of number of talkers.** The distribution of the number of talkers in
infants' input, separated by infant age group. The dashed line indicates the median number
of speakers. The number of participants which either have as many or fewer talkers in their
input as the group median is noted in the top left of each panel, the number with more
talkers than the group median is indicated on the top right.

Turning to the vowel discrimination study, infants needed on average 12.17 (SD =
5.69) trials to habituate, with a range of 4 (the preset minimum) to 24 (the preset maximum). A linear regression with number of trials to habituation as dependent measure and
age group and log number of speakers in the input as predictors did not reveal any main
effects or interactions (cf. supplementary materials on the project website).

The dependent variable in the laboratory experiment was a discrimination score,
infants' listening time to novel test trials divided by their total listening time to all (novel
and habituated) test trials. This score can range between 0 and 1, with .5 indicating no
difference between listening times to novel and previously habituated test trials. The mean
discrimination score across all age groups is .533, which is significantly above .5 (one-sample
against chance of .5 $t(137) = 3.0188$, two-tailed $p = .003$, 95% CI [.511, .554]).

To answer our key research question, we fit a linear regression declaring discrimination score as dependent measure, and age group and log transformed number of speakers in the input, as well as their interaction, as predictors. All analyses were done in R using the package "car" (Fox & Weisberg, 2011; R Core Team, 2016). Age group was a significant predictor with $F(2,132) = 5.205$, $p = 0.007$. The effect of age group emerged because younger infants performed better than older ones (all t-tests are one-sample against .5 chance, all p's are two-tailed): 4-month-olds $t(45) = 2.326$, $p = .024$, 95% CI [.505, .576]; 6-month-olds $t(45) = 4.175$, $p < .001$, 95% CI [.535, .601]; 12-month-olds $t(45) = -.533$, $p = .596$, 95% CI [.448, .530]. Neither the log transformed number of speakers ($F(1,132)$ = 1.396, $p = 0.240$) nor the interaction of the two predictors ($F(2,132) = 2.529$, $p = 0.084$) reached significance. Bayesian analyses[1] confirmed that age group alone predicts discrimination scores ($BF_{10} = 4.026$), and that adding log number of input speakers does not add explanatory power.

Although no interaction was found, in the interest of fully describing our data for readers, we carried out correlations of discrimination scores with log transformed number of speakers within each age group, as depicted in Figure 2. For the 4-month-olds, this correlation is not significant and positive ($r = .106$, CI: [-.189, .384], $p = .482$, $BF_{10} = 0.147$), for the 6-month-olds non-significantly negative ($r = -.126$, CI: [-.402, .170], $p = .403$, $BF_{10} = 0.163$), and for the 12-month-olds significantly positive ($r = .303$, CI: [.014, .545], $p = .041$, $BF_{10} = 0.920$). All BF_{10} for correlations were calculated based on Wetzels and Wagenmakers (2012).

[Insert Figure 2 about here.]

Discussion

We set out to investigate whether the number of speakers in daily life, a proxy of speaker variation in the input, predicts infants' developing ability to discriminate native vowels. Our findings show no overall relationship between discrimination ability and input variability. Follow-up analyses revealed that there was no significant link at the ages of four and six months, and a significant positive correlation at twelve months. These results cannot be viewed as confirming or contradicting previous research, since this is the first study measuring potential long-term effects of variation in infants' natural experience. Before discussing the theoretical implications of these results, we assess some potential interpretations.

We start with the non-significant relationships. Typically, a null result is viewed as inconclusive as it could follow from (a) the absence of an effect; or (b) the existence of an effect paired with failure from the experimenter to detect it. Would we have measured an effect with statistical reliability with more precise measures and increased power? We believe the answer is no. First, our sample was appropriately large to find effects similar in size to previous work. Indeed, we tested 46 children per age group, increasing sample sizes

