The Role of Multisensory Temporal Covariation in Audiovisual Speech Recognition in Noise

Jonathan H. Venezia^{1,2}, Robert Sandlin³, Leon Wojno³, Anthony D. Tran³, Gregory Hickok^{3,4},
and Virginia M. Richards³

¹VA Loma Linda Healthcare System

²Department of Otolaryngology, School of Medicine, Loma Linda University

³Department of Cognitive Sciences, University of California, Irvine

⁴Center for Language Science, University of California, Irvine

Corresponding Author:

Jonathan Venezia
VA Loma Linda Healthcare System
11201 Benton Street
Loma Linda, CA 92357

Phone: (909) 825-7084 x7292

Fax: (909) 777-3244

Email: jonathan.venezia@va.gov

Author Note

Research reported in this publication was supported by the National Institute of Deafness and Other Communication Disorders under Award Numbers R21 DC013406 (MPIs: VMR and Y. Shen) and R01 DC03681 (PI: GH). During this investigation, JHV was supported by the National Institute on Deafness and Other Communication Disorders under Award Number T32 DC010775 from the University of California, Irvine, CA, USA. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. This material is the result of work supported with resources and the use of facilities at the VA Loma Linda Healthcare System, Loma Linda, CA. The contents do not represent the views of the U.S. Department of Veterans Affairs or the United States Government.

Abstract

Static and dynamic visual speech cues contribute to audiovisual (AV) speech recognition in noise. Static cues (e.g., “lipreading”) provide complementary information that enables perceivers to ascertain ambiguous acoustic-phonetic content. The role of dynamic cues is less clear, but one suggestion is that temporal covariation between facial motion trajectories and the speech envelope enables perceivers to recover a more robust representation of the time-varying acoustic signal. Modeling studies show this is computationally feasible, though it has not been confirmed experimentally. We conducted two experiments to determine whether AV speech recognition depends on the magnitude of cross-sensory temporal coherence (AVC). In Experiment 1, sentence-keyword recognition in steady-state noise (SSN) was assessed across a range of signal-to-noise ratios (SNRs) for auditory and AV speech. The auditory signal was unprocessed or filtered to remove 3-7 Hz temporal modulations. Filtering severely reduced AVC (magnitude-squared coherence of lip trajectories with cochlear-narrowband speech envelopes), but did not reduce the magnitude of the AV advantage ($AV > A$; ~ 4 dB). This did not depend on the presence of static cues, manipulated via facial blurring. Experiment 2 assessed AV speech recognition in SSN at a fixed SNR (-10.5 dB) for subsets of Exp. 1 stimuli with naturally high or low AVC. A small effect ($\sim 5\%$ correct; high-AVC $>$ low-AVC) was observed. A computational model of AV speech intelligibility based on AVC yielded good overall predictions of performance, but over-predicted the differential effects of AVC. These results suggest the role and/or computational characterization of AVC must be re-conceptualized.

The Role of Multisensory Temporal Covariation in Audiovisual Speech Recognition in Noise

The addition of visual speech cues to an auditory speech signal confers a significant speech recognition advantage in the presence of acoustic noise, equal to a gain of roughly 4-6 dB for sentence materials (MacLeod & Summerfield, 1990; Middelweerd & Plomp, 1987) and 8-10 dB for open-set word recognition (Ross, Saint-Amour, Leavitt, Javitt, & Foxe, 2007). The source of this benefit is not well understood, as reflected by the fact that mechanistic accounts of audiovisual speech perception disagree with respect to the type of information extracted from each sensory modality and the level of processing at which this information is combined (R. Campbell, 2008; Massaro & Simpson, 2014; Rosenblum, Pisoni, & Remez, 2005; Schwartz, Basirat, Ménard, & Sato, 2012). One issue of contention is the relative importance of static (time-invariant) versus dynamic (time-variant) visual speech cues (Rosenblum & Saldana, 1998).

Dynamic visual speech cues encompass the kinematic trajectories of the facial articulators – i.e., the displacement, velocity, and acceleration of a talker’s lips, tongue, jaw, cheeks, and head (Brooke & Summerfield, 1983; K. G. Munhall, Jones, Callan, Kuratate, & Vatikiotis-Bateson, 2004) – which are available to the observer as a visual motion signal. These dynamic cues can be encoded with relatively low visual resolution. In contrast, static visual speech cues are those which can be captured in a single frame (e.g., still facial configurations), and which require detailed visual information (e.g., color/luminance contrast or texture) to distinguish between the facial articulators (e.g., lips vs. teeth and tongue). Static cues are particularly relevant to lip-reading or “speechreading”, and thus are related to the concept of visemes – essentially phoneme-like, minimally distinguishable units of visual speech (Fisher, 1968). Many viseme classes can be described by a still image (Chen & Rao, 1998), and some vowels and consonants can be identified from still photographs of faces during certain phases of

articulation (Ruth Campbell, 1986; Ruth Campbell, Landis, & Regard, 1986). In addition, static images of visible speech gestures produce neural activation in similar cortical speech networks as those evoked by videos of a moving talker (G. Calvert & Campbell, 2003). In the context of audiovisual speech recognition in noise, the magnitude of the audiovisual advantage in individual listeners is predicted by speechreading scores (MacLeod & Summerfield, 1987, 1990), and the magnitude of the audiovisual benefit for sentences is reduced by approximately half when visual speech is presented as point-light displays (i.e., without static cues; Rosenblum, Johnson, & Saldana, 1996). However, the fact that point-light stimuli produce any significant audiovisual gain at all (up to 10 dB in individual listeners) may, in fact, suggest a strong role for dynamic cues in the audiovisual speech recognition advantage (Rosenblum et al., 1996). Indeed, when videos of a talker's face are blurred sufficiently to reduce closed-set (six-item) speechreading scores by roughly half (Thomas & Jordan, 2002), the magnitude of the audiovisual benefit for sentences in noise remains unchanged (K. Munhall, Kroos, Jozan, & Vatikiotis-Bateson, 2004).

In fact, evidence that dynamic cues play a role in visual speech perception is abundant. Visual interference and/or “McGurk” illusory effects (fusion of conflicting auditory and visual phonetic content; McGurk & MacDonald, 1976) are observed with point-light stimuli (Rosenblum & Saldaña, 1996), when visual detail is removed by severe blurring or spatial quantization (MacDonald, Andersen, & Bachmann, 2000; Thomas & Jordan, 2002), and when visual speech is observed from distances up to 20 meters (Jordan & Sergeant, 2000) or in non-foveal regions of the visual field (Paré, Richler, ten Hove, & Munhall, 2003). Moreover, observers can identify particular talkers – and even learn the identity of particular talkers – from dynamic visual speech information alone (Girges, Spencer, & O'Brien, 2015; Jesse & Bartoli,

2018; Rosenblum, Niehus, & Smith, 2007; Rosenblum, Smith, Nichols, Hale, & Lee, 2006; Rosenblum et al., 2002), and the natural movement of a talker's head alone provides a significant boost to audiovisual speech recognition in noise (K. G. Munhall et al., 2004).

One theory concerning the dual contributions of static and dynamic visual speech cues is that static cues provide complementary information that can substitute for auditory speech information that is hard to hear, while dynamic visual speech cues provide correlated information that enhances heard speech by exploiting cross-modal redundancies between time-varying auditory and visual speech signals (R. Campbell, 2008; Rosenblum et al., 2005; Q. Summerfield, 1987). Regarding the latter, several recent studies have shown that facial kinematic trajectories – particularly of the lips/mouth – are often highly correlated with the acoustic speech temporal and spectral envelopes, and the fundamental frequency contour (Chandrasekaran, Trubanova, Stillitano, Caplier, & Ghazanfar, 2009; Grant & Seitz, 2000; H. Yehia, Rubin, & Vatikiotis-Bateson, 1998; H. C. Yehia, Kuratate, & Vatikiotis-Bateson, 2002). Grant and Seitz (2000) showed that the ability of a listener to detect the presence of an auditory sentence in noise is improved when a congruent visual speech signal is presented (1.6 dB average effect), and, crucially, that the size of the improvement (0.8-2.2 dB) depends on the correlation between the lip area function (size of the mouth opening over time) and the acoustic speech envelope. They discussed two possible explanations for this effect: first that the visual speech signal may provide timing cues that facilitate the detection of peaks in the auditory speech signal, and second that significant cross-modal comodulation may serve as a cue for selecting intervals that contain auditory speech.

Recent neurophysiological data may provide support for the second explanation – namely that studies show congruent visual speech improves the ability of neuronal oscillations in the

auditory cortex to track the acoustic speech envelope of a target utterance presented in steady state noise (Crosse, Di Liberto, & Lalor, 2016) or competing speech (Golumbic, Cogan, Schroeder, & Poeppel, 2013). However, a number of follow-up behavioral studies investigating the audiovisual speech detection advantage reported by Grant and Seitz (2000) suggest that the effect is not specific to speech and depends largely on visual speech “pre-cuing” the presence of upcoming auditory speech (L. E. Bernstein, Auer, & Takayanagi, 2004; Eskelund, Tuomainen, & Andersen, 2011; Kim & Davis, 2004). An additive-noise ideal observer analysis of the audiovisual speech detection advantage revealed that visual speech improves detection by increasing sampling efficiency and reducing uncertainty about the auditory speech signal, rather than decreasing internal noise as would be expected under an integrative mechanism. This is consistent with physiological data showing that visible biological (including facial and speech) motion reduces the latency of responses in the early auditory cortex when the visual signal acts as a pre-cue (Chandrasekaran, Lemus, & Ghazanfar, 2013; Stekelenburg & Vroomen, 2007; van Wassenhove, Grant, & Poeppel, 2005). In fact, pre-cuing can also explain enhanced speech tracking in the auditory cortex during audiovisual presentation, which according to one prominent explanation occurs because visual speech cues reset the phase of ongoing neuronal oscillations to ensure that auditory inputs arrive during a “high excitability state” (Schroeder, Lakatos, Kajikawa, Partan, & Puce, 2008).

These results appear to imply that correlated time-varying information across auditory and visual speech channels is important only insofar as it implies a consistent relation between the onsets of corresponding visual and auditory events (Chandrasekaran et al., 2009; Schwartz, Berthommier, & Savariaux, 2004; Schwartz & Savariaux, 2014). However, this notion was challenged by a recent magnetoencephalography study, which showed that redundant (i.e.,

correlated) time-varying information across visual (lip area function) and auditory (broadband envelope) speech channels is uniquely encoded by the left posterior superior temporal sulcus (pSTS) – a known, high-level audiovisual integration region for speech (Arnal, Morillon, Kell, & Giraud, 2009; Beauchamp, Nath, & Pasalar, 2010; G. A. Calvert, Campbell, & Brammer, 2000; Jonathan H. Venezia et al., 2017) – and that the strength of such encoding predicts performance in a competing speech task (Park, Ince, Schyns, Thut, & Gross, 2018). Moreover, a model of audiovisual-speech intelligibility was recently developed in which the time-varying visual trajectories of multiple facial effectors (e.g., lips, jaw, and cheeks) were optimally combined with the envelopes of narrowband cochlear filters to reconstruct the original envelopes of noisy speech, where better reconstructions would predict better intelligibility; this model provided a good qualitative fit to audiovisual speech recognition performance across a range of signal-to-noise ratios (Grant, Bernstein, & Summers, 2013). Therefore, the correlation between dynamic visual and auditory speech information may be a crucial contributor to audiovisual speech *recognition* in noise, which is distinct from audiovisual speech *detection* in noise (Eskelund et al., 2011; Schwartz et al., 2004).

To summarize the information presented thus far: (1) static and dynamic visual speech cues contribute to audiovisual speech recognition in noise; (2) one proposal is that static cues provide complementary information while dynamic cues provide correlated information relative to the auditory speech signal; (3) a potentially important source of correlated information is the temporal comodulation of time-varying kinematic trajectories (e.g., the lip area function) and time-varying changes in acoustic speech energy (e.g., the speech envelope); (4) such correlations have been used as an explanation for the audiovisual speech detection advantage and the audiovisual speech recognition advantage, although recent data suggest a potentially larger role

for correlated information in the latter. Despite the rather large body of literature on audiovisual speech perception, little if any work has been devoted to the systematic quantification of recognition benefits conferred by correlated audiovisual information, including the extent to which such benefits depend on the presence of static visual speech cues. Therefore, the goals of the present study were twofold: (1) measure the relative contributions of static and dynamic visual speech cues to audiovisual speech recognition in noise across a full range of SNRs; and (2) operationalize dynamic cues in terms of the degree of correlated audiovisual information.

Toward these ends, participants were asked to recognize audiovisual speech in steady-state, speech-shaped noise across a range of SNRs in the following conditions (Figure 1): (1) Clean: full range of audiovisual cues available; (2) Blurred: face of the talker was blurred with a Gaussian filter to reduce the contribution of static cues; (3) Filtered: temporal modulations from 3-7 Hz were removed from the speech envelope to reduce the temporal coherence between visual and auditory signals, thus reducing the contribution of correlated dynamic cues; and (4) Blurred/Filtered: combined facial blurring and modulation filtering to reduce the contribution of static cues and correlated dynamic cues. Separate groups of participants were tested in the blurred (2/4) and unblurred (1/3) conditions. Each participant was also tested in auditory-only versions of the Clean and Filtered speech conditions. For each condition, SNRs along the full range of a cumulative Gaussian psychometric function were sampled adaptively according to the Updated Maximum Likelihood (UML) procedure (Shen, Dai, & Richards, 2015). Psychometric functions were then fit to the data from all participants in all conditions using a Bayesian hierarchical regression procedure. The crucial dependent measure for each condition was the magnitude of the audiovisual advantage, operationalized as both the shift in the threshold of the psychometric function (in dB) and the maximum difference in percent correct at any SNR. We

also fit a ‘guess rate’ parameter that essentially measured lip-reading performance (i.e., percent correct at arbitrarily low SNRs). We hypothesized that the magnitude of the audiovisual advantage would pattern by condition as follows: $1 > 2 > 3 > 4$. Thus, the largest advantage was expected in condition 1, for which neither the auditory nor visual speech signal was corrupted. A larger advantage was expected in condition 2 (facial blurring, clean audio) relative to condition 3 (clear face, filtered audio) because the degree of facial blurring applied in condition 2 was expected to reduce static cues while largely preserving the audiovisual recognition advantage (K. Munhall et al., 2004). The smallest advantage was expected in condition 4, for which the speech signals in both sensory modalities were corrupted (facial blurring, filtered audio). In fact, no difference in either measure of audiovisual advantage was observed across conditions. This was unsurprising for condition 2 (blurring) but was not expected for conditions 3 and 4 (filtering), for which audiovisual temporal coherence was severely disrupted.

Figure 1. Experiment 1 Design. The four experimental conditions are shown in a 2x2 table indicating whether (YES) or not (NO) static and dynamic visual speech cues were expected to play a large role in the audiovisual speech recognition advantage. In the left margin of the table, examples of control and modulation-filtered auditory stimuli (A) are shown in the modulation power spectrum domain. Temporal modulation rate (Hz) is plotted along the abscissa and spectral modulation rate (cyc/kHz) along the ordinate. The color scale denotes modulation power in dB (arbitrary ref.). In the top margin of the table, examples of clear and blurred visual speech stimuli (V) are shown.

