

The sliding window of audio–visual simultaneity

Warrick Roseboom

School of Psychology, The University of Queensland,
Brisbane, Australia



Shin'ya Nishida

NTT Communication Science Laboratories, Nippon,
Telegraph & Telephone Corporation, Kanagawa, Japan



Derek H. Arnold

School of Psychology, The University of Queensland,
Brisbane, Australia



Humans exist in an environment wherein many unrelated events occur in close spatial and temporal proximity. Audio–visual timing experiments, however, have often examined only isolated pairs of sensory events. We therefore decided to assess how audio–visual timing perception would be shaped by the presence of an additional audio or visual event. We found that the point of subjective synchrony for a sensory event can be shifted away from the presence of other temporally proximate events. These interactions made audio–visual pairs seem unrelated, or asynchronous, at timings at which they had seemed synchronous when presented in isolation. This shows that the interval across which humans are insensitive to audio–visual asynchrony is not fixed, but dynamic, shaped by interactions between multiple sensory events. Importantly, we establish that these interactions can enhance the sensitivity of timing judgments. These interactions could therefore help to segregate unrelated sensory events across time. Such effects are likely to be common in the cluttered environments in which humans exist.

Keywords: timing perception, audio–visual, perceptual segregation, sensory integration, audio–visual simultaneity window

Citation: Roseboom, W., Nishida, S., & Arnold, D. H. (2009). The sliding window of audio–visual simultaneity. *Journal of Vision*, 9(12):4, 1–8, <http://journalofvision.org/9/12/4/>, doi:10.1167/9.12.4.

Introduction

Isolated audio (AUD) and visual (VIS) events can be judged as synchronous across a fairly broad range of physical offsets (~ -100 to $+200$ ms from physical synchrony; see Dixon & Spitz, 1980). This has traditionally been taken to indicate that the binding of AUD and VIS inputs is subject to a large temporal *window*—referred to as the audio–visual (AV) simultaneity window (Lewald, Ehrenstein, & Guski, 2001; Lewald & Guski, 2003; Spence & Squire, 2003). This premise implies that pairs of AUD and VIS events are only reliably judged as *asynchronous* if separated by an interval that extends beyond this window.

The AV simultaneity window has been shown to be flexible with regards to input type. For example, larger AV simultaneity windows are associated with AUD and VIS speech, relative to simple stimulus pairs such as sound bursts and light flashes (Dixon & Spitz, 1980; Spence & Squire, 2003). This flexibility has been linked to prior experience and expectation effects (Guski & Troje, 2003; Jackson, 1953; Radeau & Bertelson, 1987; Spence, 2007; Vatakis & Spence, 2007; Welch & Warren, 1980).

Previous investigations of AV simultaneity have often focused only on isolated pairs of sensory events. It is therefore unclear how the AV simultaneity window might

be shaped by the presence of additional AUD or VIS events. The possibility that the AV simultaneity window might be shaped by additional events is suggested by temporal ventriloquism, where the apparent timing of cross-modal pairs of AUD and VIS events can be drawn toward one another (Fendrich & Corballis, 2001; Morein-Zamir, Soto-Faraco, & Kingstone, 2003; Vroomen & Keetels, 2006).

To assess how the AV simultaneity window is shaped by additional VIS or AUD events, we contrasted AV simultaneity judgments in the presence and absence of additional, temporally offset, events. We found that the point of subjective synchrony for an AUD and VIS pair can be shifted *away* from the presence of additional AUD or VIS events—a temporal *segregation*. Furthermore, we found that the presence of additional events can enhance the sensitivity of timing judgments.