---

[1] Bayes Factors (BF) indicate how much more probable the data were under the one hypothesis compared to a second one hypothesis. We report BF_{10}, comparing the data to their compatibility with the alternative of an (undirected) correlation. The BF_{01} indicating the probability of the observed data compared to the null hypothesis can be found in the supplementary materials, along with explanatory figures. All BF were computed in R using the package BayesFactor (Morey & Rouder, 2015), supplementary computations and illustrations used JASP (JASP Team, 2016).
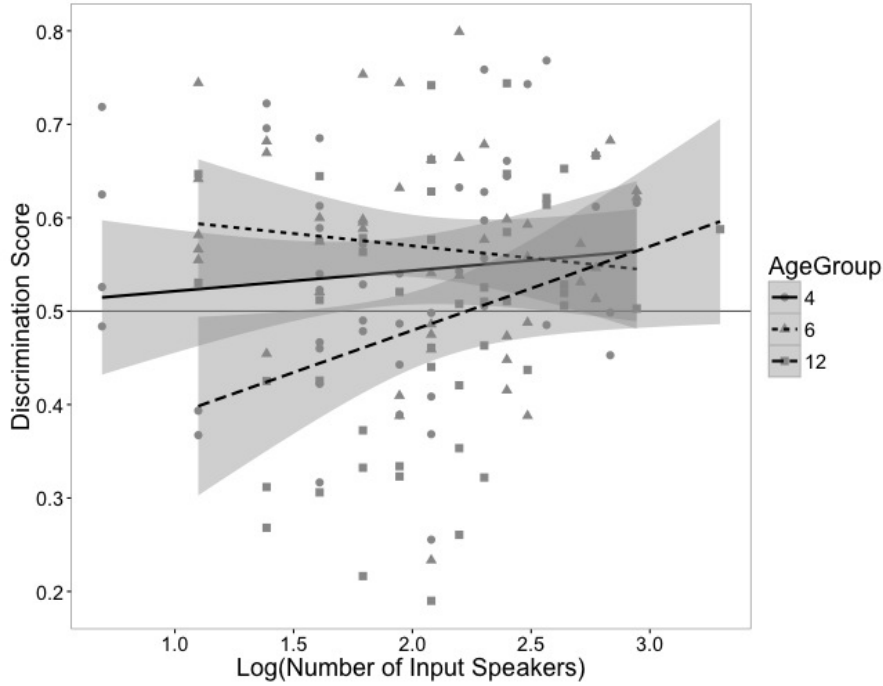
Figure 2. **Discrimination scores as a function of log-transformed number of talkers in infants' input.** Linear regression lines are superimposed, along with the respective 95% confidence intervals. Each line style and symbol corresponds to one age group; every symbol indicates one participant. The horizontal gray line marks a discrimination score of .5 (no looking preference).

from typical vowel discrimination studies, where around 23 infants are necessary to pick
275 up a general discrimination effect with 80% power (Tsuji & Cristia, 2014). Moreover, the
sample, totaling 138 infants, is also much larger than those found in previous work looking
at individual variation in infant sound discrimination (Cristia, 2011; Liu et al., 2003; Melvin
et al., 2016). The experimental method we used is also firmly rooted within the existing
literature. The habituation-dishabituation method has been used across many research
280 topics, including native and non-native vowel discrimination (Tsuji & Cristia, 2014). The
outcome measure (discrimination score) follows common practice in the recent literature on
vowel discrimination (e.g., Pons et al. 2012; see Tsuji and Cristia 2014).

The failure to detect an effect might not only be caused by the method but could also
be due to the stimuli used in the present work. Previous research pointed to the contrast
285 chosen being acquired between four and twelve months, with the former age failing to
discriminate it (Pons et al., 2012). Since we wanted to be able to measure variance in infants'
discrimination abilities without a ceiling effect, we selected this contrast. Contrary to what
one would have expected following previous work, even 4-month-olds showed an overall
discrimination response. The fact that our 4-month-olds were able to discriminate the two
290 sounds speaks against this task being overly difficult, and thus prevents an interpretation
of lack of correlation due simply to floor effects. As Figure 2 shows, there is considerable
variation in responses measured within each age group, variation that speaks against a

ceiling interpretation.

