

Listeners are maximally flexible in updating phonetic beliefs over time

David Saltzman¹ · Emily Myers¹

© Psychonomic Society, Inc. 2017

Abstract Perceptual learning serves as a mechanism for listeners to adapt to novel phonetic information. Distributional tracking theories posit that this adaptation occurs as a result of listeners accumulating talker-specific distributional information about the phonetic category in question (Kleinschmidt & Jaeger, 2015, *Psychological Review*, 122). What is not known is how listeners build these talker-specific distributions; that is, if they aggregate all information received over a certain time period, or if they rely more heavily upon the most recent information received and down-weight older, consolidated information. In the present experiment, listeners were exposed to four interleaved blocks of a lexical decision task and a phonetic categorization task in which the lexical decision blocks were designed to bias perception in opposite directions along a “s”–“sh” continuum. Listeners returned several days later and completed the identical task again. Evidence was consistent with listeners using a relatively short temporal window of integration at the individual session level. Namely, in each individual session, listeners’ perception of a “s”–“sh” contrast was biased by the information in the immediately preceding lexical decision block, and there was no evidence that listeners summed their experience with the talker over the entire session. Similarly, the magnitude of the bias effect did not change between sessions, consistent with the idea that talker-specific information remains flexible, even after consolidation. In general, results suggest that listeners are maximally flexible when considering how to categorize speech from a novel talker.

Keywords Speech perception · Spoken word recognition

Perceptual learning is an inherent component of speech perception. Talkers vary significantly in their phonetic properties (e.g., Hillenbrand, Getty, Clark, & Wheeler, 1995), and consequently listeners must adjust their mapping between acoustics and phonetic categories for each new talker that they encounter. Luckily for the listener, this variability tends to have a statistical structure that is characteristic of the talker. For instance, a given talker may have a consistently longer mean voice onset time (VOT) for voiceless stops (Allen, Miller, & DeSteno, 2003), or consistently low F2 value for vowels (Hillenbrand et al., 1995). Further, individual talkers also differ in their variability; that is, one talker may have wide variability in their productions, whereas another may produce a narrower range of acoustic values (Newman, Clouse, & Burnham, 2001).

An array of findings supports the view that listeners are sensitive to the phonetic characteristics of a given talker, and that they adjust their perceptual criteria to use this information to reach a stable phonetic percept (Clayards, Tanenhaus, Aslin, & Jacobs, 2008; Kleinschmidt & Jaeger, 2015; Kraljic & Samuel, 2007; Theodore & Miller, 2010). Many accounts of perceptual learning share the notion that talker adaptation involves tracking the statistics of a talker’s speech over time, discovering the distributional acoustic patterns associated with each novel talker, and using this information to create probabilistic maps between acoustics and linguistic representations (Maye, Weiss, & Aslin, 2008; McMurray, Aslin, & Toscano, 2009; recently formalized using a Bayesian framework in Kleinschmidt & Jaeger, 2015). Under this view, distributional information is combined with contextual information (e.g., “who is the talker,” “what word is likely in this context”) to generate a talker-specific, contextually bound probability that

✉ Emily Myers
emily.myers@uconn.edu

¹ Department of Speech, Language, and Hearing Sciences, University of Connecticut, 850 Bolton Road, Unit 1085, Storrs, CT 06269, USA

a given acoustic token will match a likely phonetic category. This class of theories predicts that changing the statistical distribution of tokens in the input will ultimately result in perceptual adaptation.

Perceptual learning paradigms (e.g., Bertelson, Vroomen, & De Gelder, 2003; Norris, McQueen, & Cutler, 2003) demonstrate many situations in which listeners quickly adapt to phonetic characteristics of a novel talker. Listeners might hear an ambiguous speech token which is resolved either by lexical context (e.g., Norris et al., 2003) or audiovisual information (e.g., Bertelson et al., 2003), accompanied by a clear version of the contrasting phonetic category. The speech stream thus contains both top-down contextual information (i.e., “in the lexical context, ‘epi_ode’, ‘s’ is the only probable interpretation of the ambiguous sound”) as well as bottom-up distributional information (i.e., listeners are exposed to a bimodal distribution of tokens—one ambiguous, one clear—that is shifted for each of the exposure conditions). Using the distributional learning framework, the effect found in perceptual learning studies can be explained as the listener pairing top-down information about phoneme identity with distributional information about the statistics of the novel talker’s input, which in turn allows for reshaping of their phonetic categories (Pajak, Fine, Kleinschmidt, & Jaeger, 2016).