General method

Apparatus

Experiments were run on a Dell Pentium 4 PC. Visual stimuli were displayed on a gamma-corrected 21-in. Samsung SyncMaster 1100p+ monitor (resolution

1024 × 768 pixels, refresh rate 120 Hz) and generated using a ViSaGe from Cambridge Research Systems (CRS). Participants viewed stimuli from 57 cm, with their head placed in a chinrest. AUD signals were generated with a TDT Basic Psychoacoustic Workstation (Tucker-Davis Technologies) and were presented diotically via Sennheiser HDA 200 headphones. AUD presentations were synchronized with the VIS display using triggers from the ViSaGe, timed to coincide with a monitor refresh. Participants' responses were recorded using a CRS CB6 Response Box.

Experiment 1A

This experiment was conducted to assess how subjective AV timing would be shaped by the presence of additional sensory events *within* the classically defined AV simultaneous window.

Method

Participants included two of the authors and six volunteers who were naïve as to the experimental purpose. All reported normal hearing and normal, or corrected-to-normal, visual acuity.

The VIS stimulus consisted of two static discs (subtending 0.8° of visual angle), colored black (~ 0 cd/m²) or white (CIE 1931 $x = 0.279$, $y = 0.293$, 120 cd/m²), separated vertically by 0.8° of visual angle (dva), and presented against a gray (60 cd/m²) background. Static discs were located above and below a central cross hair fixation point (subtending 0.5 dva). There was also either one, or two, moving discs that translated horizontally back and forth at a rate of 5 dva/sec at a periodicity of 0.5 Hz. See [Movies 1–2](#) for examples of stimulus appearance. VIS events consisted of repetitive collisions (0.5 Hz) between the moving and corresponding static discs. AUD events consisted of a transitory change (8.33 ms) in the frequency (from 400 to 900 Hz) of a persistent tone (68 dB SPL).

In the experiment, there were two conditions: Baseline and Test. In the Baseline condition (see [Movie 1](#)), a single standard visual event (Vstd) was presented. Depicted in [Figure 1A](#) is a space–time plot of a single cycle of the VIS stimulus in the Baseline condition. The dotted black line represents the trajectory of the Vstd moving disc. The solid red line depicts the position of the corresponding static disc. The distance between the broken green lines depict the range over which AUD events were presented (11 intervals presented across ± 250 ms from Vstd timing). The depicted cycle starts (-1000 ms) with the moving disc at maximal separation from its static counterpart. The moving disc then converges on the static disc, colliding




Figure 1. Space–time plots depicting a single animation cycle of the VIS stimulus during: (A) Baseline trial presentation with a single moving disc; (B) Test presentations with a preceding Vad; (C) Test presentations with a succeeding Vad. See text for further details.

after 1000 ms (0 ms). At 2000 ms (+1000 ms) they are again at maximal separation.

In the Test condition (see [Movie 2](#)), the Vstd was accompanied by an additional VIS event (Vad), which either preceded or succeeded Vstd by 100 ms. [Figure 1B](#) depicts a Test condition with a preceding Vad. The dotted black line again represents the trajectory of the Vstd, while the dotted red line represents the trajectory of the Vad. At the beginning of the cycle (−1000 ms), the moving Vad disc is 100 ms advanced from Vstd in the cycle. Therefore, the Vad pair collide after only 900 ms (−100 ms), reach maximal separation after 1900 ms, and return to their start positions after 2000 ms (+1000 ms). [Figure 1C](#) depicts the same as above, though with a succeeding Vad disc pair. In this case, the moving Vad disc does not reach maximal separation from its static counterpart until 100 ms into the cycle. It then collides after 1100 ms (+100 ms). The locations of the VIS events (above or below fixation) were randomized on a trial-by-trial basis. On each trial, the 2-sec animation cycle was presented three times or was terminated upon a participant response.

In the Baseline condition, participants were required to fixate centrally and make AV synchrony/asynchrony judgments. In the Test condition, participants judged if AUD events were synchronous with the VIS event above fixation, below fixation, with both VIS events, or with neither of the VIS events, a four-alternative-forced-choice task. Participants were not aware which of the two VIS events was the Vstd. Each of the three types of presentation (Baseline, preceding Vad and succeeding Vad) were interspersed during runs of trials in a random order. Each run of trials consisted of 132 individual trials, four trials for each of 11 auditory offsets for Baseline presentations, for Test presentations with preceding Vad events, and for Test presentations with succeeding Vad events. Each participant completed two runs of trials.