Finally, the measurement of infants' input variability has to be carefully considered
as a possible cause for a null finding. The measure of infants' environment and input,
derived from a questionnaire parents filled out in the lab, was specifically designed and
piloted for this study. Thus, it lacks extensive validation, for example with observations at
home. Validation would depend on data that is time-consuming to generate, either for the
parents, who could complete a dense diary, or for the experimenter who analyzes daylong
recordings, where trained listeners annotate speakers in a process that takes typically at
least as long as the recording itself (here, we would require several full days). Without prior
evidence that speaker variation plays an important role for infant language acquisition, it
may not be justified to allocate substantial human resources to this task. Therefore, we
chose to rely on parents' subjective judgment, and provided detailed written instruction and
assistance from the experimenter while filling out the questionnaire. Although responses
may vary between participants, the questions asked were clear and general enough that it is
unlikely they biased parents' responses. We are thus confident that we capture the typical
input of a given infant in the time frames we analyzed (a typical week, based on slots of
several hours per day). The dependent variable chosen to capture infants input, the number
of speakers in a typical week, might not be covering the time-frame that most impacts
infants' language acquisition on the sound-level. Other possibilities have been explored in
supplementary analyses based on the same questionnaire data that were used for the main
results, and none added explanatory power to the statistical models. We tested, among
other things, whether the presence of a stable speaker and the average number of people
present at any one time have an effect. Therefore, we can rule out other plausible measures
of input variability that could predict a difference in infants' language development. More
fine-grained measures, e.g., based on voice similarity, would rely on daylong recordings like
those just described, which are arguably too costly to obtain for a study exploring a new,
and hitherto unstudied, factor in language acquisition.

Overall, both the laboratory experiment and the questionnaire were designed with
care and rigor sufficient to obtain a meaningful result to address our key research question.
Taken together with the high sample size in the present work, we thus consider our mea-
sures a solid basis for our conclusions and continue to analyze and interpret our findings
accordingly.

We presented both common frequentist analyses of our data, including *p*-values, and
the emerging Bayes Factors (Wagenmakers, Morey, & Lee, 2016). The two types of statistics
complement each other in the description of our data. In the two younger age groups, we
found that the *p*-values for our key factor, talker number, did not reach the significance
threshold. From that alone, however, we cannot conclude that there is no relationship
between variability in the input and vowel discrimination. Bayes Factors can help illuminate
whether the measure was insensitive (due to noisy data and/or lack of power) or whether
the observed data are in line with the null hypothesis – in the present case implying that
there is no relationship between the number of talkers in infants' input and their ability to
discriminate native vowels. Along with the small correlation coefficients, the Bayes Factors
for the 4- and 6-month-olds provide evidence for the null hypothesis that the amount of
speakers an infant is exposed to does *not* affect vowel discrimination in our laboratory task
(at least for the vowels we have chosen and in the way we tested infants).

For the 12-month-olds, where we found a significant positive correlation, the Bayes Factor is near 1, implying that from a Bayesian perspective the evidence is inconclusive and the data are also compatible with the null hypothesis of no relationship between talker number and infant native vowel discrimination. Additionally, we believe that these results are not particularly compelling, as they seem to be driven partially by below-chance performance found among children hearing few talkers. While this pattern could be interpreted meaningfully in a fixed-familiarization study (as a familiarity preference), below-chance performance is impossible to interpret in a habituation study like ours. Moreover, as a group, 12-month-olds did not show above-chance performance, suggesting that a task that was reasonably easy at 4 and 6 months of age was not as easily solved at twelve months. In other words, although our experiment as a whole was not inappropriately built (given that even the 4-month-olds showed above-chance performance), it may have been pray to floor effects for the oldest age group. On the whole, although the correlation between performance and talker number in this age group is significant, there are sufficient reasons, both conceptual (group effect at chance, some individuals' performance inexplicably below chance) and empirical (Bayes factor at 1), to disfavor a strong interpretation whereby talker number and speech perception are robustly related.

Taking all age groups together, what are the theoretical implications of these results? In the introduction, we laid out three possibilities of input variability influencing language development, each supported by a theoretical framework and by experimental findings that pertained to infants' ability to deal with speaker variation in the laboratory in various learning and discrimination tasks. We found no evidence for a negative impact, which would have been predicted by accounts basing language acquisition on statistical computations over the raw acoustic input (Maye et al., 2002). There are two possible learning scenarios compatible with our finding of no negative impact of talker number on native sound discrimination.

The first scenario builds on an abstractionist account (Kuhl, 2004), where innate or early-acquired biases lead infants to separate talker-specific from linguistic information early on, and they thus acquire their native sound system using representations that are talker-invariant. One issue this scenario must solve is that abstraction is difficult to reconcile with current knowledge of human sound processing and automatic speech recognition attempts to emulate these abilities. As to the former, even adults are affected in their sound processing abilities by speaker variation and change (Creel, Aslin, & Tanenhaus, 2008; Kleinschmidt & Jaeger, 2015), and acoustic representations that take into account talker variation better predict human adult perception than raw representations (Richter, Feldman, Salgado, & Jansen, 2016). As to the latter, attempts to compute a representation of the acoustic signal that allows for the accurate classification of different sound categories across talkers point to the need of either substantial knowledge of the target sound categories or large computing resources.