A follow-up experiment was conducted to rule out potential unintended effects of the filtering manipulation in the first experiment. In the follow-up experiment, subsets of utterances with naturally high vs. low audiovisual temporal coherence were selected, and audiovisual speech recognition performance was compared across these subsets. Only a small difference in performance was detected between the high- and low-coherence subsets. A model of audiovisual speech intelligibility based on correlated dynamic cues predicted large effects of audiovisual temporal coherence for both experiments that were not borne out in the experimental data. These findings call into question a major role for correlated dynamic cues in the audiovisual speech recognition advantage, and at least suggest that temporal coherence must be operationalized

differently in future research aimed at characterizing the effects of coherence on speech recognition performance.

Experiment 1: Direct Manipulation of Static and Dynamic Visual Speech Cues

Methods

Participants

A total of 16 (10 female) participants (mean age = 21.4, range = 18-31) were recruited to take part in Experiment 1. Of these participants, 13/16 were right-handed and 10/16 were native speakers of English. All participants were fluent speakers of English as determined by a brief interview with an experimenter (i.e., participant reported speaking English as a second language with exposure to native speakers from an age < 10 years). Audiometric thresholds were 20 dB HL or better from 250-8000 Hz in both ears for all participants. All participants had normal or corrected-to-normal vision. Participants were students enrolled at University of California (UC) Irvine and received either course credit or \$10/hr. for their participation. These students were recruited via the UC Irvine Human Subjects Lab or paper fliers placed around the UC Irvine campus. Oral informed consent was obtained from each participant in accordance with the UC Irvine Institutional Review Board guidelines.

Stimuli

The speech stimuli used in this study were adapted from the Hearing in Noise Test (HINT; Nilsson, Soli, & Sullivan, 1994). The HINT consists of 250 sentences divided into 25 phonemically balanced lists. The HINT sentences were originally drawn from the Bamford-

Kowal-Bench (BKB) sentences (Bench, Kowal, & Bamford, 1979), and revised from the British English versions to equate length and content for American English. To generate an audiovisual version of the HINT, we recorded a female actor from the UC Irvine Drama Department while she produced the HINT sentences with a natural rhythm, intonation, and visual prosody (i.e., head motion). Sentences were produced one after another in a single recording session with the inter-sentence interval freely determined by the actor. Digital video recordings were made in a quiet room against a green screen (resolution = 1920x1080, frame rate = 30/s). The actor was filmed from the shoulders up with approximately 75% of the scene devoted to the background screen on either side of the actor. Audio was simultaneously recorded on separate microphone, digitized (44.1 kHz, 16-bit), and synced to the main video sequence in post-processing (a deliberately produced acoustic transient was manually aligned to the same acoustic event in the camera audio stream). Post-processing of the aligned audiovisual sequence was carried out using Final Cut Pro v6 (Apple, Inc). The green screen was replaced with a solid grey background using a chroma key procedure and the video sequence was cropped to 640 x 480 pixels. After cropping, the actor was centered in the middle of the screen and pictured from the neck up. The actor's face occupied approximately 300 vertical pixels and 240 horizontal pixels. Individual sentences were extracted manually from the post-processed sequence. Approximately 10-20 'still frames' (i.e., not including any speech-related visual motion) were included at the beginning and end of each clip. The individual clips were exported to Audio Visual Interleaved (.AVI) format, and the visual (still image sequence) and audio (.wav file) components of each clip were saved separately for presentation in the experiment. The audio files were normalized to equal root-mean-square amplitude after excluding silent portions ('still frames') at the beginning and end of each file. An 8.2 minute segment of steady state, speech-shaped noise

(SSN) was generated by concatenating all the audio files and filtering Gaussian noise to match the long-term spectrum of the concatenated speech signal using an overlap-add method (MATLAB `fftfilt.m`).

To reduce the contribution of static visual cues, blurred versions (BLUR) of the original still-image sequences (CLEAR) of all 250 HINT sentences were produced by applying a symmetric Gaussian low pass filter with a cutoff frequency of 5.5 cycles per (horizontal) face (i.e., Gaussian $sd = \frac{240 \text{ pixels}/5.5}{2\pi}$). Crucially, this manipulation did not significantly affect dynamic visual cues, which remain available to the observer even at relatively low visual resolution (K. Munhall et al., 2004; Rosenblum & Saldana, 1998). To reduce the contribution of correlated dynamic visual cues, the acoustic speech files were distorted in order to reduce or eliminate temporal covariation between acoustic and visual speech signals. Specifically, fluctuations in the acoustic envelope from 3-7 Hz were removed from the acoustic speech signal by means of a notch filter imposed in the 2D modulation spectrum domain (Elliott & Theunissen, 2009). Control versions (CTRL) of the modulation-filtered sentences (FILT) were generated by applying the same signal processing steps without removing the 3-7 Hz envelope fluctuations. The FILT signals were significantly distorted by this procedure but remained intelligible in quiet. However, acoustic information across the range of modulation frequencies that co-vary maximally with dynamic visual speech cues (Chandrasekaran et al., 2009) was eliminated.

Audiovisual temporal covariation is defined here in terms of the magnitude-squared coherence (henceforth “coherence”) between the acoustic speech envelope and the interlip distance function (i.e., the frame-by-frame measurement of the distance between the lips in pixels). Coherence is essentially a measure of correlation in the spectral domain – it describes the fraction of the total (audiovisual) signal power that is shared via linear relation by the interlip

distance function and the acoustic envelope at given spectral analysis frequency. A previous study (Chandrasekaran et al., 2009) calculated audiovisual coherence for several large sets of speech materials, showing that narrow-band speech envelopes are maximally coherent with the mouth area function at roughly the syllabic rate (~ 2 -7 Hz), and in acoustic frequency bands from about 300-4000 Hz. The mouth area function is highly correlated with the interlip distance function (Chandrasekaran et al., 2009; Grant & Seitz, 2000; Q. Summerfield, 1992), so these measures likely reflect the same underlying kinematic signal.

Here, the interlip distance function was measured via a semi-automated procedure utilizing the MATLAB Computer Vision Toolbox (Figure 2). For each HINT sentence, an experimenter manually identified pixel coordinates corresponding to regions of high color contrast on the outer edges of the upper and lower lips using the first video frame. Sets of seven points near the midline of the upper and lower lips were then defined by including the manually identified coordinates along with six additional points, three each at neighboring pixels to the left and right. These points were then automatically tracked frame-by-frame using the Kanade-Lucas-Tomasi feature tracking algorithm (Tomasi & Kanade, 1991). Simultaneously, the position of the talker's face was tracked frame-by-frame using the Viola-Jones face detection algorithm (Viola & Jones, 2004). On each frame, the face tracker was used to estimate the geometric distortions (caused by head motion) between the current frame and previous frames. These transformations were applied to the output of the point tracker to place the tracked lip points into a common coordinate space. Within this common space, the interlip distance was taken to be the maximum vertical pixel location of the points on the upper lip minus the minimum vertical pixel location of the points on the lower lip. In some cases, the point tracker failed to track any of the seven upper and/or lower lip points on a given frame. In these cases, an

experimenter manually re-started the tracking procedure using a new set of starting coordinates, and this was repeated until the interlip distance was successfully tracked for each frame. The interlip distance function for each HINT sentence was then manually reviewed by an experimenter by viewing the tracked lip coordinates superimposed on the actual video frame-by-frame, and cross-checking this information against the estimated value of the interlip distance function for a given frame.

Figure 2. Computation of the interlip distance function. Eight visual frames from an example AVHINT are plotted. For each frame, the yellow box around the talker's face indicates the facial position estimated by the Viola-Jones face detection algorithm. The white dots above and below the talker's lips mark features tracked by the Kanade-Lucas-Tomasi algorithm. The vertical distance (pixels; corrected for facial position) between the white dots denotes the interlip distance for a given frame. Plotted above each frame is the time-course of interlip distance (blue) with a red circle denoting the interlip distance calculated from the corresponding frame.

To measure audiovisual coherence (Figure 3A), each acoustic HINT sentence was passed through a 128-band gammatone filterbank (Institute of Sound Recording MATLAB toolbox: gammatoneFast; <http://iosr.surrey.ac.uk/software/index.php>) with center frequencies evenly

spaced on the ERB-rate scale (Moore & Glasberg, 1983) from 100-8000 Hz. Within each band, the envelope was extracted as the magnitude of the analytic signal, low-pass filtered with a 15 Hz cutoff (2000-order finite impulse response), shifted to account for the low-pass filter delay, and downsampled to 30 Hz. A measure of band-specific coherence between the downsampled acoustic envelope and the interlip distance function was then generated. Samples corresponding to ‘still frames’ at beginning and end of the signals were removed, and coherence was calculated using the MATLAB function *mscohere* with a 24-sample hamming analysis, 20 frames of overlap, and a 32-point fast Fourier transform. The resulting output vector, whose values ranged from 0-1, described the fraction of total audiovisual power shared by the acoustic and visual signals at temporal modulation frequencies from 1-15 Hz within a given gammatone frequency band. Coherence across all gammatone frequency bands can be described in two dimensions (modulation frequency x gammatone band frequency). This 2D representation was calculated for all 250 CTRL sentences and all 250 FILT sentences (Figure 3B). Notably, audiovisual coherence was drastically reduced for FILT sentences within the range of temporal modulations that had been removed from the signal (3-7 Hz). This confirms that correlated dynamic visual speech cues were significantly reduced by the modulation filtering manipulation, at least when such cues are operationalized in terms of audiovisual coherence.

Figure 3. Audiovisual coherence (AVC). (A) Schematic of the computation of AVC in the 1000 Hz-center-frequency gammatone band. Bandpass (gammatone) filtering is followed by envelope extraction and downsampling. The AVC between the band-specific envelope and the interlip distance function is then obtained by calculating the magnitude squared coherence. (B) Mean AVC of AVHINT sentences. The mean AVC across all 250 AVHINT sentences is plotted for the control (CTRL) and modulation-filtered (FILT) stimulus conditions. Temporal modulation rate (Hz) is plotted along the abscissa and gammatone band center frequency (Hz) along the ordinate. The color scale indicates the magnitude of AVC.

Procedure

Audiovisual stimuli were displayed and responses collected using Psychtoolbox-3 (Kleiner et al., 2007) in MATLAB R2015a running on Ubuntu Linux 12.04 LTS. The native sound card and 14-inch display of a System76 Galaga Pro laptop were used for stimulus delivery. Acoustic stimuli were presented diotically over Sennheiser HD 280 Pro circumaural earphones. All testing was carried out in a double-walled, sound-attenuated booth. Participants completed the experiment in a single session lasting approximately 1-1.5 hours. Prior to the experiment, participants' pure-tone thresholds were measured at audiometric frequencies from 0.25-8 kHz using an automatic screening audiometer.

The experimental task was to view/listen to audiovisual or auditory-only sentences in the presence of an SSN masker, and to reproduce via keyboard any words from the sentence that were recognized. Participants were instructed to type out what they heard after each trial, to pay attention to videos of the talker if and when they were present, and to attempt to make use of the visual information to help them understand the noisy speech signal. A trial was initiated by pressing 'Enter' on the keyboard. On each trial, a single HINT sentence was selected from the set of 250 possible sentences at random without replacement. A portion of the 8.2 min SSN file with a pseudo-randomly selected start point and duration equal to the selected HINT sentence was excised and added to the auditory sentence stimulus at a signal-to-noise ratio (SNR) determined by an adaptive procedure (see below).

On audiovisual trials, the visual and auditory components of the selected HINT sentence were presented simultaneously by synchronizing presentation of each signal to the vertical refresh of the laptop screen. The video of the talker was presented in the center of the screen against a grey background. The talker's face subtended $\sim 4^\circ$ of visual angle. On auditory-only

trials, only the auditory component of the selected HINT sentence was presented while the laptop screen remained blank with a solid grey background. After presentation of the target sentence, the participant was cued to respond by typing any words s/he understood into an edit box on the laptop screen. After the participant confirmed her/his response, the typed text remained on the screen but could not be edited. At this point, several boxes appeared on the screen, each of which was labeled using a keyword from the target sentence presented on that trial. The participant was instructed to select the keywords that were correctly identified in their typed response by clicking the corresponding buttons. By instruction, errors in tense or obvious typos were counted as correct. All words except “*the*”, “*is*”, “*for*”, “*to*”, “*a*”, “*was*”, “*an*”, “*were*”, and “*are*” were counted as keywords. The next trial began after the self-scoring process was complete.

The target auditory speech signal was always presented at an overall level of 55 dBA. The SNR on each trial was varied by adjusting the level of the SSN masker according to the updated maximum-likelihood (UML) adaptive procedure (Shen, Dai, et al., 2015) as described in Appendix 1. Briefly, the UML procedure is a Bayesian adaptive procedure that efficiently samples “sweet points” along a psychometric function to minimize the expected variance of the psychometric function parameters (i.e., to produce best-fitting parameter estimates). The psychometric function parameters are updated from trial to trial based on the listener’s current response (whether each keyword is correct or incorrect). Here, the UML was used simply to generate an optimal scheme for selecting the SNR from trial to trial. On each trial, the SNR was adjusted by sampling from one of several sweet points along the current best-fitting psychometric function. Sweet point selection was determined by an up-down adaptive procedure and each adaptive track was terminated after 16 reversals. Due to a programming error, the sweet point at the upper asymptote of the psychometric function (see Appendix 1) was not

sampled in auditory-only conditions. This may have biased the psychometric-function slope estimates in the auditory-only conditions, but this effect was consistent across the experimental manipulations of interest (CTRL vs. FILT, and BLUR vs. CLEAR). For the purpose of the UML procedure, each keyword was treated as a separate “trial”, such that more than one reversal could occur within a single sentence with multiple keywords (presented at a single SNR). In real time, the correctness of each keyword was taken to be the self-scored response generated by the participant. A semi-automated spell-check algorithm was used to validate the self-scoring. The minimum percent agreement between the original (self-scored) and validated (corrected) responses was 97%. Only the corrected responses were used in the final data analysis. For audiovisual conditions, participants performed three extra trials at the lowest possible SNR (-25 dB) at the end of the adaptive track. These trials allowed for a more robust estimate of the lower asymptote of the psychometric function (i.e., ‘lipreading’ ability).

Each participant completed two adaptive tracks in each of four conditions: (1) Auditory-only speech with CTRL sentences; (2) Audiovisual speech with CTRL sentences; (3) Auditory-only speech with FILT sentences; and (4) Audiovisual speech with FILT sentences. Thus, each participant completed eight adaptive tracks in total (~ 25 sentences per track). Half of the participants ($n = 8$) completed audiovisual trials with BLUR videos, while the remaining half of the participants completed audiovisual trials with CLEAR videos. A single adaptive track was performed in each condition (1-4) with the condition order determined by a 4x4 Latin square, and a second adaptive track was then performed in each condition with the condition order determined by a mirror-reflection of the Latin square. The same Latin square was repeated for each block of four participants within each of the BLUR/CLEAR groups. To familiarize participants with the trial structure and experiment interface, 10 auditory-only practice trials

were performed at the highest SNR (+20 dB) prior to data collection. Once data collection was complete, psychometric functions were re-fit to the entire dataset using a Bayesian hierarchical regression procedure.