Results

During Baseline trials, responses were coded as 1 (synchronous) or 0 (asynchronous). During Test trials, responses were coded as 1 (AUD–Vstd synchrony reported, or AUD synchronous with both VIS events) or 0 (AUD–Vad synchrony reported, or AUD asynchronous with both VIS events). Points of subjective synchrony (PSS) between AUD and Vstd were determined by taking the peaks of Gaussian functions fitted to the distributions of these responses (see [Figures 2B](#) and [2C](#); also see [Supplementary Figure 2](#) for distributions depicting reports of *both* VIS events as synchronous with AUD and distributions of synchrony reports with Vad).

Analysis of data from the Baseline condition revealed that the timing of preceding and succeeding Vad events (± 100 ms) fell well within the classically defined AV simultaneity window for isolated AV events (the full-width half-height of Gaussian functions fitted to distributions




Figure 2. (A) Bar plot showing average PSS shifts, for AUD and Vstd, in the presence of Vad (offset ± 100 ms from Vstd). Error bars show ± 1 standard error across eight participants. (B) Distributions of reported AUD and Vstd synchrony as a function of the physical relative AUD timing averaged across eight participants. Black data points were derived from Baseline trials; red data points from Test trials with preceding Vad events. The unbroken red vertical bar depicts Vad timing; the unbroken vertical black line represents Vstd timing. Dotted vertical bars correspond with peaks of fitted Gaussian functions; black for Baseline trials and red for Test trials. (C) As above, but with succeeding Vad events.

of reported AV synchrony in the Baseline condition; 196 ± 6 ms).

In the Baseline condition, AUD and Vstd events seemed synchronous when AUD events lagged Vstd by 16 ms (Figures 2B and 2C). In Test trials with *preceding* Vad events, AUD and Vstd seemed synchronous when AUD events lagged Vstd by 67 ms (Figure 2B). Thus, the PSS for AUD and Vstd was shifted *forward* in time (~ 50 ms; $t_7 = 6.51$, $p < 0.001$, paired samples—two tailed; Figure 2A), *away* from the presence of a preceding Vad event, relative to when no Vad was present. The magnitude of this effect is approximately half the physical temporal distance between Vstd and Vad.

Similar effects were observed in Test trials with *succeeding* Vad events. In this case, AUD and Vstd seemed synchronous when AUD events *preceded* Vstd by 9 ms (Figure 2C). Therefore, the PSS for AUD and Vstd was shifted *backward* in time (~ 25 ms, one quarter the physical temporal distance; $t_7 = -3.42$, $p = 0.011$, paired samples—two tailed; Figure 2A) relative to the Baseline condition. Here again the shift of apparent Vstd timing, as indicated by the PSS between Vstd and AUD events, was away from the succeeding additional event.

Experiment 1B

To assess the generality of effects from Experiment 1A, we examined if similar shifts could be induced with additional AUD events. The methods were similar to those described above with the following exceptions. We effectively reversed AUD and VIS roles, so there was just one VIS event and either one or two AUD events (Astd and Aad). Timing relationships were determined relative to Astd, with VIS events presented at 11 intervals ± 250 ms from Astd and, when present, Aad offset ± 100 ms from Astd (see Movie 3). When only a single AUD event was present, participants made AV synchrony/asynchrony judgments. In the presence of both Astd and Aad, participants judged if VIS events were synchronous with the first AUD event, the second AUD event, with both AUD events, or with neither of the AUD events. The order of AUD events, first or second, was determined on a trial-by-trial basis such that participants were not aware which of the two was the Astd. This experiment was completed by six participants, including two of the authors and four participants naïve as to the experimental purpose.

Results of Experiment 1B were similar to those of Experiment 1A. When Aad events preceded Astd, the PSS for Astd and VIS was shifted *forward* in time (~ 57 ms; $t_5 = 8.54$, $p < 0.001$, paired samples—two tailed).