According to the second scenario, infants might be able to *adapt* to talkers' voices, and learn sounds and words from them. Adaptation is a process by which representations and expectations are adjusted to the current situation. This process can take place on many different linguistic levels and be driven by various cues. In adult speech processing, the most common account holds that listeners use the lexical context to infer the target sounds, and adjust for differences between expected and observed sounds on the phonological and/or

lexical levels (Creel et al., 2008; Dahan, Drucker, & Scarborough, 2008; Kleinschmidt & Jaeger, 2015; McQueen, Cutler, & Norris, 2006). Which cues adult listeners employ, what level of representation is being adjusted in which way, and what factors modulate adaptation are subject to ongoing debate. This is also the case for toddlers, who are thought to employ a range of more or less linguistically-informed heuristics to adapt to unexpected pronunciations, using lexical context where available (Mulak, Best, Tyler, Kitamura, & Irwin, 2013), but also simply becoming more accepting of variable pronunciations when prompted with socially diverse visual cues (for a discussion see Schmale, Seidl, and Cristia 2015).

Independent of which mechanisms listeners employ, adaptation requires at least some short exposure phase to a specific speaker for the listener to be able to adjust to their specific voice and pronunciations. As yet, this ability has virtually not been assessed in young infants. Indeed, studies aiming to investigate how talker differences affect infants' perception aimed at blocking adaptation. Typically, syllables or words spoken by multiple talkers are presented in short succession (e.g., Dehaene-Lambertz and Peña 2001; Jusczyk et al. 1992; Kuhl 1979; Polka, Masapollo, and Ménard 2014; Seidl et al. 2014) or a talker change is introduced and infants' reactions are measured immediately (Houston & Jusczyk, 2000; van Heugten & Johnson, 2012). It would be interesting for future work to explore to role of adaptation in infants' sound discrimination, as other experimental evidence suggests that all pre-requisites are in place from early on, as follows.

To be able to adapt to talkers, infants must at least be able to pick up on visual and environmental cues and/or be sensitive to different voice characteristics that do not contain linguistic information but transport talker-identifying information. As to visual cues, infants can already match male and female voices to faces at two months (Patterson & Werker, 2003). Additionally, visual cues, such as lip movements, influence infants' sound learning already at six months (Teinonen, Aslin, Alku, & Csibra, 2008), showing the role of environmental information and infants' abilities to exploit it during and for language processing.

Moreover, even top-down, lexically-informed adaptation might be available within the first year of life. Both empirical findings and emerging theories suggest an interplay between word-level and sound-level processing throughout language acquisition. Infants as young as six months old have been shown repeatedly to know a few common words (Bergelson & Swingley, 2012; Mandel, Jusczyk, & Pisoni, 1995; Tincoff & Jusczyk, 1999), which could serve as anchor. Emerging formal accounts show that discrimination of a sound contrast is facilitated when the two sounds are used in different wordforms (Feldman et al., 2013) or associated with different objects (Gogate, Prince, & Matatyaho, 2009; Yeung & Nazzi, 2014; Yeung & Werker, 2009).

To summarize, together with the consideration of adult speech processing models and insights from automatic speech recognition mentioned previously, our results lead to the exciting proposal that infants' speech perception abilities crucially involve various mechanisms for adaptation. Such a proposal invites laboratory studies assessing infants' abilities to adapt to talkers when acquiring and adjusting sound categories, and theories of language acquisition that incorporate adaptation as an important component.

We now turn to results from the oldest age group, which show a positive correlation. Although, in our view, this result needs to be taken with great caution, we must point out

that this precise pattern was predicted by a framework where infants use variability to learn to weigh linguistic and speaker-specific information in the acoustic signal differently (Rost & McMurray, 2009; Seidl et al., 2014). Since we only find a positive correlation at 12 months, it is tempting to propose a developmental change, enabling older infants to harness the beneficial properties of variable input which points to the reliable and important aspects of the signal (as proposed by e.g., Rost and McMurray 2009). Multiple studies have remarked a sharp increase in infants' word recognition abilities around their first birthday (Bergelson & Swingley, 2012), and the common factor might lie in an advancement in categorization skills and the ability to weigh cues differently (Younger & Cohen, 1986). Nonetheless, in the absence of a significant interaction of talker number with age group in our data, and given the multiple problems we raised with the interpretation of results in this age group, we do not pursue this possibility further.

