The unity hypothesis revisited: can the gender incongruent McGurk effect be disrupted by priming?

Kennis S.T. Ma¹, Jan W.H. Schnupp^{2*}

1. The Open University, United Kingdom

2. The City University of Hong Kong, Hong Kong SAR

*Corresponding author

E-mail:wschnupp@cityu.edu.hk

Abstract

The “unity assumption hypothesis” contends that higher-level factors, such as a perceiver’s belief and prior experience, modulate multisensory integration. The McGurk illusion exemplifies such integration. When a visual velar /ga/ is dubbed with an auditory bilabial /ba/, listeners unify the discrepant signals with knowledge that open lips cannot produce /ba/ and a fusion percept /da/ is perceived. Previous research claimed to have falsified this theory by demonstrating the McGurk effect occurs even when a face is dubbed with a gender incongruent voice. But perhaps stronger evidence than just an apparent incongruence between unfamiliar faces and voices is needed to prevent perceptual unity. Here we investigated whether the McGurk illusion with gender incongruent stimuli can be disrupted by priming with appropriate pairing of face and voice. In an online experiment, 89 participants aged 18-62, were randomly allocated to experience experimental trials containing either a male or female face with incongruent gender voice. The number of times participants experienced a McGurk illusion was measured before and after a training block which familiarized them with the true pairings of face and voice. After training and priming, the susceptibility to the McGurk effects decreased significantly on average. The findings support the notion that unity assumptions modulate intersensory bias, and confirm and extend previous studies using gender incongruous McGurk stimuli.

Introduction

Audition is the dominant modality in speech perception for normal hearing individuals, as all the information necessary for perceiving speech is contained within the acoustic signals. Nevertheless, individuals do often combine both auditory and visual inputs to perceive a “best-fit utterance” which increases perceptual accuracy [1]. The McGurk effect is a compelling demonstration that speech perception is an inter-modal, rather than a unimodal, phenomenon, as it reveals a strong visual bias on auditory perception. In their classic study, McGurk and MacDonald [2] constructed incongruent visual and acoustic speech signals by dubbing a voice producing one consonant, such as /b/, over a mouth articulating a discrepant consonant, such as /g/. This induced perceivers to experience an auditory illusion, where the heard consonant may match neither the visual nor the auditory one, but instead a “fusion” consonant is perceived. For instance, many perceivers report hearing a /d/ when the visual velar phoneme /g/ is dubbed with an auditory bilabial phoneme /b/. In other instances, vision may appear to dominate. For example, when a mouth articulating a phoneme /b/ is dubbed with an auditory phoneme /d/, perceivers often report hearing the visual /b/ rather than the actual acoustic /d/ [3]. Such visual bias effects have been replicated many times, and although people appear to differ in how sensitive they are to such visual influences, the phenomenon is overall very robust [4].

Welch and Warren [5] later proposed a model of multisensory interactions, outlining conditions under which intersensory bias occurs when the brain attempts to maintain coherent perception in the face of discrepant signals from more than one modality. According to their model, the perceptual outcome of whether discrepant sensory inputs

are treated as a single multisensory object or event is determined by stimulus-driven and higher-level factors. They put forward the “unity assumption” hypothesis, which proposes that a perceiver's belief and expectation (either conscious or unconscious) determine whether two or more sensory signals are treated as a single multisensory object or event or not. According to this theory, McGurk effects should only arise if observers can plausibly make the unity assumption that the lip movement that is seen and the syllable that is heard both originate from one and the same speaker. However, Green et al. [6] later published results of experiments which used variants of the original McGurk paradigm which appeared to contradict Welch and Warren's [5] unity assumption hypothesis. Their experiments sought to investigate how robust the McGurk effect is if the visual and auditory stimuli are made “incongruent” in a manner that should make the unity assumption highly implausible. One example of such incongruency involves dubbing videos of a woman mouthing syllables for a McGurk effect with sound recording of a man pronouncing these syllables. Observers who are shown such gender incongruent audio/visual stimuli often find them peculiar, given that the mismatch between voice and face is very noticeable. Yet despite the obvious gender incongruence between voice and face, participants still unconsciously integrated the auditory and visual inputs into a single percept, resulting in a McGurk illusion. This is difficult to reconcile with Welch and Warren's framework. The gender incongruence between voice and face should provide evidence against the assumption that there is only a single speaker, and this should prevent unified, fused audio-visual perception. Nevertheless, many participants still report hearing /da/ when shown a video of a man mouthing /ga/ dubbed with the sound of a woman pronouncing /ba/.