Data Analysis

Psychometric functions were fit to the binary response outcomes (correct/incorrect) across all observations (keywords), experimental conditions, and participants within a single model using Bayesian hierarchical regression. The package *brms* (Bürkner, 2017) v2.3.1 was used to perform the regression analysis within the R statistical computing environment (R Core Team, 2013) v3.4.4. The *brms* package implements Bayesian regression models in the *Stan* programming language (Carpenter et al., 2017), which uses No-U-Turn sampling (NUTS) (Hoffman & Gelman, 2014) to estimate the posterior distribution of model parameters. Broadly, the hierarchical or “mixed model” approach to modeling psychometric functions (Houpt & Bittner, 2018; Moscatelli, Mezzetti, & Lacquaniti, 2012) aims to derive psychometric function parameters at the population level (i.e., across participants; “fixed effects”) via generalized linear regression. The psychometric function parameters for individual participants are characterized as deviations from the population-level parameters (“random effects”). These random effects are represented by model parameters that capture the variance of the population-level parameters. Co-variance among population-level parameters can also be modeled. Thus, in the Bayesian framework, we can envision individual-participant psychometric function parameters as being sampled from (potentially inter-related) population-level distributions of psychometric function parameters. Experimental manipulations (e.g., the effect of condition) can be modeled in terms of fixed or random effects – i.e., differences in psychometric function parameters by condition at the population and individual-participant levels, respectively.

The formal specification of our model is given in Appendix 2. Briefly, performance (percent keywords recognized) was characterized in terms of a cumulative Gaussian psychometric function as follows:

$$\psi = \gamma + (1 - \gamma - \lambda) * 0.5 * \operatorname{erfc}\left(\frac{-b*(SNR-a)}{\sqrt{2}}\right),$$

where ψ is the psychometric function describing percent correct keyword recognition as a function of SNR , γ is the ‘guess rate’ (essentially ‘lipreading performance’ in the context of audiovisual speech), λ is the ‘lapse rate’, a is the psychometric function threshold, b is proportional to the psychometric function slope, and erfc is the complementary error function. The lapse rate parameter, λ , was fixed at 0.02. For a and b , fixed (population-level) effects were estimated for each cell of the overall experimental design (i.e., cell means were estimated): presentation modality (auditory/audiovisual) x filter condition (CTRL/FILT) x blur condition (CLEAR/BLUR). Random (subject-level effects; deviations from the population means) were estimated for each cell of the within-subject experimental design: presentation modality (auditory/audiovisual) x filter condition (CTRL/FILT). For γ , fixed effects were estimated for each of the following design cells: presentation modality (auditory/audiovisual) x blur condition (CLEAR/BLUR). However, γ was essentially forced to zero in the auditory-only conditions by setting a very narrow prior. Therefore, random effects were estimated only for the audiovisual presentation modality (i.e., subject-level deviation from the blur-group-specific mean across all audiovisual conditions).

To estimate the Bayesian hierarchical regression model, we ran 10 independent sampling chains, each with 4000 warmup samples and 2000 post-warmup samples, resulting in 20000 total draws from the posterior. The values of `adapt_delta` and `max_treedepth` were set to 0.9 and 10,

respectively, and all parameters were initialized to a randomly selected value between -2 and 2. No divergent transitions were observed. Following Hout et al. (2018), two criteria were used to ensure the samples adequately represented the posterior distribution: (1) visual checks indicated good convergence and good mixing and the Gelman-Rubin $\hat{R}$ was less than 1.01 for all parameters; and (2) the smallest effective sample size (ESS; Carpenter et al., 2017) across all parameters was still adequately large to characterize the posterior density. In Hout et al. (2018), the minimum ESS was 650 (out of 4000 total samples). Our minimum ESS over all parameters was 7220 (out of 20000 total samples).

The primary outcome measure in this experiment was the audiovisual speech recognition advantage, which was quantified in two complementary ways: (1) the improvement in the speech reception threshold (SRT; the SNR for which proportion correct keyword recognition is 0.5) for audiovisual speech vs. auditory-only speech, expressed in dB; and (2) the maximum improvement in proportion of keywords correctly recognized (audiovisual > auditory-only) at any given SNR. The improvement in SRT was calculated by determining the 50% correct point of the psychometric function – described by the parameters a , b , and γ determined during the model-fitting step – from its analytical inverse. The maximum improvement in proportion correct was determined by subtracting the fitted values of the auditory-only psychometric function from the fitted values of the audiovisual psychometric (SNR grid from -25 to 20 dB in steps of 0.1) and determining the maximum difference across SNR steps. These metrics were calculated for each of the 20000 draws from the posterior across the four conditions in the experiment: (1) CLEAR videos with CTRL sentences; (2) BLUR videos with CTRL sentences; (3) CLEAR videos with FILT sentences; and (4) BLUR videos with FILT sentences. Differences in the outcome measures by condition were assessed by constructing posterior

distributions for each possible pairwise linear contrast of conditions. One-tailed contrasts were performed according to the hypotheses described in the Introduction: $1 > 2$, $1 > 3$, $1 > 4$, $2 > 3$, $2 > 4$, and $3 > 4$, where “ $>$ ” indicates a larger audiovisual advantage. A decision was made to reject the null hypothesis (contrast coefficient < 0) if at least 95% of the posterior distribution fell above zero. This is somewhat akin to performing a t-test by determining whether a one-tailed 95% (frequentist) confidence interval contains zero, but the Bayesian test directly addresses the posterior believability of the contrast coefficient exceeding zero given the observed data, while the frequentist test addresses the probability of observing a value of the contrast coefficient as extreme or more extreme than the actually-observed value under the null hypothesis. For all contrasts, the following information is reported: the posterior mode of the contrast coefficient (via kernel density estimation), the posterior 95% highest density interval (HDI) of the contrast coefficient, and the probability that the contrast coefficient is less than or equal to zero (p_{H-}). For directional contrasts, p_{H-} can be considered a Bayesian one-sided p-value, where decisions to reject the null hypothesis based on p_{H-} are approximately equivalent to a Bayesian hypothesis test for the direction (positive or negative) of an effect based on the Bayes factor (Marsman & Wagenmakers, 2017). A test of direction is not the same as test of existence – i.e., it does not assess whether an effect is or is not equivalent to zero, and therefore assumes that the effect is non-zero. By the criteria described above, directional contrasts were treated as “significant” (i.e., the data support the existence of a positive effect) when $p_{H-} < 0.05$.

Statistical Power and Sample Size

The primary goal of the Bayesian statistical analysis was to estimate the magnitude and precision of the effects of interest (Kruschke & Liddell, 2018), rather than to make binary decisions about their presence or absence. Therefore, sample size was selected based on the

study of Shen, Manzano & Richards (2015), which found relatively little variability in the SRT for auditory sentences in steady-state SSN – as estimated via the UML procedure – across a group of young, normal-hearing listeners ($n = 9$). The standard deviation of the SRT across listeners was approximately 1.5 dB. Based on this estimate of variability, a simulation was performed (Appendix 3A) to determine the power of the present experiment to detect a range of within- and between-subject differences in the audiovisual advantage (improvement in SRT) using the Bayesian one-sided p-value (p_H). Although binary hypothesis testing was not the primary goal of the statistical analysis, the assessment of power across different potential effect sizes provides an estimate of the minimum effect that could be reliably detected for the selected sample size ($n = 16$; 8 per listener group). Table 1 shows the estimated power for within- and between-subject effects ranging from 1-4 dB in 0.5 dB steps. Power was greater than 0.8 for within-subject effects of 2.5 dB and larger and between-subject effects of 3 dB and larger. This suggests the experiment was sufficiently powered to detect medium to large effects. These estimates of power are likely conservative because, unlike the true statistical analysis, the simulation procedure did not estimate psychometric function parameters within the same model used for statistical testing (Appendix 3A). Thus, the simulation model likely produced less precise estimates of the audiovisual advantage at the population (fixed-effects) level (Hout & Bittner, 2018).

Table 1. Power to detect differences in the audiovisual advantage (dB) using the Bayesian one-sided p-value for within-subject and between-subject contrasts.

		Effect Size (dB)						
		1	1.5	2	2.5	3	3.5	4
Within-Ss		0.27	0.44	0.62	0.82	0.92	0.98	0.99
Between-Ss		0.18	0.37	0.54	0.69	0.82	0.90	0.97

Audiovisual Intelligibility Model

A simplified version of the model described by Grant, Bernstein & Summers (2013) was used to generate predictions of audiovisual speech intelligibility across the range of SNRs tested in the experiment. This model relies exclusively on dynamic audiovisual cues. Thus, to derive an intelligibility index for a given audiovisual sentence at a given SNR, the interlip distance function and the 128-channel gammatone-frequency-specific auditory speech envelopes were extracted and matched for sampling rate as described for the estimation of audiovisual coherence under ‘Stimuli’ above. Separate estimates of the channel-specific speech envelopes were obtained for clean and noisy versions of the auditory stimulus, where the noisy version consisted of the clean stimulus combined with additive SSN at the specified SNR. Within each gammatone frequency channel, multiple linear regression was used to obtain the best-fitting prediction of the clean speech envelope using the interlip distance function and the noisy speech envelope as predictors. In fact, five interlip-distance predictors were included in the regression model consisting of the original interlip distance function and time-lagged versions of that function across a range of ± 2 frames (± 66 ms). This was done to simulate the relatively broad window of temporal integration for audiovisual speech (Grant & Greenberg, 2001; van

Wassenhove, Grant, & Poeppel, 2007). The predictive weights allocated to the auditory and visual predictors were stored for later analysis.

Once a best-fitting prediction of the clean speech envelope and the corresponding predictive weights were obtained for each gammatone frequency channel, the entire “reconstructed” gammatone frequency spectrogram was compared to the gammatone frequency spectrogram for the clean speech signal to derive an audiovisual intelligibility index as follows:

$$AV_{Index} = 1 - \frac{\|GT - \widehat{GT}\|^2}{\|GT\|^2}$$

where AV_{Index} is the audiovisual speech intelligibility index, GT is the vectorized gammatone frequency spectrogram of the clean speech signal, $\widehat{GT}$ is the best-fitting reconstruction of the clean speech signal’s gammatone frequency spectrogram, and $\|\cdot\|$ is the Euclidean norm. Unlike the calculation of audiovisual coherence (‘Stimuli’, above), which included all video frames from the start to the end of mouth movement (i.e., potentially some frames before and after the onset of acoustic speech), the calculation of AV_{Index} was restricted to frames during which at least some energy was present in the broadband acoustic speech signal. This was done because the clean auditory speech signal, GT , served as the template against which intelligibility was evaluated.

An auditory-only intelligibility index was obtained in the same way as AV_{Index} , but the interlip distance function was omitted from the regression-model predictions of the clean speech envelope. Thus, the audiovisual model predicted intelligibility by obtaining an “internal” reconstruction of the clean speech signal from an optimal weighting of the noisy auditory signal and the time-varying visual signal, and comparing that reconstruction to an “internal”

representation of the actual clean speech signal. The auditory-only model did the same but used only the noisy auditory signal to obtain a reconstruction. Audiovisual and auditory-only intelligibility indices were obtained for both CTRL and FILT sentences. To obtain estimates of audiovisual and auditory-only intelligibility indices across the range of SNRs tested in the experiment, 40 sentences were randomly selected from the stimulus set at each SNR from -25 to 20 in 1 dB steps and processed through the model as described above. The overall audiovisual and auditory-only intelligibility indices at a given SNR were obtained by computing the average index across all 40 sentences processed through the model. In addition, the overall predictive weights on the auditory and visual speech channels at each SNR were obtained by averaging the weight matrices across all 40 sentences processed through the model and taking the maximum auditory and visual weights across all 128 gammatone frequency channels.

The model-derived audiovisual and auditory-only intelligibility functions were naturally sigmoid in shape across the range of SNRs tested. Thus, for ease of display, cumulative Gaussian psychometric functions were fitted to the intelligibility functions using the same parameterization as for the experimental data. To allow direct comparison of the intelligibility functions to the experimental data, the intelligibility functions were scaled such that the auditory-only intelligibility function would range from 0-0.98. The scaling coefficients that produced this range for the auditory-only intelligibility function were used to scale the audiovisual intelligibility function such that the relative magnitudes of the functions would be maintained after scaling. Unlike the intelligibility functions, the pattern of auditory and visual predictive weights across SNRs could not be easily characterized, so for ease of display the weights were smoothed across SNRs by fitting a fourth-order polynomial function.

Results

Bayesian Statistical Model

Population-level (fixed effects) psychometric functions for each condition (auditory-only and audiovisual presentations of CTRL and FILT sentences) in each listener group (CLEAR/BLUR) are plotted in Figure 4A/B. These functions reflect the median of the posterior fitted values of the psychometric functions as derived from the posterior distributions of psychometric-function parameters estimated via Bayesian hierarchical regression. Shaded regions show the 95% Bayesian credible intervals of the fitted values. The fitted values provide an estimate of the proportion of correct keywords identified across the range of SNRs tested in the experiment (-25 to 20 dB). The most noticeable effect is the overall reduction in intelligibility produced by removing 3-7 Hz temporal modulations from the auditory speech signal (i.e., FILT; note the rightward shift of the black and magenta functions compared to the blue and red functions in Figure 4A/B). The overall reduction of intelligibility caused by modulation filtering was expected and not of major interest. Rather, the relative difference in the audiovisual advantage across CTRL (blue vs. red) and FILT (black vs. magenta) conditions (contribution of dynamic cues) and/or CLEAR (left panel) and BLUR (right panel) listener groups (contribution of static cues) was of primary interest. However, such effects of modulation filtering and/or facial blurring were not immediately clear from qualitative inspection of the psychometric functions.

Figure 4. Experiment 1 Bayesian statistical model: Psychometric functions. (A) Population-level psychometric functions for the 'clear face' group. (B) Population-level psychometric functions for the 'blurred face' group. (C) Population-level estimates of the proportion-correct gain for audiovisual speech vs. auditory speech in the 'clear face' group. (D) Population-level estimates of the proportion-correct gain for audiovisual speech vs. auditory speech in the 'blurred face' group. For all plots, signal-to-noise ratio (SNR) is plotted along the abscissa and proportion correct is plotted along the ordinate. The dashed lines indicate the posterior median of the Bayesian-model fitted values and shaded regions indicate the 95% Bayesian credible intervals. A = auditory, AV = audiovisual, ctrl = unfiltered audio, filt = modulation-filtered audio.

The SNR-by-SNR differences between audiovisual and auditory-only psychometric functions are plotted in Figure 4C/D. These functions reflect the median of the posterior differences in the fitted values (audiovisual minus auditory-only) of the psychometric functions for each condition (blue = CTRL, black = FILT) and listener group (left panel = CLEAR, right panel = BLUR). Here, the interesting aspect of the data is the peak height of the difference functions, as opposed to their global position relative to SNR. Again, no qualitative difference in

peak heights across was immediately apparent, with the maximum gain in proportion correct ranging from ~ 0.3 - 0.35 across both conditions and listener groups.

Figure 5. Experiment 1 Bayesian statistical model: Scalar metrics. (A) Population-level posterior distributions of the audiovisual advantage in terms of signal-to-noise ratio shift (dB) at the speech reception threshold (SRT). The population-level prior is also shown. (B) Population-level posterior distributions of the audiovisual advantage in terms of maximum gain in proportion correct (PC). (C) Population-level posterior distributions of the guess rate parameter. The population-level prior is also shown. For all plots, distributions are visualized as 'violin plots' where the estimated density is given by the horizontal distance between the vertical line in the center and either of the vertical curves plotted to the left/right of the center. The interquartile range of the distribution is visualized as a rectangular 'box plot' with the median indicated by a horizontal black bar. Clear = unblurred face, blur = blurred face, ctrl = unfiltered audio, filt = modulation-filtered audio.