However, when Aad events succeeded Astd, the PSS was shifted *backward* in time (~ 26 ms; $t_5 = -2.63$, $p = 0.047$, paired samples—two tailed). This shows that the

timing shifts we have identified are not limited to interactions involving two VIS events but can occur for other combinations of sensory events.

Experiment 2

The results of Experiments 1A and 1B show that the presence of an additional sensory event, *within* the classically defined AV simultaneity window, can shift the PSS for a sensory event. This shift in PSS is *away* from the timing of an additional event. As a consequence, AV pairs can seem *less* related, or asynchronous, at epochs at which they had seemed synchronous when presented in isolation. This brings about an altered simultaneity window for the AV pair.

The shifts of apparent timing in Experiments 1A and 1B may be indicative of a functional process of temporal segregation. Timing judgments may tap competitive processes such that the perceptual coupling of an AUD and a VIS event makes other events seem *less* related and therefore less synchronous. This process results in an apparent shift in timing of perceptually coupled events *away* from the timing of unpaired sensory events. This proposal prompts the possibly counter-intuitive prediction, that AV timing judgments might be more sensitive in environments containing multiple competing events, than they are in isolated settings.

An alternate possibility is that the timing shifts of Experiments 1A and 1B are indicative of a changed response criterion. For instance, qualitatively similar timing shifts could occur if participants are disinclined to report that both (clearly separated) intra-modal events are synchronous with the cross-modal event when cross-modal timing relationships are unclear. This process would result in worse, not better, AV timing sensitivity in the presence of additional sensory signals.

To assess these possibilities, Experiment 2 used a signal detection task to examine the sensitivity of AV timing judgments in the presence and absence of additional VIS events.

Method

Methods in Experiment 2 were similar to Experiment 1A with the following exceptions. Participants included one of the authors and an additional five volunteers who were naïve as to the experimental purpose. All reported normal hearing and normal, or corrected-to-normal, visual acuity.

There were two conditions. In the Baseline condition, a single VIS event (Vstd) and a single AUD event were presented. AUD events were either synchronous with Vstd events (Figure 3A), or they preceded (Figure 3B) or succeeded Vstd by 75 ms. Participants reported whether