The argument that listeners maintain distributional information for each unique talker is described in Kleinschmidt and Jaeger (2015), which follows earlier research showing that listeners do indeed maintain talker-specific information (Goldinger, 1996; Nygaard & Pisoni, 1998). Talker-specific distributional representations help to explain how perceptual learning effects persist over time (Kraljic & Samuel, 2005). For instance, in a study from Eisner and McQueen (2006), participants maintained talker-specific information over a 12-hour delay, unaffected by exposure to different speakers in the intervening period between exposure and testing, suggesting that any new mapping between acoustics and phonology was specific to the test talker. In Kraljic and Samuel (2005, Exp. 3), participants engaged in a lexically guided perceptual learning (LGPL) task in which they were first exposed to phonetically biasing information, namely ambiguous tokens embedded in an unambiguous lexical context, then exposed to unaltered exemplars of previously ambiguous sounds from the same talker, and then tested. This led to an extinction of the perceptual learning effect, which is congruent with the notion from distributional tracking theories that listeners would integrate the good exemplars into the talker-specific phonetic distribution, thus disrupting the shifted category representations that they had formed for this new talker.

One issue that has received less attention is the processes by which listeners integrate new talker-specific information (here termed “recent statistics”) with existing information that listeners have accumulated about the total talker-specific distribution of acoustic cues (here termed “global statistics”).

Kleinschmidt and Jaeger (2015) state that “in situations like a recalibration experiment where listeners encounter odd-sounding, often synthesized speech in a laboratory setting, they may have little confidence, a priori, that any of their previous experiences are directly informative” (p. 13), and thus predict listeners will be maximally flexible during these experiments as the value of previous experiences with the category in question are not believed to be informative. The results of Kraljic and Samuel (2005) appear to confirm that recently encountered statistics are given a stronger weighting; that is, if listeners heavily weight new tokens, the most recent input should more strongly shape the phonetic category. Furthermore, in a series of experiments by Van Linden and Vroomen (2007, Exps. 1–4), listeners were exposed to both lip-reading and lexically biasing information for a “t”–“p” contrast in a blocked design, and the effect of the biasing information was sampled sporadically in each block. Their results demonstrate that (1) listeners can shift their category boundaries flexibly within an experiment and also (2) use the most recent statistics when building a distribution. Contrastively, Kleinschmidt and Jaeger posit that talker-specific distributions cannot be created or maintained if a listener simply tracks the recent statistics from a talker (p. 26), and go on to demonstrate a model for how beliefs about a talker are updated over experiences (see their Fig. 17).

In the current study, we ask whether listeners are continuously flexible in their adjustment to new and conflicting phonetic information about a talker, and how this affects their ability to create a talker-specific cue distribution. One possibility (see “Global Statistics”, Fig. 1) is that participants aggregate all the input from a given talker into one unified distribution, assigning equal weight to each token in memory. In this case, a listener who hears ambiguous tokens in an “s”-biasing context, for instance, and is tested on this contrast should see the previously attested shift in phonetic category boundary. Subsequent exposure to an “sh”-biasing block, however, will simply add new tokens to the emerging “s” and “sh” distributions for the talker (see “Global Statistics”, Fig. 1), leaving the category boundary somewhere in the middle of the distribution. Under this view, a participant who had heard the “sh”-bias first would show a shift to incorporate ambiguous tokens in the “sh” category, but her categorization function after being exposed to “s” tokens next would be equivalent to the participant who heard the bias blocks in the opposite order, since both participants would have heard the full complement of stimuli by the end of the experiment. Essentially, this leads to a prediction that the order of presentation of these blocks will matter, with categorization functions equalizing after listeners have heard both “s”-biasing and “sh”-biasing blocks. The biasing effect should be even more diminished if the participant were to return and complete the same task again, as listeners should be updating their beliefs about the talker’s distribution with an aggregate of all of the information they