The surprising discovery by Green et al.'s that the McGurk is robust even if stimuli are gender incongruent led them to conclude that this falsifies Welch & Warren's theory. However, perhaps it is premature to assert that gender incongruence must necessarily prevent unity. For example, in gender incongruous McGurk stimuli there is still a precise temporal synchrony between the visual and auditory stimuli. Such precise synchronization is very unlikely to arise by mere coincidence. Synchrony may therefore act as a strong binding cue, which may induce participants to perceive a strange but nevertheless single bimodal stimulus of "incongruous gender", rather than a face and a voice which are separate and just happen to be synchronized by chance. In Welch and Warren's [5] terminology, one could argue that the precise temporal synchronization in the dubbed videos creates a "highly compelling" situation in which the assumption of unity remains strong despite the peculiar voice given the face, and "intersensory bias" can change the heard /b/ to /d/.

To shed further light on this controversy, Walker et al. [7] sought to investigate whether the susceptibility of McGurk effect in gender incongruent stimuli decreases when participants knew the speaker. The logic here is that the incongruence between face and voice in a gender mismatched McGurk stimulus should be particularly striking if the face seen in that stimulus is a familiar one, and the observers know from experience that the voice is not just peculiar because it seems gender mismatched, rather it simply isn't the voice that belongs to a face they know well. To investigate whether familiarity matters, Walker et al. presented, in each trial, a familiar or unfamiliar face that was paired with a congruent (same person) or incongruent (different person of either the same or different gender) voice. They found that participants who saw the unfamiliar faces showed no differences in McGurk effect,

regardless of whether faces were paired with congruent or incongruent gender voices. This was in line with Green et al. findings [6]. However, when the face of a familiar person was combined with a different person's voice (either same or different gender), the McGurk illusion decreased, as might be predicted by Welch and Warren's theory.

One potential shortcoming of the Walker et al. study [7] is that they did not use a within-participant design, but rather compared cohorts of participants who were familiar with the faces in the stimuli with cohorts that were not. This is problematic given that recent research [8] also demonstrated a very high degree of variability between individuals in perceiving the McGurk effect. Indeed, some individuals appear insensitive to the McGurk effect and perceive it 0% of the time when tested with a particular set of AV stimuli, while other participants perceive it 100% of the time when tested with the same stimuli. Other participants may lie somewhere in between, perceiving the McGurk illusion sometimes but not always. Furthermore, the magnitude of the McGurk effect is stimulus-dependent, in the sense that the proportion of fused syllables that is perceived may vary in idiosyncratic ways depending on which faces or voices are used to construct the stimuli [9]. Finally, it remains unclear how much familiarity is really needed to lead to a reduction in the strength of the McGurk effect with gender incongruous stimuli.

In the current study we sought to establish whether the key findings by Walker et al. [7] can be replicated under a within-participant experimental design, which can cancel out potential problems stemming from the high degree of variability between individuals by measuring the strength of the gender incongruous McGurk effect both before and after a standardized amount of familiarization and priming. Stimuli for this

study were created from a subset of the validated corpus of video and audio clips for McGurk effect research known as the “Oldenburg Audio Visual Speech Stimuli (OLAVS)” [9]. The hypothesis was that the same gender incongruent stimuli would evoke fewer McGurk illusions if the participants were familiarized with the “true” voices and faces that had previously been recombined into gender incongruent dubs. In our experiment we also reminded participants of the appropriate match between face voice by priming with a gender congruous example just prior to a re-test with the gender incongruous stimuli.