Quantitative estimates of the audiovisual advantage were obtained by from the posterior distributions of two scalar metrics: (1) the magnitude of the shift in the SRT (the SNR at which proportion correct is 0.5) for audiovisual compared to auditory-only psychometric functions; and (2) the maximum audiovisual gain in proportion correct (audiovisual minus auditory-only) at any given SNR. In addition to the audiovisual advantage, which presumably reflects a combination or integration of auditory and visual speech cues, visual-only effects of “lipreading” were estimated from the posterior distribution of the ‘guess rate’ – i.e., the lower asymptote of the psychometric function, which reflects the proportion of keywords correctly identified when the SNR is arbitrarily low. A single estimate of the guess rate was obtained for each listener group (CLEAR/BLUR) because the guess rate should not be affected by modulation filtering (i.e., collapsing across CTRL/FILT is justified). Figure 5 summarizes the posterior distributions of scalar metrics representing the audiovisual advantage (SRT shift, max audiovisual gain) and the contribution of lipreading (guess rate). Planned (directional) contrasts between different cells of the experimental design were performed as described under ‘Statistical Analysis’ above. The results of these contrasts are summarized in Table 2.

Table 2. Summary of statistical contrasts for scalar estimates of the audiovisual advantage and the effect of lipreading. SRT = audiovisual advantage in terms of change in the speech reception threshold. Max PC = audiovisual advantage in terms of the maximum gain in proportion correct. Mode = posterior mode. HDI = Bayesian highest density interval. p_H = Bayesian one-sided p-value. PC = proportion correct. $*p_H < 0.05$.

Audiovisual Advantage						
	SRT (dB)			Max PC		
	Mode	95% HDI	p_H	Mode	95% HDI	p_H
CLEAR/CTRL > BLUR/CTRL	1.79	[-0.27, 3.82]	0.04*	0.02	[-0.11, 0.20]	0.29
CLEAR/CTRL > CLEAR/FILT	0.16	[-1.91, 2.29]	0.41	0.07	[-0.04, 0.18]	0.11
CLEAR/CTRL > BLUR/FILT	0.66	[-1.75, 3.02]	0.26	0.08	[-0.07, 0.23]	0.17
BLUR/CTRL > CLEAR/FILT	-1.47	[-3.94, 0.86]	0.91	0.02	[-0.13, 0.19]	0.37
BLUR/CTRL > BLUR/FILT	-1.10	[-2.88, 0.86]	0.88	0.01	[-0.11, 0.18]	0.33
CLEAR/FILT > BLUR/FILT	0.35	[-2.13, 3.27]	0.35	0.00	[-0.15, 0.16]	0.48
Lipreading						
	Guess Rate (PC)					
	Mode	95% HDI	p_H			
CLEAR > BLUR	0.08	[-0.01, 0.21]	0.03*			

The posterior mode of the audiovisual advantage was larger in the CLEAR/CTRL condition than all other conditions in terms of both SRT and maximum gain in proportion correct (Max PC), as expected. However, the only contrast for which the Bayesian one-sided p-value (p_H) was less than 0.05 was CLEAR/CTRL > BLUR/CTRL for the SRT metric, which represents the contribution of static visual cues (i.e., the effect of facial blurring). This appears to be due to an unusually small shift in SRT for the BLUR/CTRL condition compared to all other conditions (see Figure 5A). The effect of facial blurring was expected to be, if anything, smaller than the effect of modulation filtering, but this did not appear to be the case for the SRT

metric (see, e.g., BLUR/CTRL > CLEAR/FILT and BLUR/CTRL > BLUR/FILT in Table 2).

On the other hand, all of the planned contrasts patterned in the expected direction for the Max PC metric, with the largest estimated effects due to modulation filtering (see, e.g., CLEAR/CTRL > CLEAR/FILT and CLEAR/CTRL > BLUR/FILT in Table 2), but $p_H > 0.05$ for all the Max PC contrasts so these results must be interpreted with caution. In general, there were no clear effects of either facial blurring (static cues), modulation filtering (dynamic cues), or their combination. For the majority of the contrasts, including CLEAR/CTRL > BLUR/FILT, which was expected to produce the largest effects, the data did not clearly support a positive or negative effect. There was, however, a clear effect of facial blurring on lipreading (see Table 2, guess rate, CLEAR > BLUR). Lipreading contributed significantly more in the CLEAR condition and was effectively extinguished in the BLUR condition, suggesting that facial blurring effectively removed the contribution of static visual cues.

Figure 6. Experiment 1 audiovisual speech intelligibility model: Psychometric functions. (A) Model psychometric functions. Signal-to-noise ratio (SNR) is plotted along the abscissa and the scaled intelligibility index, AV_{Index} , is plotted along the ordinate. (B) Model estimates of the proportion-correct gain for audiovisual speech vs. auditory speech. Signal-to-noise ratio (SNR) is plotted along the abscissa and the scaled intelligibility index, AV_{Index} , is plotted along the ordinate. (C) Model weights by SNR. Signal-to-noise ratio (SNR) is plotted along the abscissa and the maximum model weight is plotted along the ordinate. A = auditory, AV = audiovisual, V = visual, ctrl = unfiltered audio, filt = modulation-filtered audio.

Audiovisual Intelligibility Model

The psychometric functions predicted by the audiovisual speech intelligibility index (AV_{Index}) are plotted in Figure 6A, along with difference functions showing audiovisual minus

auditory-only (middle row), and maximum auditory and visual predictive weights across SNRs (bottom row). Figure 6A/B can be compared directly to the experimental data shown in Figure 4. Model-derived scalar metrics representing the audiovisual advantage and the contribution of lipreading are plotted in Figure 7, which can be compared directly to the experimental data shown in Figure 5. The audiovisual intelligibility model cannot be used to characterize the effects of facial blurring because it does not incorporate static visual cues, so the model predictions are more directly comparable to the experimental data for the BLUR group. The findings from the audiovisual intelligibility model can be summarized in five points as follows: (1) given the model's simplicity, its predictions are remarkably well in line with the experimental data; (2) the model under-predicts the overall effect of modulation filtering (compare rightward shift of black and magenta functions relative to blue and red functions in Figure 6A vs. Figure 4A/B); (3) the model generates a fairly accurate range of predictions for the audiovisual advantage in terms of the SRT (compare Figure 7A to Figure 5A); (4) the model greatly over-predicts the contribution of visual cues at very low SNRs (Figure 6A lower asymptotes for audiovisual conditions; Figure 7C) compared to the experimental data (Figure 4A/B, lower asymptotes for audiovisual conditions; Figure 5C); (5) the model correctly predicts a modest effect of modulation filtering on the audiovisual advantage, yet the predicted effect is still larger than that observed in the Experimental data (compare especially Figure 7A to Figure 5A). Two consequences of point 4 are that the model-derived difference functions (audiovisual minus auditory-only) decrease monotonically as SNR increases (Figure 6B), whereas these functions are bandpass in shape for the experimental data (Figure 4C/D), and, relatedly, the model-derived estimates of Max PC reflect the difference in PC at the lowest SNR. Thus, the model-derived estimates of Max PC are not particularly meaningful as an indicator of audiovisual integration

per se. However, it is interesting to note that the model predictive weights transition from a visual bias (larger weight on visual cues) to an auditory bias (larger weight on auditory cues) around an SNR of approximately -8 to -10 dB (Figure 6C, blue/red and black/magenta “crossover” points), which is roughly where the audiovisual-minus-auditory difference functions reach a peak for the experimental data (Figure 4D). Thus, the point of crossover of the model weights may be a reasonable indicator of the SNR at which audiovisual integration is maximally effective.

Figure 7. Experiment 1 audiovisual speech intelligibility model: Scalar metrics. (A) Audiovisual advantage in terms of signal-to-noise ratio shift (dB) at the speech reception threshold (SRT). (B) Audiovisual advantage in terms of maximum gain in proportion correct (PC). (C) Guess rate parameter. Ctrl = unfiltered audio, filt = modulation-filtered audio.

Discussion

In Experiment 1, the contributions of static and dynamic visual cues to the audiovisual speech recognition advantage were examined by selectively degrading these cues via facial blurring and modulation filtering, respectively, and observing the effects of these manipulations on two scalar metrics of the audiovisual advantage: the magnitude of the shift in SRT (units of dB) and the maximum gain in proportion correct (Max PC; units of PC). The results of the experiment were largely equivocal. There was weak evidence for an effect of static cues on the SRT metric and for an effect of dynamic cues on the Max PC metric, but neither of these effects suggest a major contribution to the audiovisual advantage. Facial blurring significantly reduced the contribution of lipreading at very low SNRs (i.e., the ‘guess rate’; Figure 5C; Table 2), suggesting that the contributions of static cues were successfully minimized for the BLUR group. However, when the contribution of dynamic cues was minimized by removing 3-7 Hz temporal modulations from the auditory signal and thereby severely disrupting audiovisual coherence (Figure 3), this had remarkably little effect on the audiovisual advantage even for the BLUR group for which static cues had also been degraded (Figures 4-5, Table 2, BLUR/FILT).

Thus, a fundamental question arises from Experiment 1: Why wasn’t the audiovisual advantage extinguished or greatly reduced when both static and dynamic visual cues were degraded/removed from the signal? If we consider static cues alone, this finding is not particularly surprising since previous research had shown that audiovisual speech perception is largely unaffected by facial blurring to the degree imposed in the present experiment (K. Munhall et al., 2004). Therefore, this question reduces to one about modulation filtering: Why is it that modulation filtering produced minimal effects on the audiovisual advantage? One potential explanation is that notch filtering at 3-7 Hz left residual dynamic information in the

auditory signal at frequencies below and/or above the filter cutoffs that remained correlated with the visual speech signal. This residual correlated information, although reduced compared unfiltered audiovisual speech, may have been sufficient to drive the audiovisual advantage. Another potential explanation is that modulation filtering may have created an additional avenue by which visual speech was able to influence performance. Specifically, in addition to enhancement of the auditory speech signal in noise, the information in the visual speech signal may restore some of the auditory information removed by filtering (Shahin, Kerlin, Bhat, & Miller, 2012). Such restorative effects may have offset any reductions in audiovisual enhancement caused by filtering.

Both of these explanations are supported by the results of our audiovisual intelligibility model, which predicted only a modest effect of modulation filtering on the audiovisual advantage. Although the model mischaracterized performance in several minor ways as noted in the Results, the fact that the model predicted only a small effect of modulation filtering suggests there was enough residual information in the auditory and visual speech signals to derive a significant audiovisual benefit despite modulation filtering. Therefore, the results of Experiment 1 suggest that modulation filtering was not an ideal manipulation for assessing the contribution of dynamic visual speech cues. To address this, Experiment 2 measures the effects of naturally occurring differences in audiovisual temporal coherence – i.e., absent any artificial stimulus manipulations – on audiovisual speech recognition in noise.

Experiment 2

The purpose of Experiment 2 was to further characterize the contribution of dynamic visual speech cues as measured in terms of the magnitude of audiovisual temporal coherence. Unlike Experiment 1, which directly manipulated audiovisual temporal coherence by filtering 3-7 Hz temporal modulations from the auditory speech signal, Experiment 2 examined the effect of naturally occurring differences in the magnitude of audiovisual temporal coherence (cf., Grant & Seitz, 2000). Specifically, the set of 250 audiovisual HINT sentences employed in Experiment 1 was separated into two 40-item subsets with the naturally highest and naturally lowest audiovisual temporal coherence. A group of listeners was then presented high-coherence and low-coherence audiovisual sentences in random order, and keyword recognition performance was measured in a background of steady-state SSN at a fixed SNR. The visual speech stimuli were blurred as in the BLUR condition of Experiment 1 to ensure that the contributions of visual speech would be (mostly) limited to dynamic visual cues (i.e., to reduce or eliminate contributions of lipreading). A hierarchical Bayesian regression model was used to test for differences in performance across high- and low-coherence audiovisual sentences.

Methods

Participants

A total of 10 (8 female) participants (mean age = 20.8, range = 19-23) were recruited to take part in Experiment 2. Of these participants, 9/10 were right-handed and 9/10 were native speakers of English. All participants were fluent speakers of English as determined by a brief interview with an experimenter (i.e., participant reported speaking English as a second language, in a well-immersed environment, from an age < 10 years). Audiometric thresholds were 20 dB

HL or better from 250-8000 Hz in both ears for all participants. All participants had normal or corrected-to-normal vision. Participants were students enrolled at University of California (UC) Irvine and received either course credit or \$10/hr. for their participation. These students were recruited via paper fliers placed around the UC Irvine campus. Oral informed consent was obtained from each participant in accordance with the UC Irvine Institutional Review Board guidelines.

Stimuli

The stimuli were drawn from the same 250 audiovisual HINT sentences used for Experiment 1. A 2D representation of audiovisual temporal coherence (gammatone frequency $\times$ temporal modulation rate) was derived for each HINT sentence as described for Experiment 1. A single number describing the overall audiovisual coherence for each sentence was then derived by taking the average of the coherence values across gammatone channels with center frequencies > 700 Hz and < 4750 Hz, and temporal modulation rates > 1.5 Hz and < 8 Hz. This region of the 2D audiovisual coherence space was found to produce the largest average values of audiovisual coherence across the entire set of 250 sentences. Once the overall audiovisual coherence was calculated for each sentence, the sentences were sorted in order from high to low coherence. The 40 sentences with the highest coherence (high-AVC) and the 40 sentences with the lowest coherence (low-AVC) were selected as the target stimuli for this experiment. Figure 8A plots the mean 2D representation of audiovisual coherence for high- and low-AVC items, while Figure 8B shows with histograms plotting the distribution of duration (left) and number of keywords (right) for high- and low-AVC items. The high-AVC items tended to be shorter in duration and have fewer keywords, but there was considerable overlap between the high- and low-AVC distributions. The background SSN for Experiment 2 was generated just as in

Experiment 1. A Gaussian blur was applied to the visual speech stimuli using the same parameters as in Experiment 1.

Figure 8. Experiment 2 high and low audiovisual coherence (AVC) stimulus subsets. (A) Mean AVC across 40-item subsets of high-AVC and low-AVC sentences. Temporal modulation rate (Hz) is plotted along the abscissa and gammatone band center frequency (Hz) along the ordinate. The color scale indicates the magnitude of AVC and is matched to Figure 3B. (B) Left: distributions of sentence duration for high-AVC (blue) and low-AVC items (red). Overlap shown in magenta. Right: distributions of 'number of keywords' for high-AVC (blue) and low-AVC items (red). Overlap shown in magenta.

Procedure

The procedure for Experiment 2 was identical to Experiment 1 except that speech recognition testing was carried out at a constant SNR of -10.5 dB. This SNR was chosen based on the SRT in the BLUR/CTRL condition in Experiment 1 and was found in pilot testing to produce an average performance level of approximately 50% correct for (blurred) audiovisual

speech. Only audiovisual stimuli were presented in Experiment 2. Each participant completed two blocks of 40 sentences (80 total sentences) consisting of the low-AVC and high-AVC items where the condition (high vs. low) was determined pseudo-randomly from trial to trial.

Participants were asked to self-score their responses just as in Experiment 1. A semi-automated spell-check algorithm was used to validate the self-scoring. The minimum percent agreement between the original (self-scored) and validated (corrected) responses was 92%. Only the corrected responses were used in the final data analysis. Prior to data collection, each participant performed five practice trials using sentences drawn from the middle of the AVC distribution.