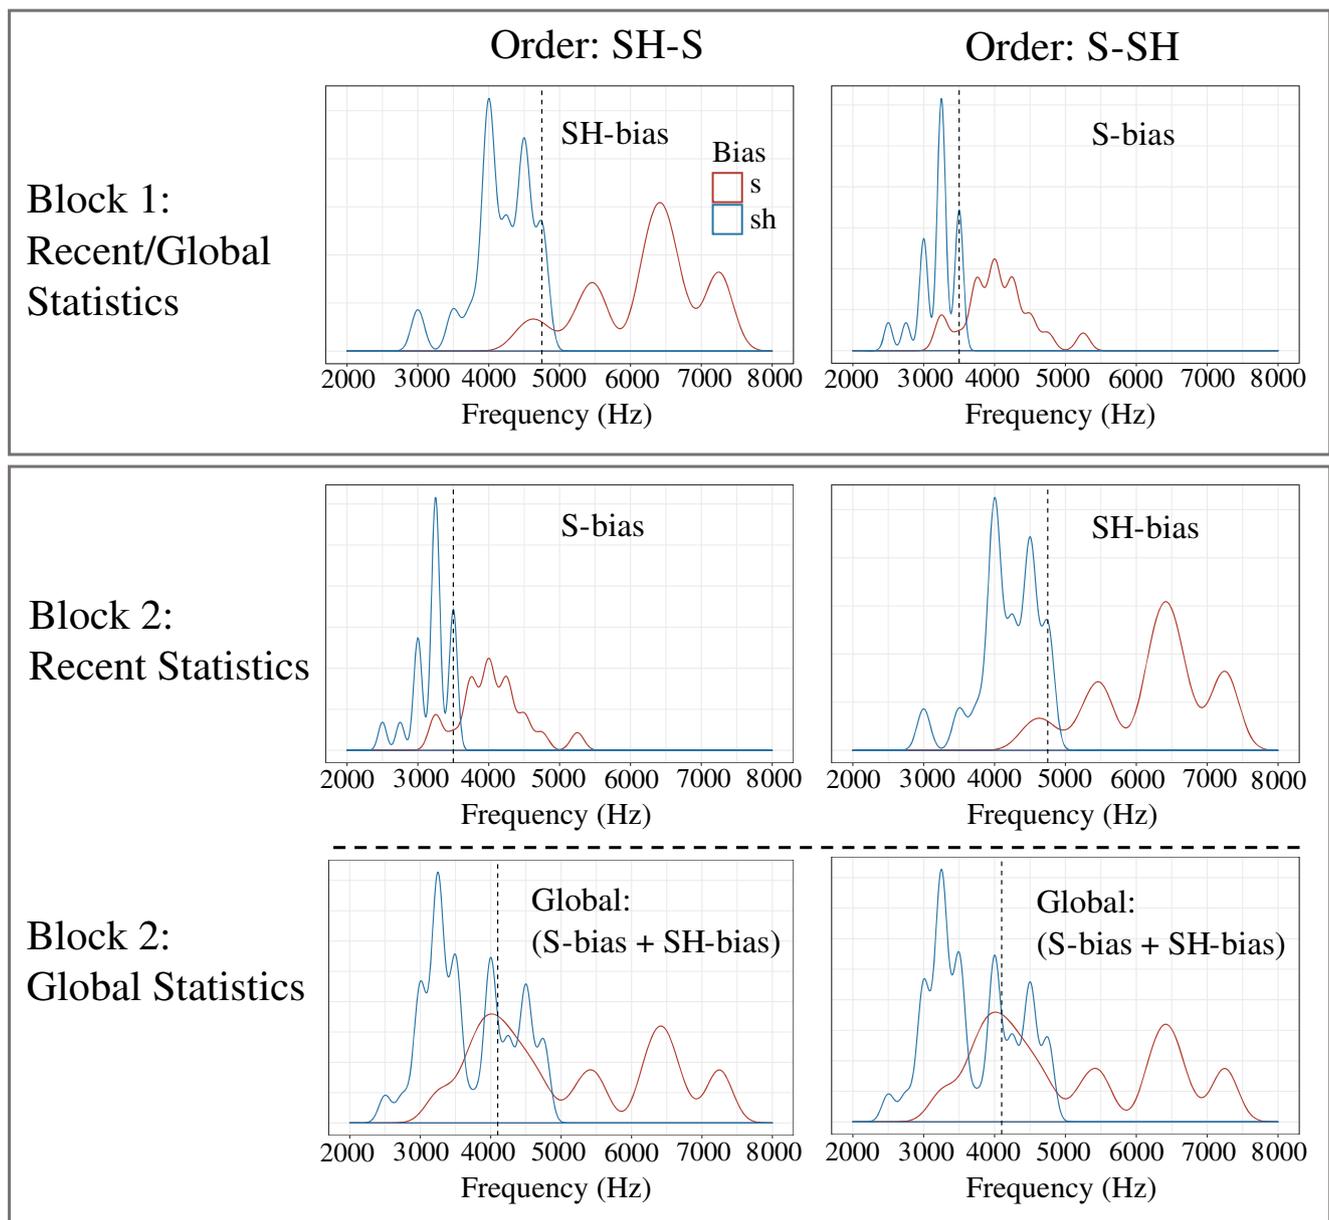


Fig. 1 Schematic showing the probability density function over the centroid frequencies of the “s” and “sh” tokens that listeners hear on each block of the LD task (see Drouin, Theodore, & Myers, 2016). Blue shows the “sh” tokens, red shows the “s” tokens. In SH-bias blocks, listeners hear naturally produced versions of the “s” tokens (red) and altered, ambiguous versions of the “sh” tokens, while the reverse is true of S-bias blocks. In each panel, a vertical dotted line indicates a hypothetical ideal boundary that minimizes miscategorizations of the

exposure set. After Block 1, both recent statistics and global statistics hypotheses predict the same boundary. However, after Block 2, the recent statistics hypothesis predicts that listeners will resolve on a boundary dictated by the immediately previous LD block (Block 2: recent statistics), while the global statistics view predicts that listeners will generate distributions over the entire set of LD stimuli, and thus both groups of participants will show the same boundary value at Block 2. (Color figure online)

received in their first exposure to the talker. In addition, the talker should now be more familiar, which should allow the listener to safely use this prior experience with the talker to inform their future experience with input from said talker. In essence, the more speech a listener hears from a talker, the “heavier” the distributional information for that talker, and the harder it should be to shift.

An alternative is that listeners are maximally flexible, easily disregard old information about a talker, looking only to the most recently encountered tokens when considering how to process incoming information (see “Recent Statistics”, Fig. 1). This would predict that listeners will shift and reshift their phonetic criteria on the basis of recent information, and that the shift for the second-encountered bias will be just as large

as that of the first set of biasing information a listener hears. Upon a second exposure to the same task, we should see listeners continue to shift and reshift their category boundaries as a result of the biasing information. Following this hypothesis, it is possible that listeners will create a very flexible talker-specific distribution (or perhaps do not create one at all, which is discussed later) and simply move around in that distributional space.