Method and materials

Design

A repeated-measures design was used to investigate whether audio-visual pairings (IV1) and priming (IV2) have an effect on the susceptibility to McGurk illusion (DV). There were two levels in IV1 (phonetically matched and mismatched) and two levels in IV2 (before and after familiarization and priming). Further details on the stimuli were given below. The dependent variable was the number of times participants report hearing the fusion sound /da/ in mismatched stimuli. Counterbalancing was used to randomly allocate half of the participants to experience incongruent male voice/ female face stimuli, and the other half to the opposite pairing. Trials were randomized within blocks.

Materials

The stimuli were 6 soundtracks and 6 videos of one male and one female speaker saying /ba/, /ga/ or /da/, which were adapted from a subset of the video and audio clips for McGurk effect research known as the “Oldenburg Audio Visual Speech Stimuli (OLAVS)” [9]. The clips were processed using ffmpeg (<https://ffmpeg.org/>) to combine video and audio tracks as required, and to crop and downsample the videos to 360x360 pixels for rapid online delivery. Custom javascript code was prepared to embed the video clips into a Qualtrics online survey, which was used to allow participants to perform the experiment using their own computers and standard web browsers.

The stimuli for 12 practice and 98 experimental trials were constructed by combining voices and faces of the /ba/, /ga/ and /da/ syllable recordings from one male and one female actor from the OLAVS set (identified in [9] as “TK05” and “TK01”), as required. For example, phonetically matched (A:/ba/ V:/ba/ or A:/ga/ V:/ga/) and mismatched (A:/ba/ V:/ga/) combinations were produced to test whether and how often the mismatch induced a McGurk fusion percept /da/ (IV1). In addition, the phonetically matched stimuli provided a simple means for assessing data quality. Correctly identifying the unambiguous syllable presented during phonetically matched stimuli ought to be easy, and frequent failures to identify these syllables correctly gives a strong indication that a participant was inattentive or otherwise unable to perform the task adequately, and that their data needed to be excluded. Meanwhile, AV combinations could be either gender incongruent for testing the strength of the McGurk effect with incongruent stimuli both at baseline (Block 1) and after familiarization and

priming (Block 3), or they could be congruent during familiarization and priming (Blocks 2 and 3, IV2, see Procedure below for further details. Trials were randomized within blocks so participants could not predict which syllable comes next.

A list of instructions, a debriefing sheet, and informed consent forms were also incorporated into the Qualtrics survey used to conduct the experiment.

Procedure

All procedures were conducted in compliance with the British Psychological Society's Code of Human Research ethics and approved by the DE300 ethical review board of the Open University. Motivated by social distancing guidelines during the Covid-19 pandemic, the experiments were conducted online with no experimenters presented. Participants used their own computers and headphones to access an online survey implemented on the Qualtrics platform which guided them through the experiment. Informed consent was obtained before the experiments. Participants were informed of the nature of the experiment and were instructed to watch trial videos attentively in a quiet room with no distractions. While this type of online format afforded the experimenters little ability to standardise stimulus presentation or monitor the experiment, these issues could be mitigated by the within-subject design of the study. Individual differences in experimental settings which arise as each participant uses their own consumer grade hardware for stimulus presentation in their own homes or offices can be factored out as each subject serves as their own control. Arguably the

online format also increases the ecological validity of the study, as data were collected in the subjects' "natural habitat".

Before the experiment, participants answered a questionnaire asking for voluntary information about participants' age and gender, and asked them to confirm self-reported normal hearing and normal or corrected-to-normal vision. To preserve the participants' anonymity, no further identifying personal information was collected. Before the start of the experiment proper, participants were guided through 12 practice trials to familiarise them with the trial format and user interface. For the trials, the experimental video and audio clips were embedded into online survey questions using custom written javascript code.

A

Hit the "Play" button below or press Alt+A, then select the spoken syllable you heard:

Click here or press Alt+A to Play.