Data Analysis

To test for differences in the proportion of keywords correctly identified for high-AVC vs. low-AVC items, Bayesian hierarchical logistic regression was implemented in *brms* (Bürkner, 2017) v2.3.1 within the R statistical computing environment (R Core Team, 2013) v3.4.4. The statistical model for Experiment 2 was similar to that for Experiment 1 except that only a single variable was included as a fixed and random factor (*condition*: high- vs. low-AVC) and the link function (see Appendix 2) was the logit function. Also, an intercept term was included in the model such that the intercept term reflected performance in the baseline condition (low-AVC) and the *condition* coefficient reflected the difference in performance between groups (high-AVC minus low-AVC).

To estimate the Bayesian hierarchical regression model, we ran 10 independent sampling chains, each with 2000 warmup samples and 4000 post-warmup samples, resulting in 40000 total draws from the posterior. The values of *adapt_delta* and *max_treedepth* were set to 0.95 and 10, respectively, and all parameters were initialized to a randomly selected value between -2 and 2. No divergent transitions were observed. Following Houpt et al. (2018), two criteria were used to

ensure the samples adequately represented the posterior distribution: (1) visual checks indicated good convergence and good mixing and the Gelman-Rubin $\hat{R}$ was less than 1.01 for all parameters; and (2) the smallest effective sample size (ESS; Carpenter et al., 2017) across all parameters was still adequately large to characterize the posterior density. The minimum ESS over all parameters was 8368 (out of 40000 total samples).

The primary outcome measure in this experiment was the difference in proportion of correct keywords identified for high-AVC vs. low-AVC items. This difference was calculated for each of the 40000 draws from the posterior by converting the fixed-effects regression coefficients to proportion correct for each condition via the inverse logit function and then subtracting the proportion correct for low-AVC from the proportion correct for high-AVC. The resulting posterior distribution of difference scores was treated as the distribution for a one-tailed contrast coefficient (high-AVC > low-AVC). A decision was made to reject the null hypothesis (contrast coefficient < 0) if at least 95% of the posterior distribution fell above zero. The following information is reported for the statistical contrast: the posterior mode of the contrast coefficient (via kernel density estimation), the posterior 95% highest density interval (HDI) of the contrast coefficient, and the probability that the contrast coefficient is less than zero (p_H).

In addition to the statistical test for the direction of the effect (high-AVC > low-AVC), evidence for the existence of the effect (high-AVC - low-AVC $\neq$ 0) was assessed by computing the Bayes Factor via model comparison. A null statistical model was generated by removing the fixed effect of *condition* from the model and estimating the model coefficients just as for the full model. A Bayes factor reflecting the relative goodness-of-fit of the models to the data was computed via bridge sampling (Gronau et al., 2017) in *brms*. As computed, the Bayes factor reflects a ratio as follows:

$$BF_{FULL} = \frac{P(D|M_{FULL})}{P(D|M_{NULL})}$$

where D is the data, M_{FULL} is the full statistical model, M_{NULL} is the null statistical model, and BF_{FULL} is the Bayes such that a ratio above 1 favors the full model and a ratio below 1 favors the null model.

Statistical Power and Sample Size

As for Experiment 1, sample size was selected based on the study of Shen, Manzano & Richards (2015), from which it can be inferred that the standard deviation in proportion correct for a fixed SNR near the mean SRT (i.e., 50% correct) is 0.15, where the standard deviation in the SRT was ~ 1.5 dB and the average slope of the psychometric function was 0.1/dB. Based on this estimate of variability, a simulation was performed (Appendix 3B) to determine the power of the present experiment to detect within-subject differences in proportion correct (high-AVC > low-AVC), across several assumed effect sizes, using the Bayesian one-sided p-value (p_H). Like Experiment 1, binary hypothesis testing was not the primary goal of the statistical analysis, but the assessment of power across different potential effect sizes provides an estimate of the minimum effect that could be reliably detected for the selected sample size ($n = 10$). Table 3 shows the estimated power for differences in proportion correct ranging from 0.05-0.15 in steps of 0.025. Power was greater than 0.9 for effects of 0.075 and larger.

Table 3. Power to detect differences in proportion correct (PC) using the Bayesian one-sided p-value

		Effect Size (PC)				
		0.050	0.075	0.100	0.125	0.150
Power		0.68	0.93	> 0.99	> 0.99	> 0.99

Audiovisual Intelligibility Model

The audiovisual intelligibility index, AV_{Index} , was computed just as for Experiment 1 except that, for each SNR processed through the model, AV_{Index} was calculated separately for the 40 high-AVC items and the 40 low-AVC items, as opposed to sampling 40 items at random from the overall set of 250 audiovisual HINT sentences. As for Experiment 1, AV_{Index} was calculated across a range of SNRs so that predicted psychometric functions could be visualized. Additionally, even though no auditory-only experimental data were collected in Experiment 2, auditory-only intelligibility was estimated within the model.

Results

Bayesian Statistical Model

Population-level (fixed-effects) posterior distributions of proportion correct in the high-AVC and low-AVC conditions are plotted in Figure 9. Performance was slightly better in the high-AVC than the low-AVC condition. The posterior distribution of the contrast between conditions (high-AVC > low-AVC) had a modal value of 0.046 (HDI: [-0.002, 0.088], $p_H = 0.03$). The Bayes factor for the fixed effect of condition was 1.60, which Kass and Raftery (1995) suggest is “not worth more than a bare mention” in terms of evidence in favor of the full model.

Figure 9. Experiment 2 Bayesian statistical model: Proportion correct. Population-level posterior distributions of proportion correct at a fixed signal-to-noise ratio of -10.5 dB are visualized as violin plots (see Figure 5) for high-AVC (blue) and low-AVC (black) items.

Audiovisual Intelligibility Model

Psychometric functions predicted by the audiovisual speech intelligibility index (AV_{Index}) for high-AVC and low-AVC items in a background of SSN are plotted in Figure 10A, along with difference functions showing audiovisual minus auditory-only (Figure 10B), and the maximum auditory and visual predictive weights across SNRs (Figure 10C). Not surprisingly, just as for Experiment 1, the model predicts a large audiovisual gain in proportion correct even at very low SNRs. This gain is larger for high-AVC than low-AVC items. Again, the difference functions (audiovisual minus auditory-only) decrease monotonically with SNR. The visual predictive weights are higher for high-AVC than low-AVC items at SNRs lower than approximately -4 dB. At higher SNRs, equal visual-predictive weight is placed on high-AVC and low-AVC items. The crossover points of the visual and auditory predictive weight functions for high-AVC and low-AVC items occur at appx. -10 dB and -14 dB, respectively. Figure 11A shows summary metrics indicating the predicted magnitude of the audiovisual advantage in dB at the SRT for

high-AVC and low-AVC items, and Figure 11B shows the predicted proportion correct for audiovisual speech in the high-AVC and low-AVC conditions at the same SNR used in experimental testing (-10.5 dB). The predicted audiovisual advantage is ~ 2.5 dB larger for high-AVC than low-AVC items, and the predicted proportion correct at -10.5 dB SNR is ~ 0.10 larger for high-AVC than low-AVC items. Thus, the predicted audiovisual gain is approximately twice that observed in the experimental data. However, the model predictions are fairly accurate in comparison to the experimental data in terms of the overall proportion correct at -10.5 dB SNR (~ 0.5).

Figure 10. Experiment 2 audiovisual speech intelligibility model: Psychometric functions. (A) Model psychometric functions. Signal-to-noise ratio (SNR) is plotted along the abscissa and the scaled intelligibility index, AV_{Index} , is plotted along the ordinate. (B) Model estimates of the proportion-correct gain for audiovisual speech vs. auditory speech. Signal-to-noise ratio (SNR) is plotted along the abscissa and the scaled intelligibility index, AV_{Index} , is plotted along the ordinate. (C) Model weights by SNR. Signal-to-noise ratio (SNR) is plotted along the abscissa and the maximum model weight is plotted along the ordinate. A = auditory, AV = audiovisual, V = visual, high-AVC = items with high audiovisual coherence, low-AVC = items with low audiovisual coherence.

Figure 11. Experiment 2 audiovisual speech intelligibility model: Scalar metrics. (A) Audiovisual advantage in terms of signal-to-noise ratio shift (dB) at the speech reception threshold (SRT). (B) Proportion correct at the tested signal-to-noise ratio of -10.5 dB. (C) Guess rate parameter. HighAVC = items with high audiovisual coherence, lowAVC = items with low audiovisual coherence.

Post-hoc Cross-correlation Analysis

One possible confound in the comparison of high-AVC to low-AVC sentences is that the duration of a sentence – defined here as the number of frames from the onset to the offset of visible lip movements – can influence the magnitude of AVC. This is clear from Figure 8, which shows that high-AVC items tend to have a shorter duration. Indeed, the computation of magnitude squared coherence relies on estimation of (cross-)spectral density using Welch’s method of averaging over short, windowed segments of the input signals (Welch, 1967). When the total number of samples is low, as is the case here, the number of averages can vary considerably with the duration of the signal, and fewer averages will bias the estimate of

magnitude squared coherence to be larger. It is possible that such bias is present within the subsets of high-AVC and low-AVC items selected for this experiment. To address this, we computed a revised measure of audiovisual temporal synchrony using the cross-correlation, which does not necessitate windowing and therefore is not subject to the same duration bias as the magnitude squared coherence. It should be noted that the cross-correlation is not completely free of duration bias because spuriously high or low correlations are more likely to occur when the number of samples is very low, but this was of little concern given the relatively large number of samples in the full-duration (i.e., un-windowed) AVHINT sentences.

Computation of the cross-correlation measure proceeded just as computation of AVC – i.e., gammatone filtering, envelope extraction, and downsampling of the auditory signal (see Figure 3A) – but rather than computing the magnitude squared coherence of the band-specific envelope with the interlip distance function, the cross-correlation of these signals at lags of ± 5 frames was instead computed. This range of lags was chosen to capture most of the rather broad audiovisual window of temporal integration for speech (Grant & Seitz, 2000; van Wassenhove et al., 2007). Note, negative lags here correspond to the visual speech signal leading the auditory speech signal in time, which commonly yields a slightly larger correlation between the signals (Grant & Seitz, 2000). Also, unlike computation of AVC, computation of the cross-correlation included only frames for which some energy was present in the broadband auditory speech signal.

In the end, a 2D representation (128 gammatone bands x 11 time lags) describing the temporal correlation (in Pearson r units) between the band-specific acoustic speech envelope and the interlip distance function at different temporal offsets was produced. This 2D representation was calculated separately for the 40-item subsets of high-AVC and low-AVC items (Figure 12).

The following can be concluded from visual inspection of the mean “audiovisual cross-correlograms” of high-AVC and low-AVC items: (1) the largest correlations were observed at lags of -2 to 0 frames for both high-AVC and low-AVC items; (2) the largest correlations were observed at gammatone center frequencies of ~ 700 -5000 Hz (i.e., the same range of frequencies where AVC was maximal); (3) in the region of maximal correlations, the average correlation of high-AVC items was about 0.25 higher than the average correlation of low-AVC items; and (4) nonetheless, the division between high-AVC and low-AVC items was less clear in the cross-correlation than it was in the AVC metric itself (compare Figure 12 to Figure 8A).

Figure 12. Audiovisual cross-correlation (AVXc) of high-AVC and low-AVC sentence subsets. The mean AVXc of high-AVC (left) and low-AVC (middle) items is plotted, along with the difference in the means (right). Cross-correlation lag is plotted along the abscissa (negative = visual leading) and gammatone band center frequency (Hz) along the ordinate. The color scale indicates the magnitude of AVXc in Pearson r units.

Therefore, a single scalar metric summarizing the audiovisual cross-correlation of each high-AVC and low-AVC item was extracted by averaging the values of the cross-correlogram over the region of maximal r values (lag of -1, gammatone bands > 700 Hz and < 4750 Hz). These audiovisual cross-correlation (AVXc) values were then sorted from low to high, and a

median split was performed to create new subsets of high-AVXc and low-AVXc items. The distributions of high-AVXc and low-AVXc items (Figure 13A) were largely non-overlapping, but again were less well separated than the completely non-overlapping distributions of high-AVC and low-AVC items (not pictured, but see Figure 8A). However, 82.5% of the items classified as high-AVC and low-AVC remained in the “high” and “low” groups, respectively, when defined in terms of AVXc. Nevertheless, we computed a new Bayesian statistical model comparing performance on high-AVXc vs. low-AVXc items, which was otherwise identical to the original model comparing performance on high-AVC vs. low-AVC items. Population-level (fixed-effects) posterior distributions of proportion correct in the high-AVXc and low-AVXc conditions are plotted in Figure 13B. Performance was essentially no different between conditions. The posterior distribution of the contrast between conditions (high-AVXc > low-AVXc) had a modal value of 0.004 (HDI: [-0.037, 0.044], $p_H = 0.44$). The Bayes factor for the fixed effect of condition was 0.21, which in fact corresponds to a Bayes factor of 4.8 in favor of the null hypothesis (i.e., modest evidence in support of the conclusion that performance did not differ between high-AVXc and low-AVXc items).

Figure 13. Experiment 2 post-hoc cross-correlation analysis. (A) Newly-defined distributions of items with high (blue) and low (red) audiovisual cross-correlation (AVXc) based on a median split. (B) Bayesian statistical model of performance for high-AVXc and low-AVXc sentence subsets. Population-level posterior distributions of proportion correct at a fixed signal-to-noise ratio of -10.5 dB are visualized as violin plots (see Figure 5).

Discussion

Experiment 2 measured the effect of naturally occurring, item-wise differences in audiovisual temporal coherence (high-AVC vs. low-AVC) on audiovisual speech recognition in steady-state SSN. Experiment 1 had shown, at best, small effects of an artificial manipulation designed to reduce audiovisual coherence (temporal modulation filtering), leading to concerns

that the manipulation was either not completely effective or that reductions in overall intelligibility caused by modulation filtering obscured any effects of filtering on the audiovisual speech recognition advantage. The motivation behind Experiment 2 was to determine whether naturally occurring differences in audiovisual coherence would produce larger, more reliable, or more-easily interpretable changes in the audiovisual advantage. The hypothesis was that audiovisual speech recognition in noise – operationalized here as the proportion of keywords correctly identified at -10.5 dB SNR – would be better for high-AVC than low-AVC sentences.

In fact, Experiment 2 largely mirrored Experiment 1 in that only weak evidence was found to support an effect of audiovisual coherence on speech recognition performance. On the one hand, the data supported a directional effect of audiovisual coherence (high-AVC > low-AVC) based on the Bayesian one-sided p-value ($p_H < 0.05$). On the other hand, the posterior mode of the effect was only 0.046, and the Bayes factor based on model comparison (i.e., of models with and without a fixed effect of audiovisual coherence) was only 1.60, indicating only a trivial difference between the models. Moreover, the proportion-correct difference between high- and low-AVC items predicted by the AV_{Index} (0.10) was twice the difference observed in the experimental data. This is worth noting because the AV_{Index} is based entirely on the temporal relationship between the interlip distance function (i.e., dynamic visual speech cues) and the auditory speech envelope. As in Experiment 1, the AV_{Index} tended to over-predict the contribution of dynamic visual speech cues at very low SNRs, and this effect was more pronounced for high- vs. low-AVC items (Figure 10A). This over-prediction at low SNRs could be responsible for the difference between the model predictions and the experimental data at the tested SNR of -10.5 dB.