To test these alternatives, in the current study, we manipulated lexical bias within participants, otherwise closely following methods of Kraljic and Samuel (2005). Listeners were exposed to four interleaved blocks of a lexical decision task and a phonetic categorization task (see Fig. 2) in which the lexical blocks were designed to bias perception in opposite directions. Listeners also returned several days later for a second session, in which they completed the identical task from their first session.

If listeners behave per the global statistics hypothesis, performance on the final phonetic categorization task in Session 1 will be the same regardless of whether the lexical-decision block immediately preceding it was “s”-biasing or “sh”-biasing, and therefore we should see a main effect of order (i.e., a shift in the predicted direction after the first biasing block, then an equilibration of the effect after being exposed to the opposite-direction bias). Also per this hypothesis, the main effect of bias should be significantly reduced (or extinguished) during session 2, leading to a bias by session interaction. However, if listeners behave per the recent statistics hypothesis, we should see a large boundary shift in their categorization functions following each biasing block, regardless of the order of the blocks. This effect should also reproduce during the second session.

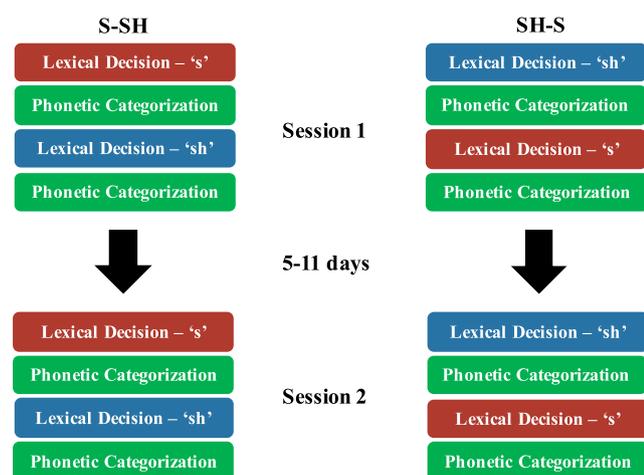


Fig. 2 Experimental schedule. Participants were assigned to either the S-SH group or SH-S group (see text for details). In each session, lexical decision (red: s-biasing block, blue: sh-biasing block) blocks alternated with phonetic categorization (green) blocks. Participants returned after 5–11 days to repeat the identical experimental procedure. (Color figure online)

Method

Participants

Seventy-four undergraduates (ages 18–23 years, $M = 18.83$) were recruited from the University of Connecticut. All participants indicated that they were monolingual English speakers with normal hearing. A written informed consent was obtained from every participant in accordance with the guidelines of the University of Connecticut Institutional Review Board. Participants received course credit for their participation.