Taking all three age groups together, it is clear that our data provide no support for the interpretation that talker variation *negatively* impacts vowel discrimination in real life. Nonetheless, the evidence from speech technology for variability degrading unsupervised category clustering (Bergmann et al., 2016), and for adult perception being best modeled when representations take into account talker variation (Richter et al., 2016) is so compelling, that we believe it is worth exploring it further. It is possible that more fine-grained predictors could possibly uncover a negative link between input variability and the development of infants' discrimination ability at a level that was not possible in the present study. Ideally, this would be done with daylong recordings that sample over several days or weeks (Koorathota, Morton, Amatuni, & Bergelson, 2016). From these, one could extract voice characteristics of all speakers in infants' input as well as measurement of the times different talkers spoke in succession, which could for example influence adaptation. Gathering such data was not feasible in the scope of the present study, but others may also find this prospect interesting, as these data could further enrich our understanding of the impact of infants' linguistic environment on their language development. Similarly, the strongest evidence in favor of interactions between talker variability and speech perception emerge from studies where generalization is studied (e.g., Houston and Jusczyk 2000). We purposefully chose to use a standard within-speaker discrimination task to tap native vowel categorization rather than generalization skills, and future research might address the interesting prediction that infants exposed to few talkers are less able to generalize to novel speakers than their more experienced peers.

If the conclusion that there is no strong relationship between speaker number and vowel discrimination holds, this would indicate a startling asymmetry between speech and non-speech category formation. Forming categories based on more or less variable input is an issue studied across domains of developmental research, with a large body of literature covering complex effects of variability on visual category formation (Eimas, Quinn, & Cowan, 1994; Oakes & Spalding, 1997). In the visual domain, the complexity and range of exemplars presented to infants has been shown to affect their ability to construct categories and generalize to new tokens (e.g., Younger and Cohen 1986). If acoustic development were to mirror visual categorization, we would have expected that younger infants are distracted by more talker variability, as multiple features vary. Continuing in their first year, older infants might have been able to deal with talker variation and the oldest age group would then, somewhat in line with our data, harness variability in a more robust fashion. However,

our data do not align with these predictions, as we see no support for a negative relationship between talker variability and vowel discrimination at any of the three ages we tested.

Based on our data, we suggest that acoustic category formation follows a different path from visual categorization of objects and events. This might be due to a number of factors that future research has to disentangle. First, even if infants try to apply the same domain-general category formation skills to their visual and auditory input, the learning outcomes may be different because information is more transient as visual input changes less rapidly than the speech signal. This may predict that category formation in sign languages would follow the pattern extrapolated from visual category work more closely, and does not align with the one we observe in spoken language acquisition. Second, infants might rely more on multi-modal information during speech processing compared to visual categorization, again because of the differences in stability across domains. A common spoken label influences infants' visual categorization (Fulkerson & Waxman, 2007), showing that infants can and do combine cues across modalities. We propose conversely that speech processing has to rely on some cross-modal cues, for example to track the environment and/or who is talking. Both of these accounts, the mono-modal and the multi-modal one, could explain differences in visual and acoustic categorization ability and development, but they interestingly differ in their predictions on the impact of cross-modal cues.

To sum up, the present study sought to assess a theoretically crucial question that has previously not been investigated experimentally: To what extent do infants' early experiences with multiple talkers shape their vowel discrimination skills? Our findings do not align with expectations drawn from a line of theories of infant perceptual category learning which assume that categories are acquired on the basis of input statistics. In the case of speech categories, those statistics become confusing as more talkers provide the input. Instead, our data support views where talker variation does not greatly impact infant vowel category acquisition, inviting further experimental and theoretical work. We suggest exploring specifically mechanisms of adaptation in the first year of life.

Acknowledgments

Contributions

Anonymized

References

Bergelson, E., & Swingley, D. (2012). At 6–9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences*, *109*(9), 3253–3258.

Bergmann, C., Cristia, A., & Dupoux, E. (2016). Discriminability of sound contrasts in the face of speaker variation quantified. In *Proceedings of the 38th annual meeting of the cognitive science society*.

Creel, S. C., Aslin, R. N., & Tanenhaus, M. K. (2008). Heeding the voice of experience: The role of talker variation in lexical access. *Cognition*, *106*(2), 633–664.