A possible confound was noted in which the magnitude of AVC could be spuriously inflated for short duration sentences. To address this, audiovisual temporal synchrony was redefined in terms of the cross-correlation (AVXc), the magnitude of which should not be driven by duration. In fact, the correlation between duration and AVXc was -0.13 across all of the sentences tested in this experiment. When new sets of high-AVXc and low-AVXc items were constructed (median split), 82.5% of high(low)-AVC items remained in the newly-defined high(low)-AVXc sets, which suggests that the duration confound did not produce large distortions in AVC. However, the separation between the distributions of high-AVXc and low-AVXc items (Figure 12A) was not as large as the separation between the distributions of high-AVC and low-AVC items (Figure 8A). This may reflect some contribution of duration to the estimates of AVC. Therefore, performance was re-assessed in terms by comparing the mean proportion correct for high-AVXc vs. low-AVXc items, but no difference in performance was detected (posterior mode = 0.004, Bayes factor = 4.8 in favor of the null hypothesis). This result should be interpreted with caution because the set of 80 hint sentences used in this experiment was selected to be maximally different in terms of AVC, not AVXc. Indeed, the modes of the high-AVXc and low-AVXc distributions were separated by only ~ 0.1 (Pearson r units; Figure 13A). On the other hand, the means of the high-AVXc and low-AVXc distributions were separated by 0.28, which is comparable to the difference in AVXc that has been observed to modulate the magnitude of the audiovisual speech detection advantage (Grant & Seitz, 2000). At the very least, we can conclude that larger and/or more consistent differences in AVXc would be required to produce an effect on the audiovisual speech recognition advantage.

One potential criticism of Experiment 2 is that we did not report on differences in auditory-only performance for high/low-AVC or high/low-AVXc items. Therefore, any

differences in audiovisual speech recognition performance reflect auditory-only plus audiovisual differences. However, since testing was carried out at a rather unfavorable SNR (-10.5 dB) chosen to maximize the influence of the visual speech information on performance (Holmes, 2007, 2009; Ross et al., 2007; Stein, Stanford, Ramachandran, Perrault Jr, & Rowland, 2009; Stevenson & James, 2009) and thus the ability to detect performance differences based on AVC/AVXc, auditory-only speech recognition performance at the tested SNR can be assumed to be quite poor (proportion correct ~ 0.2 as inferred from Experiment 1). Item-wise differences in auditory-only performance would likely be unreliable when tested so close to the performance floor. One option would have been to assess auditory-only performance at a different SNR that matches the overall proportion correct for audiovisual trials (0.5), but this assumes that item-wise differences do not interact with SNR. Moreover, we can reason that auditory-only differences in performance were not likely to have influenced the audiovisual results because: (1) the subsets of high/low-AVC/AVXc items were large enough ($n = 40$) that any idiosyncratic differences in item-wise performance should “average out” between the groups; (2) the audiovisual intelligibility model, AV_{Index} , did not predict any differences in auditory-only performance for high-AVC vs. low-AVC items (Figure 10); and (3) there is no *a priori* reason to expect that auditory-only intelligibility should be systematically related to AVC/AVXc. Regarding the latter point, if such a systematic relation between auditory-only intelligibility and AVC/AVXc were present and, moreover, responsible for the small and/or negligible effect of coherence/correlation (e.g., if auditory-only intelligibility was systematically lower for high coherence/correlation items), we would conclude that an increase in AVC/AVXc confers no practical benefit to audiovisual speech recognition performance because high AVC/AVXc would be offset by low auditory-only intelligibility.

General Discussion

In this study, two experiments were performed to investigate the influence of multisensory temporal covariation on audiovisual speech recognition in steady-state, speech-shaped noise. Here, speech recognition was assessed using AVHINT sentences with keyword scoring. Experiment 1 used modulation filtering to remove 3-7 Hz modulations from the envelope of the acoustic speech signal, which significantly reduced the audiovisual temporal coherence of the multisensory speech signal (Figure 3B). This manipulation was expected to reduce the magnitude of the audiovisual advantage for speech recognition in noise. However, modulation filtering produced no effect on the magnitude of the audiovisual speech recognition advantage measured in terms of an SNR shift (dB) at the SRT (Figure 5A), and produced, if anything, only a small decrease in the maximum proportion-correct gain for audiovisual versus auditory-only speech recognition (Figure 5B). Removal of static visual cues – operationalized in Experiment 1 as Gaussian facial blurring with a cutoff of 5.5 cycles/face – significantly impaired lipreading performance (Figure 5C) but did not appreciably reduce the magnitude of the audiovisual advantage, nor did it interact with the effect of modulation filtering on the audiovisual advantage. Experiment 2 measured differences in audiovisual speech recognition performance across naturally occurring subsets of AVHINT sentences with high or low audiovisual temporal coherence. This was intended to provide interpretational clarity relative to the modulation filtering manipulation of Experiment 1, which, by introducing acoustic distortions, may have produced several (potentially opposing) effects on the audiovisual speech recognition advantage. In Experiment 2, naturalistic differences in audiovisual temporal coherence – operationalized in terms of either magnitude squared coherence (AVC; Figure 8) or cross-correlation (AVXc; Figure 12) – did not produce appreciable differences in audiovisual

speech recognition performance (Figures 9 and 13). An audiovisual speech intelligibility model based entirely on temporal covariation between the visual and auditory speech signals predicted larger effects of modulation filtering (Experiment 1; Figure 7) and high vs. low temporal coherence (Experiment 2; Figure 11) than those observed in the empirical data.

Together, these results challenge the widely held assumption that dynamic (time-varying) visual speech cues influence audiovisual speech recognition primarily via their cross-modal redundancy (or correlation) with the time-varying auditory speech signal (R. Campbell, 2008; Chandrasekaran et al., 2009; Grant et al., 2013; Rosenblum et al., 2005; A. Summerfield, Campbell, Dodd, & Burnham, 1987). In fact, seminal work by Summerfield (1979) showed that a visual stimulus consisting of a circle whose diameter was varied in correspondence to the acoustic speech envelope provided no benefit to speech recognition (sentences in a single-talker background) at an SNR of -12 dB, while videos of the target talker's face provided an increase of ~ 40% in recognition performance. The null effect of the modulated circle stimulus was recently replicated using modern presentation equipment to add luminance cues for changes in the speech envelope and to ensure precise audiovisual temporal alignment (Strand, Brown, & Barbour, 2019). This demonstrates that correlated temporal information alone, when delivered through the visual modality as a non-speech signal, is not sufficient to improve speech recognition in noise. One potential criticism of Summerfield's finding, including the recent replication, is that the acoustic speech envelope was derived from the broadband acoustic signal, whereas the 'internal' representation of the speech envelope can be heavily modified by narrowband cochlear filtering (Ghitza, 2001). In the present study, audiovisual temporal correlation (AVC, AVXc) was computed from acoustic speech envelopes within narrow (gammatone) frequency bands (Figures 3A, 8, 12). Nonetheless, in line with Summerfield's work, we were unable to find a

systematic effect of narrowband audiovisual temporal correlation on speech recognition in steady-state noise.

The present results should not, however, be interpreted to suggest that temporal covariation between visual and auditory speech signals plays absolutely no role in speech recognition in noise. This can be demonstrated trivially by shifting the auditory and visual speech signals to be completely out of temporal alignment (i.e., asynchronous by ~ 400 ms or more), at which point no benefit will be conferred by the presence of the visual speech signal (Grant & Greenberg, 2001; van Wassenhove et al., 2007). Clearly some temporal correspondence must be present between visual and auditory speech signals in order to derive an audiovisual benefit, and, as reviewed in the Introduction above, this cannot be attributed entirely to temporally coincident occurrences of static visual speech cues (Jordan & Sergeant, 2000; MacDonald et al., 2000; K. Munhall et al., 2004; Rosenblum & Saldaña, 1996; Thomas & Jordan, 2002). Therefore, the present results should serve to “rule out” narrowband temporal covariation between the interlip distance function and the acoustic speech envelope as a key factor in the audiovisual speech recognition advantage, and to suggest several new directions of research that follow naturally from this conclusion.

One direction to pursue from here is to examine different ways of characterizing multisensory temporal covariation. First, it is possible that the degree of local (short-term) temporal covariation is more important than the degree of global (full-utterance) temporal covariation. Here, we have derived scalar metrics based on AVC and AVXc that are integrated across many frequency bands and the entire duration of each AVHINT utterance. However, just as models of speech intelligibility can benefit from integrating the short-term SNR of noisy speech across multiple observation intervals (Bronkhorst, Bosman, & Smoorenburg, 1993;

Rhebergen & Versfeld, 2005; Rhebergen, Versfeld, & Dreschler, 2006; Shen, Manzano, et al., 2015), measures of audiovisual temporal covariation could benefit from combining multiple observations of AVC or AVXc across short-duration time windows. It is certainly possible for an utterance to have a high local AVC/AVXc at crucial moments in the utterance, yet low or moderate AVC/AVXc overall. Grant and Seitz (2000) show that such local audiovisual correlations (333 ms time may be particularly important to the audiovisual speech detection advantage. Local correlations could potentially be worked into current models of audiovisual speech intelligibility (Grant et al., 2013; Grant & Bernstein, 2019), including the AV_{Index} presented here. The AV_{Index} provided remarkably good predictions of overall audiovisual-speech intelligibility given its simplicity, suggesting that tweaks designed to address the model's faults as noted in Experiments 1 and 2 (e.g., over-prediction of intelligibility at very low SNRs) could yield a very accurate and informative model. Recent machine learning models designed to perform audiovisual speech enhancement suggest that it is possible to derive local, frame-by-frame estimates of acoustic speech features directly from short segments of concurrent visual speech (Adeel, Gogate, & Hussain, 2018; Adeel, Gogate, Hussain, & Whitmer, 2018).

Second, it is possible that audiovisual temporal covariation should be quantified only after a transformation of the auditory (e.g., speech envelope) and visual (e.g., interlip distance) speech time-courses has been applied. One sensible candidate for such a transformation is the first temporal derivative, which is commonly used to augment spectral acoustic features in automatic speech recognition algorithms (Hunt, 1999). Recent data show that neuronal populations in the human auditory do not encode moment-by-moment changes in the amplitude envelope of speech, but rather detect local maxima (i.e., onset edges) in the first derivative of the envelope (Oganian & Chang, 2018). These local maxima often occur during consonant-vowel

transitions and, indeed, some of our own work shows that visual speech contributes maximally to audiovisual fusion (McGurk & MacDonald, 1976) during peaks in the first derivative of the interlip distance function (i.e., lip velocity) occurring at a consonant-vowel transition (Jonathan H Venezia, Thurman, Matchin, George, & Hickok, 2015). The temporal relation between visual and auditory “edges” could be particularly important to the audiovisual speech recognition advantage, and could be related to audiovisual “pre-cueing” effects observed to be crucial to the audiovisual speech detection advantage (L. E. Bernstein et al., 2004; Eskelund et al., 2011; Kim & Davis, 2004). Thus, a lagged measure of audiovisual temporal covariation such as AVXc is strongly suggested for quantifying covariation between the first derivatives of dynamic auditory and visual speech signals.

Another potential direction is to consider the role of audiovisual temporal covariation on speech recognition in modulated backgrounds (e.g., competing speech). One possible explanation for the small or negligible effects of AVC/AVXc observed here is that contributions of dynamic visual speech cues may have been minimized by the use of steady-state noise as the background masker. This would be the case if dynamic visual speech cues primarily benefit audiovisual speech recognition by providing release from modulation masking and/or informational masking as in a background of competing talkers (Francart, van Wieringen, & Wouters, 2011; Shen, Manzano, et al., 2015). Support for this possibility is provided by recent neurophysiological and behavioral studies (J. G. W. Bernstein & Grant, 2009; Golumbic et al., 2013; Park et al., 2018). Moreover, effects of local audiovisual temporal correlation, as discussed above, would likely be maximized in a modulated background, suggesting that these lines of investigation could or should be combined.

Finally, a direction that should certainly be pursued by future studies investigating the role of multisensory temporal covariation on audiovisual speech recognition is to consider the trajectories of facial effectors other than the lips. Here, we have characterized the dynamic visual speech signal entirely in terms of the interlip distance function. While this approach was justified by studies showing a clear relation between the lip area function or interlip distance function and the narrowband acoustic speech envelope (Chandrasekaran et al., 2009; Grant & Seitz, 2000), and by behavioral studies showing that the lips of particular importance to the perception of visual speech (Gurler, Doyle, Walker, Magnotti, & Beauchamp, 2015; Lander & Capek, 2013; McGrath, 1985), modeling studies have shown that the best predictions of the acoustic speech signal from the visual speech signal are derived by considering multiple facial effectors (e.g., lips, jaw, tongue, and cheeks; Grant et al., 2013; Grant & Bernstein, 2019; H. C. Yehia et al., 2002). It is possible that considering other facial effectors in the computation of metrics such as AVC and AVXc would enhance the ability to detect the effect of these variables on audiovisual speech recognition, particularly for large sets of utterances that may contain idiosyncratic contributions from the different facial effectors across items. In the context of Bayesian analysis as implemented here, this could translate to larger and/or more precise estimates of the effect.

To conclude, we have shown in two experiments based on exogenous manipulation of audiovisual speech stimuli and naturalistic variation within the same set of audiovisual speech stimuli that the magnitude of multisensory temporal covariation (AVC, AVXc) has little to no effect on the audiovisual speech recognition advantage. We have also shown that a simple model of audiovisual speech intelligibility based entirely on multisensory temporal covariation tends to over-predict the contribution of such covariation to audiovisual speech recognition

performance, yet the model provides rather good predictions of performance overall. These findings call into the assumption that multisensory temporal covariation plays a significant role in audiovisual speech recognition, but also suggest a number ways to refine and extend the set of experimental hypotheses that follow from this assumption, including consideration of local vs. global temporal covariation, temporal derivatives of dynamic auditory and visual speech signals, speech recognition in modulated backgrounds, and contributions from facial effectors other than the lips.

References

- Adeel, A., Gogate, M., & Hussain, A. (2018). Contextual Audio-Visual Switching For Speech Enhancement in Real-World Environments. *arXiv preprint arXiv:1808.09825*.
- Adeel, A., Gogate, M., Hussain, A., & Whitmer, W. M. (2018). Lip-reading driven deep learning approach for speech enhancement. *arXiv preprint arXiv:1808.00046*.
- Arnal, L. H., Morillon, B., Kell, C. A., & Giraud, A. L. (2009). Dual neural routing of visual facilitation in speech processing. *J Neurosci*, 29(43), 13445-13453.
doi:10.1523/JNEUROSCI.3194-09.2009
- Beauchamp, M. S., Nath, A. R., & Pasalar, S. (2010). fMRI-guided transcranial magnetic stimulation reveals that the superior temporal sulcus is a cortical locus of the McGurk effect. *The Journal of Neuroscience*, 30(7), 2414-2417.
- Bench, J., Kowal, Å., & Bamford, J. (1979). The Bkb (Bamford-Kowal-Bench) Sentence Lists for Partially-Hearing Children. *British Journal of Audiology*, 13(3), 108-112.
doi:10.3109/03005367909078884
- Bernstein, J. G. W., & Grant, K. W. (2009). Auditory and auditory-visual intelligibility of speech in fluctuating maskers for normal-hearing and hearing-impaired listeners. *The Journal of the Acoustical Society of America*, 125(5), 3358-3372.
- Bernstein, L. E., Auer, E. T., & Takayanagi, S. (2004). Auditory speech detection in noise enhanced by lipreading. *Speech Communication*, 44(1), 5-18.
- Bronkhorst, A. W., Bosman, A. J., & Smoorenburg, G. F. (1993). A model for context effects in speech recognition. *The Journal of the Acoustical Society of America*, 93(1), 499-509.