Stimuli

Stimuli for the lexical decision (LD) task were taken from Myers and Mesite (2014).¹ These items consisted of 200 total words, 100 filler nonwords, 60 filler real words, 20 critical “s” words, and 20 critical “sh” words. The critical words were real words containing either “s” or “sh” in a word-medial position. Acoustically modified versions of these words were created by replacing the “s” or “sh” with an ambiguous, 50%–50% blend of the two sounds. Further details about stimuli can be found in Myers and Mesite (2014). In the “s”-biasing condition, listeners heard words containing the ambiguous blend (“?”) in “s”-containing words and unaltered versions of the “sh”-words (e.g., “epi?ode,” “flourishing”). In the “sh”-biasing condition, the ambiguous blend appeared in “sh”-words and listeners also heard unaltered versions of the “s”-words (e.g., “flouri?ing,” “episode”).

Items for the phonetic categorization (PC) task consisted of a seven-step continuum from *sign* to *shine*, which were created in PRAAT (Boersma and Weenink, 2017) by blending (through waveform averaging) fricatives derived from the words *sign* and *shine* at different proportions from 20% “s”–80% “sh” to 80% “s”–20% “sh.” The blended fricatives were then inserted into the *sign* frame. The *sign*–*shine* continuum was pilot tested to ensure consistent perception of the endpoints of the continuum. The same talker was used for the LD and PC stimuli.

Procedure

The experiment took place over two sessions (see Fig. 2). In Session 1, participants engaged in alternating blocks of an LD task and PC task. LD blocks contained lexical information that biased listeners to perceive an ambiguous phoneme as either “s” or “sh,” and PC blocks tested the effects of having heard this biasing information. Participants were randomly assigned to either the S-SH group (in which the first LD block contained “s”-critical words and the second the “sh”-critical words) or the SH-S group (the reverse order). The PC task was identical across groups. In Session 2, participants returned

¹ See their Methods and Materials subsections Stimulus Selection and Stimulus Construction.

between 5 and 11 days later ($M = 7.10$ days, $SD = 1.19$ days) and completed the identical task as in Session 1, with the identical ordering of biasing LD conditions.

In the LD task, participants were asked to indicate whether the stimulus was a word or a nonword by pressing a corresponding key on the keyboard as quickly and accurately as possible. In the PC task, participants were asked to indicate whether they perceived the stimulus as *sign* or *shine*, which they did by pressing a corresponding key on the keyboard as quickly as possible. Each categorization task consisted of eight repetitions of each of the seven tokens from the *sign* to *shine* continuum presented in random order, for a total of 56 trials per PC block. Response options were counterbalanced in both tasks

Results

Session 1

Data from 19 participants were excluded for experimenter error ($n = 4$) or displaying poor categorization at the endpoints of the continuum (defined as less than 80% accuracy at each endpoint) averaged across the two categorization tasks from both sessions ($n = 15$). After exclusion, 55 participants remained for analysis.

A generalized-linear mixed-effects model with a logit-link was performed in the R statistical computing language (R Development Core Team, 2014) using the `glmer` command from the `lme4` package (Bates, Mächler, Bolker, & Walker, 2014). A backward-stepping selection heuristic was used and a model with continuum step (centered), biasing condition, order of presentation, and their interactions as fixed effects, and by-subject random slopes and intercepts for continuum step and biasing condition were justified by the data. All predictors except continuum step were contrast coded.

As expected, a significant main effect of continuum step was revealed, such that participants were more likely to indicate that they heard *sign* as the proportion of “s” in the

fricative blend increased ($b = 4.40$, $SE = 0.18$, $z = 23.43$, $p < .001$; Fig. 3). In addition, a significant main effect of biasing condition was found, reflecting increased *sign* responses immediately following a “s” LD block and decreased *sign* responses immediately following a “sh” LD block ($b = -0.93$, $SE = 0.10$, $z = -9.53$, $p < .001$), thus replicating the classic LGPL effect (Kraljic & Samuel, 2005; Norris et al., 2003). Notably, no effect of order of presentation was found ($p = .86$). No interactions between any fixed effects were significant.