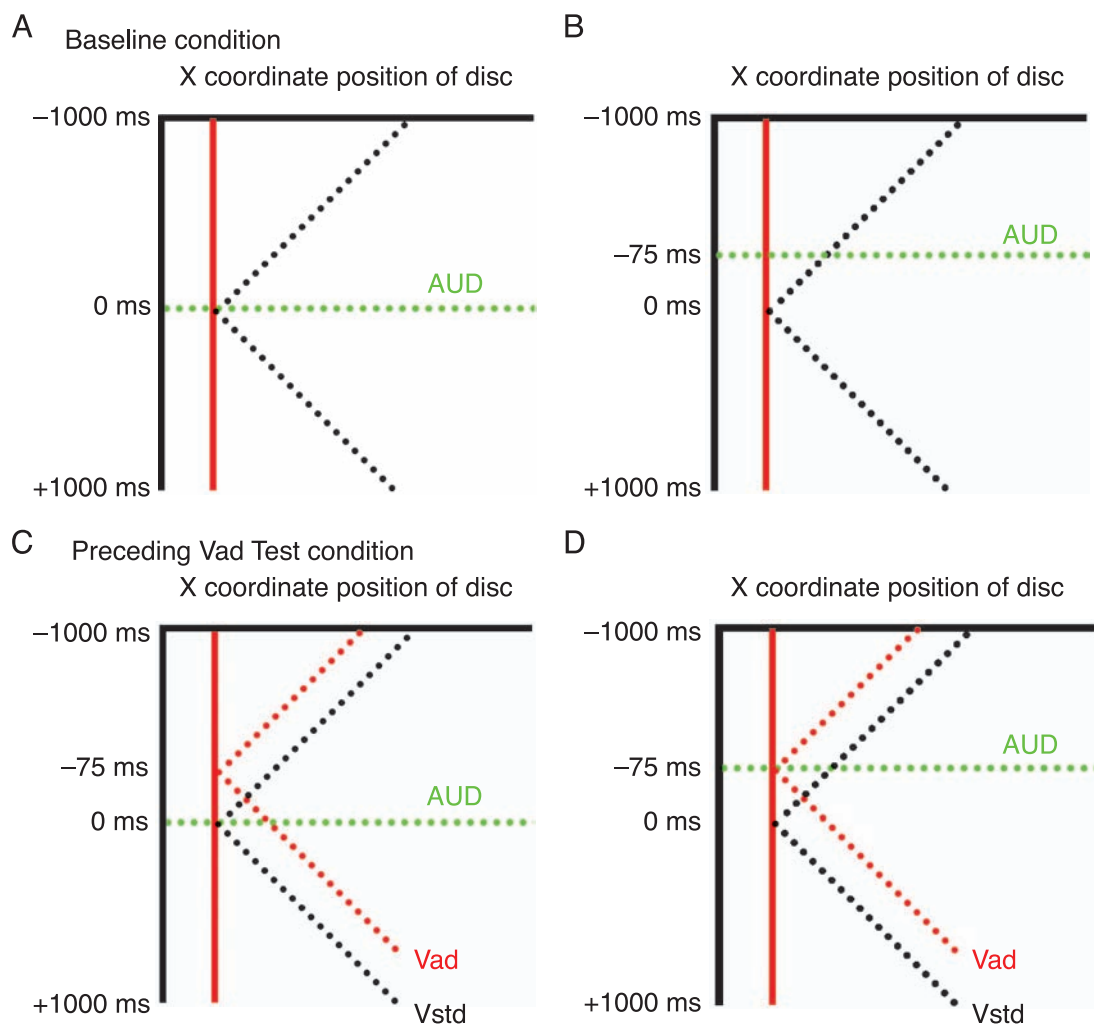


Figure 3. (A–B) Space–time plots depicting trials from Experiment 2 that contained a single VIS event. Dotted black lines depict the trajectories of moving discs. Solid red lines depict the position of static discs. Note the reversal in the trajectories of moving discs when they collide with static discs. Dotted green lines depict AUD event timings. (C–D) As above, but for trials containing Vad events. The trajectory of Vad is represented by dotted red lines. See main text for further descriptions.

AUD and Vstd events had been synchronous or asynchronous. In the Test condition, an additional VIS event (Vad) was presented. These either preceded or succeeded Vstd events by 75 ms. When Vad was presented, AUD events were either synchronous with Vstd (Figure 3C) or with the Vad (Figure 3D). In this case, participants reported which of the two alternative VIS events had been synchronous with the AUD event.

A run of trials consisted of 150 individual trials, 25 synchronous isolated AUD and Vstd events, 25 asynchronous isolated AUD and Vstd events, 25 asynchronous AUD and Vstd with preceding Vad, 25 asynchronous AUD and Vstd with succeeding Vad, and 50 synchronous AUD and Vstd with preceding or succeeding Vad. Each participant completed two runs of trials. Given the difficulty of the task, we provided feedback as to the accuracy of responses to maintain participant motivation and maximize performance.

Results and discussion

Individual scores were converted into measures of hit rate (HR), false alarm rate (FAR), and sensitivity (d') as per signal detection theory (SDT; Green & Swets, 1966; Macmillan & Creelman, 1991).

Performance in Baseline trials is indicative of sensitivity for discriminating between synchrony and asynchrony for isolated AV pairs. Performance in Test trials reflects the sensitivity for deciding which of two candidate VIS events is synchronous with AUD.