Cristia, A. (2011). Fine-grained variation in caregivers' /s/ predicts their infants' /s/ category. *The Journal of the Acoustical Society of America*, *129*(5), 3271–3280.

Dahan, D., Drucker, S. J., & Scarborough, R. A. (2008). Talker adaptation in speech perception: Adjusting the signal or the representations? *Cognition*, *108*(3), 710–718.

Dehaene-Lambertz, G., & Peña, M. (2001). Electrophysiological evidence for automatic phonetic processing in neonates. *Neuroreport*, *12*(14), 3155–3158.

Eimas, P. D., Quinn, P. C., & Cowan, P. (1994). Development of exclusivity in perceptually based categories of young infants. *Journal of Experimental Child Psychology*, *58*(3), 418–431.

Elsner, M., Goldwater, S., Feldman, N., & Wood, F. (2013). A joint learning model of word segmentation, lexical acquisition, and phonetic variability. *Proceedings of EMNLP*, 42–54.

Feldman, N. H., Griffiths, T. L., Goldwater, S., & Morgan, J. L. (2013). A role for the developing lexicon in phonetic category acquisition. *Psychological Review*, *120*(4), 751.

Fox, J., & Weisberg, S. (2011). *An R companion to applied regression* (Second ed.). Thousand Oaks CA: Sage. Retrieved from `http://socserv.socsci.mcmaster.ca/jfox/Books/Companion`

Fulkerson, A. L., & Waxman, S. R. (2007). Words (but not tones) facilitate object categorization: Evidence from 6-and 12-month-olds. *Cognition*, *105*(1), 218–228.

Gendrot, C., & Adda-Decker, M. (2005). Impact of duration on F1/F2 formant values of oral vowels: An automatic analysis of large broadcast news corpora in French and German. *Variations*, *2*(22.5), 2–4.

Gogate, L. J., Prince, C. G., & Matatyaho, D. J. (2009). Two-month-old infants' sensitivity to changes in arbitrary syllable–object pairings: The role of temporal synchrony. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(2), 508.

Hillenbrand, J., Getty, L. A., Clark, M. J., & Wheeler, K. (1995). Acoustic characteristics of american english vowels. *The Journal of the Acoustical Society of America*, *97*(5), 3099–3111.

Houston, D. M., & Jusczyk, P. W. (2000). The role of talker-specific information in word segmentation by infants. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(5), 1570–1582.

Huttenlocher, J., Haight, W., Bryk, A., Seltzer, M., & Lyons, T. (1991). Early vocabulary growth: Relation to language input and gender. *Developmental Psychology*, *27*(2), 236–248.

JASP Team. (2016). (JASP Version 0.7.5.5, Computer software)

Jusczyk, P. W., Pisoni, D. B., & Mullennix, J. (1992). Some consequences of stimulus variability on speech processing by 2-month-old infants. *Cognition*, *43*(3), 253–291.

Kleinschmidt, D. F., & Jaeger, T. F. (2015). Robust speech perception: Recognize the familiar, generalize to the similar, and adapt to the novel. *Psychological Review*, *122*(2), 148–203.

Koorathota, S., Morton, S., Amatuni, A., & Bergelson, E. (2016). *6 & 7-month-olds' noun input: Human and automated corpus analyses.* (Presented at the International Conference on Infant Studies)

Kuhl, P. K. (1979). Speech perception in early infancy: Perceptual constancy for spectrally

dissimilar vowel categories. *The Journal of the Acoustical Society of America*, *66*(6), 1668–1679.

Kuhl, P. K. (2004). Early language acquisition: Cracking the speech code. *Nature Reviews Neuroscience*, *5*(11), 831–843.

Liu, H.-M., Kuhl, P. K., & Tsao, F.-M. (2003). An association between mothers' speech clarity and infants' speech discrimination skills. *Developmental Science*, *6*(3), F1–F10.

Mandel, D. R., Jusczyk, P. W., & Pisoni, D. B. (1995). Infants' recognition of the sound pattern of their own names. *Psychological Science*, *6*(5), 314–317.

Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition*, *82*(3), B101–B111.

McQueen, J. M., Cutler, A., & Norris, D. (2006). Phonological abstraction in the mental lexicon. *Cognitive Science*, *30*(6), 1113–1126.