Session 2

A generalized-linear mixed-effects model with a logit-link was performed, with the same fixed effects structure as Session 1, and by-subject random slopes and intercepts for continuum step, biasing condition, and their interaction. As found in Session 1, there was a significant main effect of continuum step ($b = 4.95$, $SE = 0.25$, $z = 19.41$, $p < .001$; Fig. 3) and of biasing condition ($b = -1.01$, $SE = 0.10$, $z = -9.78$, $p < .001$). Again, there was no effect of order of presentation ($p = .14$). All interactions were nonsignificant.

Stability over time

Data from Session 1 and Session 2 were combined to examine how listeners update their beliefs about a talker over time. A new generalized-linear mixed-effects model with a logit-link was performed, with continuum step, biasing condition, session number, order of presentation, and their interactions as fixed effects and by-subject random slopes for continuum step (centered), biasing condition, session number, and their interaction. All predictors except continuum step were contrast coded. As expected from Sessions 1 and 2, there was a significant main effect of continuum step ($b = 4.70$, $SE = 0.18$, $z = 26.33$, $p < .001$; Fig. 3), as well as biasing condition ($b = -0.96$, $SE = 0.07$, $z = -13.01$, $p < .001$). Order was again nonsignificant ($p = .41$).

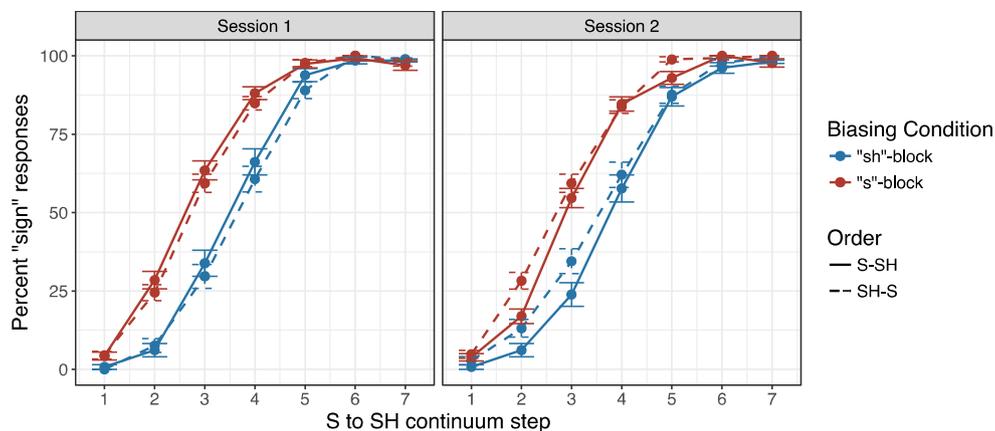


Fig. 3 Data from the phonetic categorization task. Order (S-SH, SH-S) was a between-subjects factor. Biasing condition indicates the type (“sh”-biasing or “s”-biasing) of the immediately-preceding LD block. Error bars reflect standard error of the mean. (Color figure online)

Of interest is whether participants were equally likely to shift the category boundary in response to lexical information in the second session compared to the first. There was no significant main effect of session ($p = .74$). A Session \times Order interaction approached significance ($p = .06$), which can be seen in the trend toward more *sign* responses overall in the SH-S order during Session 2. Critically, all other interactions were nonsignificant, and the lack of a Session \times Bias effect suggests that the magnitude of the bias effect was equivalent across sessions.

Discussion

Distributional learning accounts of speech typically do not specify precisely how individual episodes are aggregated over time in order to inform perceptual learning (Kleinschmidt & Jaeger, 2015; Maye et al., 2008; McMurray et al., 2009). If listeners had simply summed all of the information they had gleaned about the talker during the session, then the second PC block would show an equivalent boundary value for the S-SH and SH-S groups. Instead, in the current study, listeners used biasing lexical information to shift their category boundary first in one direction (e.g., toward the *sign* end of the continuum) and then, when confronted with the opposite bias (e.g., toward the *shine* end of the continuum), back in the other direction. Crucially, there was no evidence that the order of presentation of these blocks (S-SH vs. SH-S) affected categorization responses within a session. This result suggests that listeners use a relatively short temporal window of integration when they are considering how to interpret the speech of a talker, strongly weighting recent biasing information instead of building a session-long distributional scheme for the talker. However, it is possible that with increased exposure to phonetically biasing material within each task, selective adaptation would begin to occur and the size of the biasing shift would diminish significantly (Vroomen, van Linden, de Gelder, & Bertelson, 2007).