- ☐ Ba
- ☐ Ga
- ☐ Da

Block 1: Measure strength of McGurk illusion with gender incongruous stimuli at baseline.

20 phonetically matched and 20 mismatched AV syllables.

B

Hit the "Play" button below or press Alt+A, then watch the video and count the number of times the person in the video on the LEFT said: "Da", ignoring the speaker on the right. When the videos disappear, select your answer.

Click here or press Alt+A to Play.

- ☐ 3
- ☐ 5
- ☐ 7

Block 2: train participants on the "real" pairing of voices and faces. 18 videos (~30 s long) showing two faces side by side, both uttering, at the same time, syllables /ba/, /ga/ or /da/ at random intervals. Participants had to count how many times the face that they had encountered in Block 1 uttered a given syllable while ignoring the other face.

C

Hit the "Play" button below or press Alt+A, then select the SECOND spoken syllable you heard:

Click here or press Alt+A to Play.

- ☐ Ba
- ☐ Ga
- ☐ Da

Block 3: Measure strength of McGurk illusion with gender incongruous stimuli after training and with priming. Participants were shown two clips in quick succession, the first showing a gender congruent utterance, the second the same face bit with a gender incongruent utterance. They reported the syllable heard during the 2nd, mismatched interval.

Fig 1: Screenshots illustrating the stimuli presented to the subjects during each of the 3 blocks. The experiment first measured the strength of the gender incongruent McGurk illusion in each participant (A), then engaged the participants in a task that encouraged them to form a correct association between the faces and voices used in the experiment (B), and finally retested the participants with gender incongruent stimuli, but in this final, post-training block subjects were additionally reminded of the correct face-voice pairing by priming just prior to the presentation of each test stimulus (C)

Block 1 measured participants' baseline performance with gender incongruent stimuli. Participants were instructed to watch randomly interleaved presentations of 20 phonetically matched and 20 mismatched AV syllables. In each trial, a voice producing /ba/ or /ga/ was dubbed with a ~ 2 s long, gender incongruent face articulating either /ba/ or /ga/. Participants reported whether they heard /ba/, /ga/ or /da/ using their computer mouse or keyboard (A three-alternative forced choice design). Percepts of /da/ were counted as fusion percepts.

Block 2 familiarized participants with the “real” (congruent) voices and faces that had been recombined to generate the incongruent Block 1 stimuli. Participants watched 18 videos, each ~30 s long, showing two faces side by side, each uttering, simultaneously, syllables /ba/, /ga/ or /da/ a random number of times at random partly overlapping intervals. All these utterances were gender congruent and AV matched, so that each face spoke with its actual voice. Participants had to count how many times the face that they had encountered in Block 1 uttered either /ba/, /ga/ or /da/, while ignoring the other face, which spoke with the voice from Block 1. This task therefore

required participants to associate the Block 1 face with its true voice, and to dissociate it from the incongruent Block 1 voice.

Block 3 re-measured the strength of the McGurk effect with the gender incongruent stimuli from Block 1. Participants again watched 40 phonetically matched or mismatched AV stimuli and reported what they heard (/ba/, /ga/ or /da/) from the second clip. However, to remind the participants of the familiarization in Block 2, in each trial, participants were shown two clips in quick succession. First, a ~2 s gender congruent stimulus reminded participants of the face that actually belonged to the voice in the upcoming target (priming). This was followed immediately by a gender incongruent clip identical to the ones used in Block 1. Participants were instructed to use the mouse cursor to select the syllable heard during the second clip (3-AFC).

The entire experiment took ~25 minutes. At the end of the experiment, the purpose of the experiment was explained in greater detail with a debriefing sheet. Participants were given the opportunity to ask any questions by email and a personal number was assigned to each participant to enable them to withdraw their data.

Data Analysis

The 3-AFC responses were coded as a “fusion percept” variable. Participants reporting that they heard the syllable “da” was scored as a “fusion”, and hence an incidence of a McGurk illusion. Other reported auditory percepts (“ba” or “ga”) were scored as unfused. The main objective of this study was to determine whether the proportion of fused auditory percepts reported during gender incongruent stimuli decreases after familiarization and priming with the correct voice and face matches.