- Brooke, N. M., & Summerfield, Q. (1983). Analysis, synthesis, and perception of visible articulatory movements. *Journal of Phonetics*.
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1), 1-28.
- Calvert, G., & Campbell, R. (2003). Reading speech from still and moving faces: the neural substrates of visible speech. *Cognitive Neuroscience, Journal of*, 15(1), 57-70.
- Calvert, G. A., Campbell, R., & Brammer, M. J. (2000). Evidence from functional magnetic resonance imaging of crossmodal binding in the human heteromodal cortex. *Current Biology*, 10(11), 649-658.
- Campbell, R. (1986). The lateralization of lip-read sounds: A first look. *Brain Cogn*, 5(1), 1-21.
- Campbell, R. (2008). The processing of audio-visual speech: empirical and neural bases. *Philos Trans R Soc Lond B Biol Sci*, 363(1493), 1001-1010. doi:10.1098/rstb.2007.2155
- Campbell, R., Landis, T., & Regard, M. (1986). Face recognition and lipreading: A neurological dissociation. *Brain*, 109(3), 509-521.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., . . . Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1).
- Chandrasekaran, C., Lemus, L., & Ghazanfar, A. A. (2013). Dynamic faces speed up the onset of auditory cortical spiking responses during vocal detection. *Proceedings of the National Academy of Sciences*, 110(48), E4668-E4677.
- Chandrasekaran, C., Trubanova, A., Stillittano, S., Caplier, A., & Ghazanfar, A. A. (2009). The natural statistics of audiovisual speech. *PLoS computational biology*, 5(7), e1000436.

- Chen, T., & Rao, R. R. (1998). Audio-visual integration in multimodal communication. *Proceedings of the IEEE*, 86(5), 837-852.
- Crosse, M. J., Di Liberto, G. M., & Lalor, E. C. (2016). Eye can hear clearly now: inverse effectiveness in natural audiovisual speech processing relies on long-term crossmodal temporal integration. *Journal of Neuroscience*, 36(38), 9888-9895.
- Elliott, T. M., & Theunissen, F. E. (2009). The modulation transfer function for speech intelligibility. *PLoS computational biology*, 5(3), e1000302.
- Eskelund, K., Tuomainen, J., & Andersen, T. S. (2011). Multistage audiovisual integration of speech: Dissociating identification and detection. *Experimental Brain Research*, 208(3), 447-457.
- Fisher, C. G. (1968). Confusions among visually perceived consonants. *Journal of Speech, Language, and Hearing Research*, 11(4), 796-804.
- Francart, T., van Wieringen, A., & Wouters, J. (2011). Comparison of fluctuating maskers for speech recognition tests. *International journal of audiology*, 50(1), 2-13.
- Ghitza, O. (2001). On the upper cutoff frequency of the auditory critical-band envelope detectors in the context of speech perception. *The Journal of the Acoustical Society of America*, 110(3), 1628-1640.
- Girges, C., Spencer, J., & O'Brien, J. (2015). Categorizing identity from facial motion. *The Quarterly Journal of Experimental Psychology*, 68(9), 1832-1843.
- Golumbic, E. Z., Cogan, G. B., Schroeder, C. E., & Poeppel, D. (2013). Visual input enhances selective speech envelope tracking in auditory cortex at a “cocktail party”. *The Journal of Neuroscience*, 33(4), 1417-1426.

- Grant, K. W., Bernstein, J. G., & Summers, V. (2013). Conclusion: predicting speech intelligibility by individual hearing-impaired listeners: the path forward. *Journal of the American Academy of Audiology*, 24(4), 329.
- Grant, K. W., & Bernstein, J. G. W. (2019). Toward a Model of Auditory-Visual Speech Intelligibility *Multisensory Processes* (pp. 33-57): Springer.
- Grant, K. W., & Greenberg, S. (2001). *Speech intelligibility derived from asynchronous processing of auditory-visual information*. Paper presented at the AVSP 2001-International Conference on Auditory-Visual Speech Processing.
- Grant, K. W., & Seitz, P.-F. (2000). The use of visible speech cues for improving auditory detection of spoken sentences. *The Journal of the Acoustical Society of America*, 108(3), 1197-1208.
- Gronau, Q. F., Sarafoglou, A., Matzke, D., Ly, A., Boehm, U., Marsman, M., . . . Steingroever, H. (2017). A tutorial on bridge sampling. *Journal of mathematical psychology*, 81, 80-97.
- Gurler, D., Doyle, N., Walker, E., Magnotti, J., & Beauchamp, M. (2015). A link between individual differences in multisensory speech perception and eye movements. *Attention, Perception, & Psychophysics*, 77(4), 1333-1341.
- Hoffman, M. D., & Gelman, A. (2014). The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1), 1593-1623.
- Holmes, N. P. (2007). The law of inverse effectiveness in neurons and behaviour: multisensory integration versus normal variability. *Neuropsychologia*, 45(14), 3340-3345.
- Holmes, N. P. (2009). The principle of inverse effectiveness in multisensory integration: some statistical considerations. *Brain topography*, 21(3-4), 168-176.

- Houpt, J. W., & Bittner, J. L. (2018). Analyzing thresholds and efficiency with hierarchical Bayesian logistic regression. *Vision research*, 148, 49-58.
- Hunt, M. J. (1999). Spectral signal processing for ASR. *Proc. ASRU'99*.
- Jesse, A., & Bartoli, M. (2018). Learning to recognize unfamiliar talkers: Listeners rapidly form representations of facial dynamic signatures. *Cognition*, 176, 195-208.
- Jordan, T. R., & Sergeant, P. (2000). Effects of distance on visual and audiovisual speech recognition. *Language and Speech*, 43(1), 107-124.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the american statistical association*, 90(430), 773-795.
- Kim, J., & Davis, C. (2004). Investigating the audio-visual speech detection advantage. *Speech Communication*, 44(1), 19-30.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, 36(14), 1.1-16.
- Kruschke, J. K., & Liddell, T. M. (2018). The Bayesian New Statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychonomic Bulletin & Review*, 25(1), 178-206.
- Lander, K., & Capek, C. (2013). Investigating the impact of lip visibility and talking style on speechreading performance. *Speech Communication*, 55(5), 600-605.
- MacDonald, J., Andersen, S., & Bachmann, T. (2000). Hearing by eye: how much spatial degradation can be tolerated? *Perception*, 29(10), 1155-1168.
- MacLeod, A., & Summerfield, Q. (1987). Quantifying the contribution of vision to speech perception in noise. *British Journal of Audiology*, 21(2), 131-141.

- MacLeod, A., & Summerfield, Q. (1990). A procedure for measuring auditory and audiovisual speech-reception thresholds for sentences in noise: Rationale, evaluation, and recommendations for use. *British Journal of Audiology*, 24(1), 29-43.
- Marsman, M., & Wagenmakers, E.-J. (2017). Three insights from a Bayesian interpretation of the one-sided P value. *Educational and Psychological Measurement*, 77(3), 529-539.
- Massaro, D. W., & Simpson, J. A. (2014). *Speech perception by ear and eye: A paradigm for psychological inquiry*: Psychology Press.
- McGrath, M. (1985). *An examination of cues for visual and audio-visual speech perception using natural and computer-generated faces*. University of Nottingham.
- McGurk, H., & MacDonald, J. (1976). Hearing lips and seeing voices.
- Middelweerd, M. J., & Plomp, R. (1987). The effect of speechreading on the speech-reception threshold of sentences in noise. *J Acoust Soc Am*, 82(6), 2145-2147.
- Moore, B. C. J., & Glasberg, B. R. (1983). Suggested formulae for calculating auditory-filter bandwidths and excitation patterns. *The Journal of the Acoustical Society of America*, 74(3), 750-753.
- Moscatelli, A., Mezzetti, M., & Lacquaniti, F. (2012). Modeling psychophysical data at the population-level: the generalized linear mixed model. *Journal of Vision*, 12(11), 26-26.
- Munhall, K., Kroos, C., Jozan, G., & Vatikiotis-Bateson, E. (2004). Spatial frequency requirements for audiovisual speech perception. *Perception & Psychophysics*, 66(4), 574-583.
- Munhall, K. G., Jones, J. A., Callan, D. E., Kuratate, T., & Vatikiotis-Bateson, E. (2004). Visual prosody and speech intelligibility head movement improves auditory speech perception. *Psychological Science*, 15(2), 133-137.

- Nilsson, M., Soli, S. D., & Sullivan, J. A. (1994). Development of the Hearing in Noise Test for the measurement of speech reception thresholds in quiet and in noise. *The Journal of the Acoustical Society of America*, 95(2), 1085-1099.
- Oganian, Y., & Chang, E. F. (2018). A speech envelope landmark for syllable encoding in human superior temporal gyrus. *bioRxiv*, 388280.
- Paré, M., Richler, R. C., ten Hove, M., & Munhall, K. G. (2003). Gaze behavior in audiovisual speech perception: The influence of ocular fixations on the McGurk effect. *Perception & Psychophysics*, 65(4), 553-567.
- Park, H., Ince, R. A. A., Schyns, P. G., Thut, G., & Gross, J. (2018). Representational interactions during audiovisual speech entrainment: Redundancy in left posterior superior temporal gyrus and synergy in left motor cortex. *PLoS biology*, 16(8), e2006558.
- Rhebergen, K. S., & Versfeld, N. J. (2005). A speech intelligibility index-based approach to predict the speech reception threshold for sentences in fluctuating noise for normal-hearing listeners. *The Journal of the Acoustical Society of America*, 117(4), 2181-2192.
- Rhebergen, K. S., Versfeld, N. J., & Dreschler, W. A. (2006). Extended speech intelligibility index for the prediction of the speech reception threshold in fluctuating noise. *The Journal of the Acoustical Society of America*, 120(6), 3988-3997.
- Rosenblum, L. D., Johnson, J. A., & Saldana, H. M. (1996). Point-light facial displays enhance comprehension of speech in noise. *Journal of Speech, Language, and Hearing Research*, 39(6), 1159-1170.
- Rosenblum, L. D., Niehus, R. P., & Smith, N. M. (2007). Look who's talking: recognizing friends from visible articulation. *Perception*, 36(1), 157-159.

Rosenblum, L. D., Pisoni, D., & Remez, R. (2005). Primacy of multimodal speech perception.

Handbook of speech perception, 51-78.

Rosenblum, L. D., & Saldana, H. M. (1998). Time-varying information for visual speech

perception. *Hearing by eye II*, 61-81.

Rosenblum, L. D., & Saldaña, H. M. (1996). An audiovisual test of kinematic primitives for

visual speech perception. *Journal of Experimental Psychology: Human Perception and Performance*, 22(2), 318.

Rosenblum, L. D., Smith, N. M., Nichols, S. M., Hale, S., & Lee, J. (2006). Hearing a face:

Cross-modal speaker matching using isolated visible speech. *Perception & Psychophysics*, 68(1), 84-93.

Rosenblum, L. D., Yakel, D. A., Baseer, N., Panchal, A., Nodarse, B. C., & Niehus, R. P. (2002).

Visual speech information for face recognition. *Perception & Psychophysics*, 64(2), 220-229.

Ross, L. A., Saint-Amour, D., Leavitt, V. M., Javitt, D. C., & Foxe, J. J. (2007). Do you see what

I am saying? Exploring visual enhancement of speech comprehension in noisy environments. *Cerebral Cortex*, 17(5), 1147-1153.

Schroeder, C. E., Lakatos, P., Kajikawa, Y., Partan, S., & Puce, A. (2008). Neuronal oscillations

and visual amplification of speech. *Trends in Cognitive Sciences*, 12(3), 106-113.

Schwartz, J.-L., Basirat, A., Ménard, L., & Sato, M. (2012). The Perception-for-Action-Control

Theory (PACT): A perceptuo-motor theory of speech perception. *Journal of Neurolinguistics*, 25(5), 336-354.

Schwartz, J.-L., Berthommier, F., & Savariaux, C. (2004). Seeing to hear better: evidence for

early audio-visual interactions in speech identification. *Cognition*, 93(2), B69-B78.

- Schwartz, J.-L., & Savariaux, C. (2014). No, There Is No 150 ms Lead of Visual Speech on Auditory Speech, but a Range of Audiovisual Asynchronies Varying from Small Audio Lead to Large Audio Lag. *PLoS Comput Biol*, 10(7), e1003743. doi:citeulike-article-id:13320829
doi: 10.1371/journal.pcbi.1003743
- Shahin, A. J., Kerlin, J. R., Bhat, J., & Miller, L. M. (2012). Neural restoration of degraded audiovisual speech. *Neuroimage*, 60(1), 530-538.
- Shen, Y., Dai, W., & Richards, V. M. (2015). A MATLAB toolbox for the efficient estimation of the psychometric function using the updated maximum-likelihood adaptive procedure. *Behavior research methods*, 47(1), 13-26.
- Shen, Y., Manzano, N. K., & Richards, V. M. (2015). Psychometric functions for sentence recognition in sinusoidally amplitude-modulated noises. *The Journal of the Acoustical Society of America*, 138(6), 3613-3624.
- Stein, B. E., Stanford, T. R., Ramachandran, R., Perrault Jr, T. J., & Rowland, B. A. (2009). Challenges in quantifying multisensory integration: alternative criteria, models, and inverse effectiveness. *Experimental Brain Research*, 198(2-3), 113-126.
- Stekelenburg, J. J., & Vroomen, J. (2007). Neural correlates of multisensory integration of ecologically valid audiovisual events. *Journal of Cognitive Neuroscience*, 19(12), 1964-1973.
- Stevenson, R. A., & James, T. W. (2009). Audiovisual integration in human superior temporal sulcus: inverse effectiveness and the neural processing of speech and object recognition. *Neuroimage*, 44(3), 1210-1223.

- Strand, J. F., Brown, V. A., & Barbour, D. L. (2019). Talking points: A modulating circle reduces listening effort without improving speech recognition. *Psychonomic Bulletin & Review*, 26(1), 291-297.
- Summerfield, A., Campbell, R., Dodd, B., & Burnham, D. (1987). Some preliminaries to a theory of audiovisual speech processing. *Hearing by Eye II: The Psychology of Speechreading and Auditory-Visual Speech*, 58-82.
- Summerfield, Q. (1979). Use of visual information for phonetic perception. *Phonetica*, 36(4-5), 314-331.
- Summerfield, Q. (1987). Some preliminaries to a comprehensive account of audio-visual speech perception.
- Summerfield, Q. (1992). Lipreading and audio-visual speech perception. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 335(1273), 71-78.
- Team, R. C. (2013). R: A language and environment for statistical computing.
- Thomas, S. M., & Jordan, T. R. (2002). Determining the influence of Gaussian blurring on inversion effects with talking faces. *Perception & Psychophysics*, 64(6), 932-944.
- Tomasi, C., & Kanade, T. (1991). Detection and tracking of point features.
- van Wassenhove, V., Grant, K. W., & Poeppel, D. (2005). Visual speech speeds up the neural processing of auditory speech. *Proc Natl Acad Sci U S A*, 102(4), 1181-1186.
doi:10.1073/pnas.0408949102
- van Wassenhove, V., Grant, K. W., & Poeppel, D. (2007). Temporal window of integration in auditory-visual speech perception. *Neuropsychologia*, 45(3), 598-607.
doi:10.1016/j.neuropsychologia.2006.01.001

- Venezia, J. H., Thurman, S. M., Matchin, W., George, S. E., & Hickok, G. (2015). Timing in audiovisual speech perception: A mini review and new psychophysical data. *Attention, Perception, & Psychophysics*, 1-19.
- Venezia, J. H., Vaden Jr, K. I., Rong, F., Maddox, D., Saberi, K., & Hickok, G. (2017). Auditory, visual and audiovisual speech processing streams in superior temporal sulcus. *Front Hum Neurosci*, 11, 174.
- Viola, P., & Jones, M. J. (2004). Robust real-time face detection. *International journal of computer vision*, 57(2), 137-154.
- Welch, P. (1967). The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, 15(2), 70-73.
- Yehia, H., Rubin, P., & Vatikiotis-Bateson, E. (1998). Quantitative association of vocal-tract and facial behavior. *Speech Communication*, 26(1), 23-43.
- Yehia, H. C., Kuratate, T., & Vatikiotis-Bateson, E. (2002). Linking facial animation, head motion and speech acoustics. *Journal of Phonetics*, 30(3), 555-568.