Qualitative changes to learned phonetic information may emerge over time, especially after sleep (Earle & Myers, 2015). In particular, sleep-mediated consolidation appears to stabilize learned phonetic information and protects this information from interference (Earle & Myers, 2015; Fenn, Nusbaum, & Margoliash, 2003). If these same principles operate in this paradigm, it follows that distributional information that listeners heard during Session 1 would become stabilized overnight, yielding a lessened ability to respond to distributional learning in Session 2. There was no evidence of this—in fact, listeners displayed an equivalent bias shift in the second session compared to the first session. This finding follows that of Eisner and McQueen (2006), where no change in learning was found between participants who were tested on the same day, but with a 12-hour delay, or on the next day

after sleeping. One caveat should be noted: Because the phonetic information that was provided to listeners in the first session was essentially inconsistent or erratic, it is possible that listeners adapt a conservative strategy in interpreting the speech of the talker, and they do not settle into any particular phonetic boundary for that talker. This could come about from bottom-up mechanisms (the distribution is too broad and shallow for the system to settle) or from top-down mechanisms (the talker is viewed somehow as unreliable; see, for instance, Kraljic, Samuel, & Brennan, 2008).

Kleinschmidt and Jaeger (2015) also discuss the possibility that listeners may not always form talker-specific beliefs, especially in situations where there is no expectation they would be useful again in the future (such as in a laboratory experiment). While this hypothesis could explain the present findings, participants were instructed during the consent process that the two sessions would be identical. Given that these instructions came before an explanation of the experiment's procedure, and were not wholly explicit, it is possible the participants assumed there was no reason to form talker-specific beliefs for a transient situation. Nevertheless, multiple studies of perceptual learning have found that talker-specific phonetic distributions persist over time, implying that participants may form them regardless of the situation (Eisner & McQueen, 2006; Kraljic & Samuel, 2005). Future research should explore the effect of top-down instructions on listener's willingness to create talker-specific distributions.

An auxiliary question that this study allows us to answer is whether individuals are consistent in the size of the boundary shift that they display across sessions. A secondary analysis showed that there was no significant relationship between the size of the biasing effect across sessions.² Individual differences in language learning are becoming of increasing interest to explain the gulf in outcomes between learners. For instance, incidental language learning paradigms have found that factors such as declarative learning abilities, procedural memory, some learning styles, personality factors, and sequence learning can have an effect on learning performance (Granena, 2013; Grey, Williams, & Rebuschat, 2015; Hamrick, 2015). More relevant to distributional tracking theories is the idea that statistical learning may be a skill unto itself, with accompanying individual differences. Siegelman and Frost (2015) found that their participants' performance on a series of statistical learning tasks was stable at an individual level across time. Surprisingly, in the current study, participants were not consistent in the size of the perceptual learning effect across sessions. Future research will be directed at discovering other mediating factors that explain this lack of correlation.

² Mean difference in percentage of *sign* responses following the “s” and “sh” blocks was calculated for each participant and session as an estimate of the size of the biasing effect. There was no significant correlation ($r = .12, p = .37$) in the bias effect size between Session 1 and Session 2.

Conclusion

The ability to rapidly adjust to novel information about a talker is perhaps not surprising—a statistical account in which listeners weighted all tokens from a talker equally would generate the prediction that it would be extremely difficult to adapt to the speech of very well-known talkers if a new perturbation or disruption was introduced. It would mean, for instance, that listeners who had only seen Meryl Streep playing American roles would struggle when confronted with her Polish-accented English in *Sophie's Choice*, or that one might fail to understand the distorted speech of one's parent after dental surgery. It is an empirical question whether this is the case for very familiar talkers. However, for the novel talkers in the current experiment, listeners appear to be maximally flexible to the most recent biasing information they are provided with, and do not create particularly rigid talker-specific beliefs.

The results of the current study urge a more detailed specification of memory in current models of distributional learning. Our findings particularly point to the flexibility that listeners show in accumulating phonetic information about a talker. Distributional accounts must allow (a) new evidence to have an outsize effect on learning or (b) for local distributions to be formed on the basis of contextual information, for instance, allowing for the accumulation of information in units as long as the lexical decision block.

Acknowledgements This work was supported by NIH NIDCD grant R01 DC013064 to EBM. The views expressed here reflect those of the authors and not the NIH or the NIDCD. We would like to thank Rachel Theodore for very helpful comments on an earlier version of this manuscript, and Julia Drouin for her contributions to the acoustic analysis.