The within-subject design allowed us to examine this question by comparing “fusion” scores at baseline (Block 1) with those obtained after training and priming (Block 3). Two statistical tests were used to assess the statistical significance of these results. A 2x2 repeated measure factorial ANOVA with independent variables “Block” (two levels: either “Block 1 baseline” or “Block 2, primed”) and “AV match” (two levels: mismatch between auditory and visual utterance for McGurk stimuli, or match for control stimuli) would confirm an influence of priming on the likelihood of experiencing a McGurk illusion as a significant interaction effect. In addition, a sign-test was used to compare numbers of fusion percepts before and after priming, confirming the results of the ANOVA using a non-parametric approach.

Participants

One hundred and sixty-nine participants were recruited from online platforms, including the Open University’s student forum, the McGill University Auditory List, social media, as well as word of mouth. Participants received no payment or other incentives to take part in the experiment, and 65 participants dropped out before completing the experiment. Another 15 participants were excluded for failing to report the correct syllables in the control audio only trials or failing to perform at a high level (>85% correct) for matched AV trials (compare Fig 2). Of the remaining 89 participants, 10 did not report their age or gender, but were included for all analyses except for the age range calculation (ages 18 to 62 years, mean 36.85, SD 11.75) or gender counting (63 females and 25 males). All participants reported normal or corrected-to-normal

vision and normal hearing. Participants gave informed consent and were naïve regarding the specific aims of the study before debriefing.

Fig 2: Proportion of AV matched trials in which the syllables were correctly identified by each candidate. Correct identification of such syllables ought to be fairly easy, and is a necessary precondition for the experiment. Data from participant achieving <85% correct (pooled over blocks 1 and 3) were therefore excluded from further analysis.

Results

Fig 3: A) Each line connects the % of fusion percepts reported by one participant in blocks 1 and 3 respectively. A small random jitter ($\pm 3\%$) is added to spread out overlapping data points. **B)** Distribution of changes in % fusion percepts between block 1 and 3 for all participants.

Fig 3A shows the proportion of AV mismatched, incongruent stimuli that gave rise to McGurk fusion percepts in blocks 1 and 3 respectively for all participants. Lines of different colors connect the proportions of fusion percepts in each block for each participant. Several features are noteworthy: first, participants differed greatly in the proportion of fusions perceived, exhibiting the full range, from 0 to 100% fused percepts. Second, while for many participants the proportion of fused percepts hardly changed between blocks 1 and 3, others exhibited large changes. Third, decreases in the proportion of fusion percepts appear to be more frequent than increases. These observations are confirmed in Fig 3B, which shows the distribution of changes in % fusion percepts (3rd minus 1st block) in histogram form. In 34 cases, the number of

fusion percepts decreased, compared to 12 increases and 43 cases of zero change. The histogram thus has a heavier negative tail.

An analysis of the data using a 2x2 repeated measure factorial ANOVA confirmed significant effects, both for AV pairing ($F(1, 88) = 57.57, p < .001, \eta^2 = .395$), and for Block ($F(1, 88) = 5.27, p = .024, \eta^2 = .056$). Participants perceived significantly more fusion percepts /da/ when presented with discrepant visual and acoustic inputs (mean=6.72, SD=8.1), compared to matched AV pairings (mean=.09, SD=.33). After familiarization Block 2, participants heard overall significantly fewer fusion percepts /da/ (mean=5.35; SD=7.62), compared to the reported fusion percepts before priming (mean=6.72; SD=8.1). Also, there was a significant interaction between AV pairing and priming ($F(1, 88) = 7.1, p = .009, \eta^2 = .075$) such that participants reported fewer fusion percept /da/ for the discrepant AV pairing after priming. A sign test further confirmed that the overall downward trend in fusion percepts after priming was statistically highly significant ($p = .0016$)

Post-hoc sign tests (bonferroni corrected) revealed highly significant differences in the number of observed fusion percepts for all (six) pairwise comparisons except for AV matched stimuli before and after priming ($p = 1$), as these were close to zero at the outset and did not significantly change with priming. Participants perceived on average significantly more fusion percepts /da/ when presented with discrepant visual and acoustic inputs, compared to matched AV pairings ($p < .001$). After familiarization Block 2, participants heard on average significantly fewer fusion percepts /da/ during AV mismatched trials than before priming ($p = .009$) (see also Fig 4.).