Meints, K., & Woodford, A. (2008). *Lincoln Infant Lab Package 1.0: A new programme package for IPL, Preferential Listening, Habituation and Eyetracking.* Retrieved from `http://www.lincoln.ac.uk/psychology/babylab.htm` (WWW document: Computer software & manual)

Melvin, S. A., Brito, N. H., Mack, L. J., Engelhardt, L. E., Fifer, W. P., Elliott, A. J., & Noble, K. G. (2016). Home environment, but not socioeconomic status, is linked to differences in early phonetic perception ability. *Infancy*.

Morey, R. D., & Rouder, J. N. (2015). BayesFactor: Computation of Bayes Factors for Common Designs [Computer software manual]. Retrieved from `https://CRAN.R-project.org/package=BayesFactor` (R package version 0.9.12-2)

Mulak, K. E., Best, C. T., Tyler, M. D., Kitamura, C., & Irwin, J. R. (2013). Development of phonological constancy: 19-month-olds, but not 15-month-olds, identify words in a non-native regional accent. *Child Development*, *84*(6), 2064–2078.

Oakes, L. M., & Spalding, T. L. (1997). The role of exemplar distribution in infants' differentiation of categories. *Infant Behavior and Development*, *20*(4), 457–475.

Patterson, M. L., & Werker, J. F. (2003). Two-month-old infants match phonetic information in lips and voice. *Developmental Science*, *6*(2), 191–196.

Polka, L., Masapollo, M., & Ménard, L. (2014). Who's talking now? Infants' perception of vowels with infant vocal properties. *Psychological Science*, *25*(7), 1448–1456.

Pons, F., Albareda-Castellot, B., & Sebastián-Gallés, N. (2012). The interplay between input and initial biases: Asymmetries in vowel perception during the first year of life. *Child Development*, *83*(3), 965–976.

Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. L. (1993). Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception*, *22*(4), 463–475.

R Core Team. (2016). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from `https://www.R-project.org/`

Richter, C., Feldman, N. H., Salgado, H., & Jansen, A. (2016). A framework for evaluating speech representations. In J. Trueswell, A. Papafragou, D. Grodner, & D. Mirman (Eds.), *Proceedings of the 38th annual meeting of the cognitive science society.* Cognitive Science Society.

Rost, G. C., & McMurray, B. (2009). Speaker variability augments phonological processing

in early word learning. *Developmental Science*, *12*(2), 339–349.

Schmale, R., Seidl, A., & Cristia, A. (2015). Mechanisms underlying accent accommodation in early word learning: Evidence for general expansion. *Developmental science*, *18*(4), 664–670.

Seidl, A., Onishi, K. H., & Cristia, A. (2014). Talker variation aids young infants' phonotactic learning. *Language Learning and Development*, *10*(4), 297–307.

Swingley, D. (2009). Contributions of infant word learning to language development. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *364*(1536), 3617–3632.

Teinonen, T., Aslin, R. N., Alku, P., & Csibra, G. (2008). Visual speech contributes to phonetic learning in 6-month-old infants. *Cognition*, *108*(3), 850–855.

Tincoff, R., & Jusczyk, P. W. (1999). Some beginnings of word comprehension in 6-month-olds. *Psychological Science*, *10*(2), 172–175.

Tsuji, S., & Cristia, A. (2014). Perceptual attunement in vowels: A meta-analysis. *Developmental Psychobiology*, *56*(2), 179–191.

van Heugten, M., & Johnson, E. K. (2012). Infants exposed to fluent natural speech succeed at cross-gender word recognition. *Journal of Speech, Language, and Hearing Research*, *55*(2), 554–560.

Wagenmakers, E.-J., Morey, R. D., & Lee, M. D. (2016). Bayesian benefits for the pragmatic researcher. *Current Directions in Psychological Science*, *25*(3), 169–176.

Wetzels, R., & Wagenmakers, E.-J. (2012). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic Bulletin & Review*, *19*(6), 1057–1064.

Yeung, H. H., & Nazzi, T. (2014). Object labeling influences infant phonetic learning and generalization. *Cognition*, *132*(2), 151–163.

Yeung, H. H., & Werker, J. F. (2009). Learning words' sounds before learning how words sound: 9-month-olds use distinct objects as cues to categorize speech information. *Cognition*, *113*(2), 234–243.

Younger, B. A., & Cohen, L. B. (1986). Developmental change in infants' perception of correlations among attributes. *Child Development*, *57*(3), 803–815.