AV timing sensitivity was greater when there were two VIS events ($d' = 2.75$, $SEM = 0.31$) as opposed to just one ($d' = 0.99$, $SEM = 0.41$; $t_5 = 4.01$, $p = 0.01$, paired samples—two tailed). Thus, participants were better at deciding which of two VIS events had coincided with an AUD event than they were at deciding whether or not isolated VIS and AUD events had been coincident. These

results show that an objective measure of AV timing sensitivity can be *enhanced* by the presence of an additional VIS event.

In trials with additional VIS events, AUD was always synchronous with one or the other, creating two AV timing relationships. However, in trials without additional VIS events, AUD could precede, succeed, or be synchronous with VIS. It was therefore possible that the apparent AV timing sensitivity improvement in [Experiment 2](#) was a consequence of contrasting a task with just two AV timing relationships to one containing three. However, we obtained a qualitatively similar data set when we changed the later task, such that AUD either preceded or was synchronous with VIS (see [Supplementary experiment](#)). The apparent improvement in AV timing sensitivity must therefore arise because the presence of an additional VIS event provides information that makes timing judgments more precise.

General discussion

We have shown that the apparent timing of sensory events can be modulated by the presence of additional events located within the classically defined AV simultaneity window ([Experiments 1A](#) and [1B](#)). These interactions can make AV pairs seem *asynchronous*, in the presence of an additional sensory event, at timing relationships at which they had seemed synchronous when presented in isolation. As such, these data suggest that rather than representing a fixed interval during which detection of AV *asynchrony* is impossible, the AV simultaneity window is dynamic, shaped by interactions between temporally proximate sensory events.

The presence of these interactions could be indicative of a process of sensory segregation across time. Such a process would only be functionally advantageous if it improves timing sensitivity. The results of [Experiment 2](#) are consistent with this. We found that AV timing sensitivity can be *enhanced* by the presence of additional sensory events, relative to when AV events are presented in isolation.

Our data show that sensory events interact, shaping apparent timing. For instance, if a VIS event seems synchronous with an AUD event, other offset VIS events will not. This is true even if the offset of the other VIS event is small, such that it would seem synchronous with the AUD event if presented in isolation. Most importantly, our data show that the interactions involving additional sensory events do not just impact on response criteria but actually improve the accuracy of timing judgments. This interaction between temporally proximate events dictates that timing sensitivity in settings with multiple sensory events cannot simply be predicted on the basis of data obtained with isolated sensory events.