Fig. 4: Bonferroni corrected post hoc test results. P -values were computed with pairwise sign tests comparing each condition and Bonferroni corrected by multiplying by 6 and capping at 1. They are shown in the brackets above the bars which indicate mean fusion percept counts across the whole cohort in each condition. Gray brackets denote non-significant differences.

Discussion

This study found a significant reduction in the susceptibility to the McGurk illusion in gender incongruent stimuli after priming a face with an appropriate same-gender voice. Before priming, some participants frequently perceived the fusion percept /da/ when presented with discrepant visual /ga/ and acoustic /ba/, and reported the acoustic syllables correctly in most audio-visual matched trials. After priming, these participants heard fewer fusion percept /da/ when presented with discrepant audio-visual inputs. The size of the observed effect was medium.

These results support the hypothesis that familiarization and priming reduced the number of fusion percept specifically for the discrepant A:/ba/V:/ga/ pairings, but not for phonetically matched AV inputs. Bear in mind though that the number of fusions for the matched inputs was close to zero, as expected, and large reductions could not be expected. After knowing a particular voice belongs to a particular face, participants were on average less likely to form assumptions of unity when that same voice was dubbed with an incongruent, opposite gender face. Priming here served as a short-term context cue that strongly reminded participants that a particular voice belongs to a particular face, and thus less assumption of unity was formed when the same voice was dubbed with incongruent opposite gender face. The present study demonstrated that knowledge and experience (i.e. top-down factors) influenced multisensory integration and modulated intersensory bias in humans speech perception and these findings suggested that immediate perceptual response to intersensory discrepancy can be induced or reversed by priming.

However, note that not all participants experienced such a priming effect. Indeed, Fig 3B shows that for many (43 out of 89) participants, the number of fusion percepts did not change at all, and for 12 it went up, compared to 34 for whom it went down. Thus, even though there was a robustly significant effect of familiarization of priming on average of the whole cohort, for a substantial number of participants the McGurk effect remained stubbornly robust even after they had been both trained and primed just seconds earlier that the voice did not match the face. Also noteworthy is that 36 participants, about 40% of the total, did not experience a single McGurk illusion during the entire experiment. Individual differences are therefore very substantial. Nevertheless, both the non-parametric sign test and the ANOVA performed gave a highly significant result, providing conclusive evidence in support of the hypothesis that familiarization and priming on average reduce the McGurk effect with incongruous stimuli, considerable subject-to-subject variability in the susceptibility to and malleability of the McGurk illusion notwithstanding.

This study was partly inspired by Green et al. [6] who asked their participants to watch clips containing discrepant phonetic information that were either gender congruent or incongruent, and found that participants still integrated the auditory and visual inputs into a single percept regardless of the obvious gender incongruence between the voice and face. Our study supports Green et al.'s observation that McGurk illusions are common even if stimuli are gender incongruous, but it also demonstrates that McGurk illusions with incongruent gender stimuli reduced after priming, which is partly at odds with Green et al.'s conclusion that the gender incongruous McGurk effect falsifies Welch and Warren's [5] unity assumption hypothesis .