References

- Allen, J. S., Miller, J. L., & DeSteno, D. (2003). Individual talker differences in voice-onset-time. *The Journal of the Acoustical Society of America*, *113*, 544–552.
- Bates D, Mächler M, Bolker B, Walker S (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*, 1–48.
- Bertelson, P., Vroomen, J., & De Gelder, B. (2003). Visual recalibration of auditory speech identification: A McGurk aftereffect. *Psychological Science*, *14*, 592–597.
- Boersma, P., & Weenink, D. (2017). Praat: Doing phonetics by computer (Version 6.0.26) [Computer program]. Retrieved from <http://www.praat.org/>
- Clayards, M., Tanenhaus, M. K., Aslin, R. N., & Jacobs, R. A. (2008). Perception of speech reflects optimal use of probabilistic speech cues. *Cognition*, *108*, 804–809.
- Drouin, J. R., Theodore, R. M., & Myers, E. B. (2016). Lexically guided perceptual tuning of internal phonetic category structure. *The Journal of the Acoustical Society of America*, *140*(4), EL307–EL313.
- Earle, F. S., & Myers, E. B. (2015). Sleep and native language interference affect non-native speech sound learning. *Journal of Experimental Psychology: Human Perception and Performance*, *41*, 1680–1695.
- Eisner, F., & McQueen, J. M. (2006). Perceptual learning in speech: Stability over time. *The Journal of the Acoustical Society of America*, *119*, 1950–1953.
- Fenn, K. M., Nusbaum, H. C., & Margoliash, D. (2003). Consolidation during sleep of perceptual learning of spoken language. *Nature*, *425*, 614–616.
- Goldinger, S. D. (1996). Words and voices: episodic traces in spoken word identification and recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 1166–1183.
- Granena, G. (2013). Individual differences in sequence learning ability and second language acquisition in early childhood and adulthood. *Language Learning*, *63*, 665–703.
- Grey, S., Williams, J. N., & Rebuschat, P. (2015). Individual differences in incidental language learning: Phonological working memory, learning styles, and personality. *Learning and Individual Differences*, *38*, 44–53.
- Hamrick, P. (2015). Declarative and procedural memory abilities as individual differences in incidental language learning. *Learning and Individual Differences*, *44*, 9–15.
- 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*, 3099–3111.
- 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*, 148.
- Kraljic, T., & Samuel, A. G. (2005). Perceptual learning for speech: Is there a return to normal? *Cognitive Psychology*, *51*, 141–178.
- Kraljic, T., & Samuel, A. G. (2007). Perceptual adjustments to multiple speakers. *Journal of Memory and Language*, *56*, 1–15.
- Kraljic, T., Samuel, A. G., & Brennan, S. E. (2008). First impressions and last resorts: How listeners adjust to speaker variability. *Psychological Science*, *19*, 332–338.
- Maye, J., Weiss, D. J., & Aslin, R. N. (2008). Statistical phonetic learning in infants: Facilitation and feature generalization. *Developmental Science*, *11*, 122–134.
- McMurray, B., Aslin, R. N., & Toscano, J. C. (2009). Statistical learning of phonetic categories: Insights from a computational approach. *Developmental Science*, *12*, 369–378.
- Myers, E. B., & Mesite, L. M. (2014). Neural systems underlying perceptual adjustment to non-standard speech tokens. *Journal of Memory and Language*, *76*, 80–93.
- Newman, R. S., Clouse, S. A., & Burnham, J. L. (2001). The perceptual consequences of within-talker variability in fricative production. *The Journal of the Acoustical Society of America*, *109*, 1181–1196.
- Norris, D., McQueen, J. M., & Cutler, A. (2003). Perceptual learning in speech. *Cognitive Psychology*, *47*, 204–238.
- Nygaard, L. C., & Pisoni, D. B. (1998). Talker-specific learning in speech perception. *Attention, Perception, & Psychophysics*, *60*, 355–376.
- Pajak, B., Fine, A. B., Kleinschmidt, D. F., & Jaeger, T. F. (2016). Learning additional languages as hierarchical probabilistic inference: insights from L1 processing. *Language Learning*, *10*, 900–944.
- Siegelman, N., & Frost, R. (2015). Statistical learning as an individual ability: Theoretical perspectives and empirical evidence. *Journal of Memory and Language*, *81*, 105–120.
- Theodore, R. M., & Miller, J. L. (2010). Characteristics of listener sensitivity to talker-specific phonetic detail. *The Journal of the Acoustical Society of America*, *128*, 2090–2099.
- Vroomen, J., van Linden, S., De Gelder, B., & Bertelson, P. (2007). Visual recalibration and selective adaptation in auditory–visual speech perception: Contrasting build-up courses. *Neuropsychologia*, *45*, 572–577.