Appendix 1

The UML is a Bayesian extension of the maximum likelihood (ML) adaptive procedure described by Green (1990), in which stimulus sampling occurs at a “sweet point” that minimizes the expected variance of the threshold parameter, a , for a given psychometric function. Later work showed that two additional sweet points, close to the lower and upper asymptotes of the psychometric function, are required to minimize the expected variance of the psychometric function slope parameter, b (Brand & Kollmeier, 2002; Kaernbach, 2001; King-Smith & Rose, 1997). The sweet points for a and b can be determined analytically for a given (assumed) shape of the psychometric function (e.g., cumulative Gaussian). The UML is a Bayesian minimum-variance procedure that rapidly estimates psychometric function parameters by: (a) setting reasonable priors on the parameter values; (b) generating a best-fit psychometric function (given the priors) from a single observation at a fixed signal value; (c) sampling a new observation at a signal value determined by the sweet points of the current best-fitting psychometric function; (d) updating the psychometric function fit with the new observation so as to minimize the variance of the parameter estimates; and (e) continuing until a convergence criterion is met. To prevent large changes in signal strength from trial to trial, the sweet point location (e.g., upper asymptote, mean, or lower asymptote) is determined using an adaptive staircase procedure (Levitt, 1971).

In the present implementation of the UML procedure, the psychometric function was characterized by a Gaussian cumulative distribution function as follows:

$$\psi = \gamma + (1 - \gamma - \lambda) * (1 + \operatorname{erf}\left(\frac{-SNR+a}{\sqrt{2*b^2}}\right))/2,$$

where ψ is the psychometric function describing percent correct keyword recognition as a function of SNR , γ is the ‘guess rate’ (essentially ‘lipreading performance’ in the context of

audiovisual speech), λ is the ‘lapse rate’, a is the psychometric function threshold, b is inversely proportional to the psychometric function slope, and erf is the error function. The lapse rate parameter, λ , was fixed at 0.02. The threshold and slope parameters, a and b , were estimated from the data by sampling from three sweet points at the center and upper/lower asymptotes of the psychometric function. The guess rate parameter, γ , was estimated from the data for the audiovisual speech conditions by sampling from a fourth sweet point at the lowest possible signal value (SNR = -25 dB). For the auditory-only speech conditions, γ was fixed at zero and priors were set on a and b as follows:

$$a \sim N(10, 20); -15 < a < 10$$

$$b \sim N(2, 5); 1 < b < 10.$$

For the audiovisual speech conditions, priors were set on a and b as follows:

$$a \sim N(10, 20); -15 < a < 10$$

$$b \sim N(2, 5); 1 < b < 10$$

$$\gamma \sim N(0.2, 0.2); 0 < \gamma < 0.5.$$

For all conditions, the SNR was set to an initial value of 20 dB at the start of each UML track and was limited to the range [-25, 20]. Approximately 20-30 sentences were presented in each UML track. Each sentence contained 2-6 keywords. After each sentence, the best-fitting psychometric function was updated once for each keyword in the sentence, all of which were presented at the same SNR. The SNR for the next sentence was then chosen by sampling the appropriate sweet point (determined by a simple up-down procedure) from the psychometric function with parameters determined by the final-keyword update from the preceding sentence. This process continued until 16 sweet-point reversals were obtained.

- Brand, T., & Kollmeier, B. (2002). Efficient adaptive procedures for threshold and concurrent slope estimates for psychophysics and speech intelligibility tests. *The Journal of the Acoustical Society of America*, 111(6), 2801-2810.
- Green, D. M. (1990). Stimulus selection in adaptive psychophysical procedures. *The Journal of the Acoustical Society of America*, 87(6), 2662-2674.
- Kaernbach, C. (2001). Slope bias of psychometric functions derived from adaptive data. *Perception & Psychophysics*, 63(8), 1389-1398.
- King-Smith, P. E., & Rose, D. (1997). Principles of an adaptive method for measuring the slope of the psychometric function. *Vision research*, 37(12), 1595-1604.
- Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *The Journal of the Acoustical Society of America*, 49(2B), 467-477.

Appendix 2

Formally, *brms* implements Bayesian hierarchical regression in terms of the following model (notation due to Bürkner, 2017):

$$y_i = D(f(\eta_i, \theta)),$$

where y_i is the i th response, η_i is a linear combination of the predictors, f is the inverse link function, D is a probability distribution characterizing the distribution of responses (i.e., the model “family”), and θ describes family-specific parameters that do not vary across responses (if any). The parameter η_i can be expanded to:

$$\eta_i = \mathbf{X}_i\beta + \mathbf{Z}_i\mu,$$

where $\mathbf{X}$ and $\mathbf{Z}$ are the population-level and participant-level design matrices, respectively, and β and μ are vectors representing the fixed- and random-effect coefficients, respectively. The random effects μ are assumed be distributed multivariate normal with mean zero. Random-effect parameters associated with different participants (indexed by j) are assumed to be independent yielding:

$$\mu_j \sim N(0, \mathbf{V}),$$

where $\mathbf{V}$, the covariance matrix, is estimated from the data. In practice, $\mathbf{V}$ is parameterized in terms of a correlation matrix $\mathbf{\Omega}$ and a vector of standard deviations σ via:

$$\mathbf{V} = \mathbf{D}(\sigma) \cdot \mathbf{\Omega} \cdot \mathbf{D}(\sigma),$$

where $\mathbf{D}$ denotes the diagonal matrix with diagonal elements σ and $\cdot$ denotes matrix multiplication. Priors are specified for the parameters σ and $\mathbf{\Omega}$. The default priors in *brms* are,

for σ , a half Student-t distribution with 3 degrees of freedom and a scale parameter that depends on the standard deviation of the response after applying the link function (minimally 10), and, for Ω , the LKJ correlation prior (Lewandowski, Kurowicka, & Joe, 2009) with parameter $\zeta = 1$ (i.e., uniform over correlation matrices of the respective dimensionality). The defaults were employed here.

Crucial to the hierarchical model used in this study are the specifications of the inverse link function, f , and the fixed and random effects design matrices, $\mathbf{X}_i$ and $\mathbf{Z}_i$. Given that our response variable was binary, the responses were assumed to be Bernoulli distributed. The inverse link function was the cumulative Gaussian psychometric function with parameters a (threshold), b (slope), γ (guess rate), and λ (lapse rate), as described in the main paper body. The value of the parameter λ was fixed at 0.02 across participants and conditions. A cumulative normal transformation was applied to the guess rate parameter, γ , such that it could take any real value within the model. This was found to improve the computational stability of the model. Design matrices specifying fixed ($\mathbf{X}_i$) and random ($\mathbf{Z}_i$) effects on the remaining parameters were specified as follows (using R code):

$$a \sim 0 + \text{blur} : \text{filt} : \text{modality} + (0 + \text{filt} : \text{modality} \mid \text{Participant})$$

$$b \sim 0 + \text{blur} : \text{filt} : \text{modality} + (0 + \text{filt} : \text{modality} \mid \text{Participant})$$

$$\gamma \sim 0 + \text{blur} : \text{modality} + (0 + \text{modalityAV} \mid \text{Participant}),$$

where *blur* is a between-participant factor specifying whether or not the talker's face was blurred, *filt* is a within-participant factor specifying whether or not of 3-7 Hz temporal modulations were removed from the auditory signal, *modality* is a within-participant factor specifying the presentation modality (auditory-only vs. audiovisual), *modalityAV* is a within-

participant factor specifying that an effect should be estimated only for audiovisual presentation modality, ‘:’ indicates an interaction, and ‘0’ indicates suppression of the intercept term.

Random effects are specified as $(\dots | \textit{Participant})$ and fixed effects are specified outside the parentheses. As specified, the fixed and random effects design matrices contain a single column for each cell in the experimental design given the specified combination of effects, with no intercept column. For example, the fixed effects term $0 + \textit{blur:filt:modality}$ indicates that the estimated fixed effects parameters, β , will reflect the population-mean values of the given psychometric function parameter for each of the 8 cells implied by the specified interaction (2 blur conditions x 2 filt conditions x 2 presentation modalities). The same interpretation can be given to the random effects terms, except that the across-participant standard deviations, σ , are estimated for each design cell implied by the specified interaction. Specification of the design matrices in this way avoids the troubling step of setting priors on deviation variables (i.e., regression coefficients that reflect deviations from a reference condition or category). The correlations, Ω , between all pairs of parameters/cells specified by the random-effects interaction terms of a and b were estimated, but no correlations were estimated for the γ parameters or between the γ parameters and the a and b parameters.

Priors on all fixed effects parameters were specified based on the expected range of psychometric function shapes that might best characterize the experimental data:

$$a \sim N(-5, 10)$$

$$b \sim N(0, 2), \quad 0 < b < \infty$$

$$\Phi^{-1}(\gamma_A) \sim N(-5, 0.001)$$

$$\Phi^{-1}(\gamma_{AV}) \sim N(-2, 5),$$

where Φ^{-1} is the inverse of the normal cumulative distribution function, γ_A is the guess rate for auditory-only conditions, and γ_{AV} is the guess rate for audiovisual conditions. The narrow prior on γ_A essentially forced that parameter to zero. The priors on a and b were applied equivalently for all of the experimental design cells.

Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1), 1-28.

Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of multivariate analysis*, 100(9), 1989-2001.

Appendix 3

A. Model of Statistical Power for Experiment 1

Statistical power to detect the direction of effects using the Bayesian one-sided p-value was determined by a simulation of the experiment. In the simulation, the audiovisual advantage (improvement in SRT) for individual listeners was sampled from (assumed) fixed-effects distributions and data from two simulated groups of listeners were analyzed using Bayesian hierarchical linear regression in *brms*. This was an oversimplification of the true analysis, which combined estimation of psychometric function parameters and statistical testing in a single model-fitting step (Appendix 2). However, it was not computationally feasible to simulate the true analysis with enough random permutations to estimate power reliably. Therefore, the audiovisual advantage (dB) was sampled from separate two-dimensional Gaussian distributions representing the two listeners groups in the study (CLEAR, BLUR), each yielding two paired observations (CTRL, FILT). The standard deviations of the Gaussian distributions were assumed to be 2.121 dB, reflecting of sum of Gaussians random variables with standard deviations of 1.5 dB (Shen, Manzano, & Richards, 2015), and the correlation between conditions was assumed to be 0.3, providing only a modest reduction in variability for within-subject contrasts. The mean audiovisual advantage in the CLEAR/CTRL condition was assumed to be 5 dB (MacLeod & Summerfield, 1990; Middelweerd & Plomp, 1987), and different size effects ranging from 1-4 dB in 0.5 dB steps were simulated for within- (e.g., CTRL > FILT) and between-subject (e.g., CLEAR > BLUR) contrasts. When within-subject effects were assessed, the between-subject effect of BLUR was assumed to be 1 dB. When between-subject effects were assessed, the within-subject effect of FILT was assumed to be 2.5 dB.

Simulated observations from the two listener groups were then submitted to a hierarchical linear regression analysis in *brms* with the following structure of fixed and random effects:

$$AVad \sim 0 + blur:filt + (0 + filt | Participant),$$

where *blur* is a between-participant factor specifying a simulated listener's group membership, *filt* is a within-participant factor specifying which condition a simulated listener was tested in, ':' indicates an interaction, and '0' indicates suppression of the intercept term. Random effects are specified as (... | *Participant*) and fixed effects are specified outside the parentheses. To estimate the Bayesian hierarchical regression model, we ran 4 independent sampling chains, each with 1000 warmup samples and 1000 post-warmup samples, resulting in 4000 total draws from the posterior. The default settings in *brms* were otherwise accepted.

A simulated within-subject effect was counted as significant if 95% of the posterior observations for the directional contrast of fixed-effects coefficients within a given listener group were greater than zero (pA simulated between-subject effect was counted as significant if 95% of the posterior observations for the directional contrast of fixed-effects coefficients across listener groups were greater than zero (ppB. *Model of Statistical Power for Experiment 2* Statistical power to detect the direction of the effect of interest using the Bayesian one-sided p-value was determined by a simulation of the experiment. In the simulation, the proportion correct in each condition (high-AVC, low-AVC) was sampled for individual listeners from an assumed fixed-effect distribution, and simulated trial data from all the listeners were analyzed using Bayesian hierarchical logistic regression in *brms*. Specifically, subject-level proportion correct was sampled from a two-dimensional Gaussian distribution yielding two paired observations (high-AVC, low-AVC). The standard deviations of the 2D Gaussian distribution were both set to 0.15 (Shen et al., 2015), and the correlation between conditions was assumed to be 0.9, corresponding to a standard deviation in the difference score (high-AVC minus low-AVC) of ~ 0.067. A very high correlation was assumed because performance in this experiment was assessed at a fixed

SNR, such that the primary source of between-subject variation in proportion correct – regardless of the within-subject condition – is the location of that fixed SNR on a listener’s psychometric function.

Once the simulated proportion correct values for each listener were drawn from the assumed fixed effects distributions, they were converted into binary trial-wise performance (correct, incorrect) where the number of trials in each condition was set to the average number of keywords across high-AVC and low-AVC items (152 trials). The simulated observations were then submitted to a hierarchical logistic regression analysis in *brms* with the following structure of fixed and random effects:

$$Resp \sim 1 + condition + (1 + condition | Participant),$$

where *Resp* is the binary dependent variable, *condition* is a within-participant factor specifying high-AVC or low-AVC, ‘:’ indicates an interaction, and ‘1’ indicates the intercept term.

Random effects are specified as (... | *Participant*) and fixed effects are specified outside the parentheses. To estimate the Bayesian hierarchical regression model, we ran 4 independent sampling chains, each with 1000 warmup samples and 1000 post-warmup samples, resulting in 4000 total draws from the posterior. The default settings in *brms* were otherwise accepted.

A simulated within-subject effect was counted as significant if 95% of the posterior observations for the directional contrast ‘high-AVC > low-AVC’ were greater than zero (pp

- MacLeod, A., & Summerfield, Q. (1990). A procedure for measuring auditory and audiovisual speech-reception thresholds for sentences in noise: Rationale, evaluation, and recommendations for use. *British Journal of Audiology*, 24(1), 29-43.
- Middelweerd, M. J., & Plomp, R. (1987). The effect of speechreading on the speech-reception threshold of sentences in noise. *J Acoust Soc Am*, 82(6), 2145-2147.
- Shen, Y., Manzano, N. K., & Richards, V. M. (2015). Psychometric functions for sentence recognition in sinusoidally amplitude-modulated noises. *The Journal of the Acoustical Society of America*, 138(6), 3613-3624.