Welch and Warren theorized that intersensory bias formed only under certain conditions, which include “temporal synchrony” and “general experience” [5]. In the McGurk illusion, the temporal synchrony between seen lip movements and heard audio is thought to create a “highly compelling” situation, in which the assumption of unity is very strong. However, “general experience” ought to tell observers that male voices and female faces do not go together. Green et al. [6] therefore interpreted the existence of gender incongruent McGurk illusions as evidence against Welch and Warren's theoretical framework, which posited the importance of top-down influences shaped by experience in fusing sensory data within and across modalities into unitary events [5]. A wealth of research has manipulated various factors that may lead to different levels of the intersensory bias and demonstrated that the unity assumption can be induced explicitly by experimenter's instruction [10], or evoked by stimuli with highly congruent properties and characteristics [11], temporal synchrony [12], or when perceivers have learned to associate multimodal inputs from prior experience, either long-term knowledge [13] or short-term contextual cues [14].

However, Green et al.'s objection [6] may simply underestimate how strongly human perception weights temporal synchrony as a binding cue, and therefore the quality and quantity of “general experience” that might be required to overrule it. Synchrony is such a powerful cue in the auditory domain that it can bind frequency components presented to opposite ears into a single harmonic sound [15], or powerfully override evidence for spatial disparity between visual and auditory sources in the “ventriloquism effect” [10]. Thus, very large amounts of fairly specific “general experience” may be needed to overrule powerful temporal synchronicity cues which favour perceptual unity.

Walker et al. [7] were the first to show that this may be the case, by demonstrating that gender incongruous McGurk illusions are less likely to be observed when participants were familiar with the faces used in the construction of the incongruous stimuli, and the experience that told them that the voice and face were mismatched therefore went beyond just a general sense that the voice was somehow incongruous given the gender of the face. The effect was modest, with the mean percentage of fusion percepts for gender incongruent stimuli declining from 28% for unfamiliar faces to 21% for familiar faces, but these numbers are not very different from those observed here, with 34% of incongruent, mismatched stimuli evoking fusion percepts before priming, and this number declining to 27% after priming. The present study thus similarly concludes that the susceptibility of McGurk effect in gender incongruent stimuli will on average decrease if participants are trained to dissociate a face from a mismatched, incongruous voice and are powerfully reminded of the incongruence through priming. Like Walker et al.'s study [7], our results therefore support Welch and Warren's theory of intersensory interactions [5] overall, but the large individual differences and the relatively modest overall effect size we observed nevertheless point to the fact that top-down influences are at times perhaps surprisingly weak compared to the very strong bottom-up perceptual binding cues that are created by temporal synchrony.

As well as supporting Welch and Warren's theory [5], our study has gone some way to overcome the methodological shortcomings of previous research. In the original study of McGurk and McDonald [2], 98% of adult participants reported hearing the fused percept of /d/ when a visual bilabial phoneme /g/ dubbing with a auditory velar phoneme /b/, but later studies found a high degree of variability across individuals [8], and this variability may interact with the specifics of the video and audio clips used to

construct the stimuli [9]. By using a within-participants design in which standardised amounts of experience were provided to familiarize participants with the congruous face-voice pairings, individual differences could be excluded as a potential confound. Also, by crowdsourcing data collection on the internet, this study was able to recruit a comparatively large number of participants. Walker et al.'s study [7] had 18 participants in their experimental group and 18 in the control group. Our study, with 89 participants tested in both the control and the experimental conditions, has greater statistical power than any previous exploration of this phenomenon.

Since the experiment was carefully designed and controlled from the outset, this study has a high internal validity and a high level of population validity, which provided robust evidence with cross-population generalizability and thus the findings can be generalized to wider populations, in as far as phenomena marked by such high inter-subject variance lend themselves to generalization. Furthermore, the exclusion criterion of requiring at least 85% correct responses in matched AV trials ensured the integrity of the data.

On the other hand, there are several limitations of this study. Firstly, conducting the experiment online due to social distancing restrictions may add to the variance in the data given that each participant used their individual equipment, precluding careful calibration and standardisation of listening conditions. Similarly, the use of individual computers precludes monitoring of participants' eye positions, and one cannot be certain to what extent participants followed the instruction to fixate on the speaker's face. This is unlikely to be a big problem though given that the moving face would have

been the most salient stimulus in their environment, and that the proportions of McGurk illusions observed are comparable to those reported in previous studies.

Overall, this study complements previous studies which support the unity assumption hypothesis, providing a clear example of the top-down modulation of human multisensory perception. Follow-on studies could map out the relative weighting of bottom-up synchronicity cues versus top-down influences of incongruence or experience by testing whether decreasing the precision of the temporal synchrony between audio and video clips allows for stronger top-down effects.

Funding:

Supported by Hong Kong General Research Fund grant Nr 11100518

References

1. Schnupp J, Nelken I, King A. *Auditory neuroscience: making sense of sound*. Cambridge. The MIT Press; 2011.
2. McGurk H, MacDonald J. Hearing lips and seeing voices. *Nature (London)*. 1976; 264(5588): 746–748. doi: 10.1038/264746a0.
3. MacDonald J, McGurk H. Visual influences on speech perception processes. *Perception & psychophysics*. 1978; 24(3): 253–257. doi: 10.3758/BF03206096.
4. Tiippana K. What is the McGurk effect. *Frontiers in psychology*. 2014; 5: 725–725. doi: 10.3389/fpsyg.2014.00725.
5. Welch RB, Warren DH. Immediate perceptual response to intersensory discrepancy. *Psychological bulletin*. 1980; 88(3): 638–667. doi: 10.1037/0033-2909.88.3.638.
6. Green, KP, Kuhl PK, Meltzoff AN, Stevens EB. Integrating Speech Information across Talkers, Gender, and Sensory Modality: Female Faces and Male Voices in the McGurk Effect, *Perception & psychophysics*. 1991; 50(6): 524–536. doi: 10.3758/BF03207536.
7. Walker S, Bruce V, O'Malley C. Facial Identity and Facial Speech Processing: Familiar Faces and Voices in the McGurk Effect. *Perception & psychophysics*. 1995; 57(8): 1124–1133. doi: 10.3758/BF03208369.
8. Mallick DB, Magnotti JF, Beauchamp MS. Variability and stability in the McGurk effect: contributions of participants, stimuli, time, and response type. *Psychonomic bulletin & review*. 2015; 22(5): 1299–1307. doi: 10.3758/s13423-015-0817-4.
9. Stropahl M, Schellhardt S, Debener S. McGurk stimuli for the investigation of multisensory integration in cochlear implant users: the Oldenburg Audio Visual Speech Stimuli (OLAVS). *Psychonomic bulletin & review*. 2017; 24(3): 863–872. doi: 10.3758/s13423-016-1148-9.
10. Warren DH, Welch RB, McCarthy TJ. The role of visual-auditory “compellingness” in the ventriloquism effect: implications for transitivity among the spatial senses. *Percept Psychophys*. 1981; 30: 557–564. doi: 10.3758/BF03202010.
11. Chen YC, Spence C. Assessing the Role of the 'Unity Assumption' on Multisensory Integration: A Review'. *Frontiers in psychology*. 2017; 8: 445–445. doi: 10.3389/fpsyg.2017.00445.
12. Munhall KG, Bribble P, Sacco L, Ward M. Temporal constraints on the McGurk effect. *Perception & Psychophysics*. 1996; 58(3): 351–362. doi: 10.3758/BF03206811.

13. Doehrmann O, Naumer MJ. Semantics and the multisensory brain: How meaning modulates processes of audio-visual integration. *Brain research*. 2008; 1242: 136–150. doi: 10.1016/j.brainres.2008.03.071.
14. Nahorna O, Berthommier F, Schwartz JL. Audio-visual speech scene analysis: characterization of the dynamics of unbinding and rebinding the McGurk effect. *The Journal of the Acoustical Society of America*. 2015; 137(1): 362–377. doi: 10.1121/1.4904536.
15. Beerends JG, Houtsma AJM. Pitch identification of simultaneous dichotic two-tone complexes. *The Journal of the Acoustical Society of America*. 1986; 80(4): 1048–1056. doi: 10.1121/1.393846.