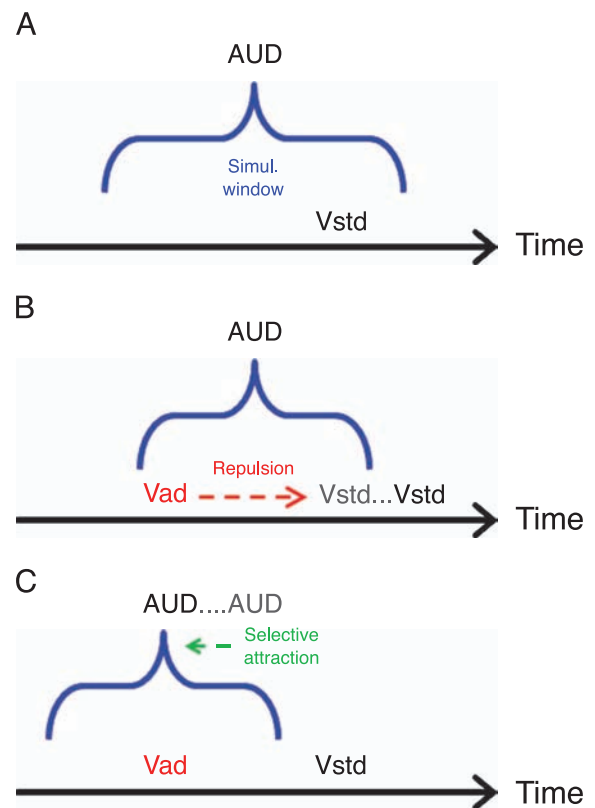


Figure 4. Graphical depictions of plausible explanations for our timing shifts. (A) When only one VIS (V_{std}) and one AUD event are present, the two events seem synchronous across a range of physical offsets (the Simultaneity window, depicted in blue). (B–C) When an additional VIS event is present (V_{ad}), two scenarios are possible. (B) The presence of V_{ad} may repel the apparent timing of V_{std} , possibly shifting it out of the AUD Simultaneity window. Here depictions in bold represent apparent timings, while physical timings are in gray. (C) Alternatively, AUD may be selectively attracted toward the timing of V_{ad} , again placing V_{std} beyond the range of the AUD Simultaneity window.

The shifts of AV timing perception in this study could have at least two causes. As depicted in [Figures 4A](#) and [4B](#), the shifts could have been brought about by a temporal repulsion between intra-modal sensory signals. Accordingly, the presence of an additional sensory event would *repel* the apparent timing of other proximate events in the same sensory modality, resulting in the offset intra-modal signal being judged as synchronous with delayed cross-modal signals.

Another possibility is that the shifts are brought about by competitive interactions, wherein a cross-modal signal is brought into apparent temporal correspondence with just one of two possible events. This is our favored hypothesis. It implies a *selective* temporal ventriloquism. This does not necessitate changes to the times at which perceptual information becomes available. For instance, speeded manual responses, triggered by event detections, might be unaffected as perceptual detection times are unchanged

(see Nishida & Johnston, 2002). Instead, when more than one intra-modal event is present, cross-modal signals might be grouped via a *selective attraction* (as illustrated in Figures 4A and 4C) and determined on the basis of relative temporal proximity.

Our data place an important constraint on plausible explanations for our apparent timing shifts. Any explanation of our data must involve a shift of the classically defined AV simultaneity window. As can be seen in Figures 2B and 2C, AUD and Vstd synchrony was reported more frequently ~ 250 ms *after* preceding Vad events, and ~ 250 ms *before* succeeding Vad events, relative to when AUD and Vstd events were presented in isolation.

It has been suggested that the shifts in PSS reported in Experiments 1A and 1B could be accounted for by a Bayesian estimation of cross-modal event timing in relation to the two intra-modal signals, with the prior distributions of AV signal timing acquired during a run of trials. (see Miyazaki, Yamamoto, Uchida, & Kitazawa, 2006). However, an explanation based on a process of Bayesian estimation cannot account for our observed improvements in timing sensitivity. In Experiment 2 we found that participants could reliably discern the physically synchronous AV pairing when the candidate events were offset by just 75 ms. However, participants could not reliably determine whether an isolated AV pairing was synchronous or asynchronous at the same physical offset. This increase in temporal sensitivity is also evident in the subjective timing data of Experiments 1A and 1B. Here the distributions of Vstd and AUD synchrony reports were narrowed in the presence of an additional sensory event (see Figures 2B and 2C). These observations place a further constraint on any plausible interpretation of our data.

Our data demonstrate that timing judgments, concerning pairs of AUD and VIS events, are not independent of other temporally proximate AUD or VIS events. Instead, there is an interaction that can enhance the precision of timing judgments. We suggest that this is due to competitive cross-modal temporal grouping. We propose that the temporally most proximate cross-modal pair are perceptually grouped, making them seem synchronous. As a consequence, the grouped events seem *asynchronous* with other sensory events, particularly when the other sensory event is clearly offset from one of the two grouped events, as was the case in our experiments. This proposition is consistent with previous findings showing that the grouping of signals *within* a sensory modality can weaken grouping *across* sensory modalities (Keetels, Stekelenburg, & Vroomen, 2007). Similarly, our results demonstrate that the grouping of one cross-modal pair can impair alternate groupings. We acknowledge that this interpretation is speculative, and will be looking to verify it in future experiments.

The perceptual events in our experiments were highly abstract, allowing for tight control in an experimental setting. However, we anticipate that the processes identified

will be relevant in real-world settings. For instance, making unrelated sensory events seem less related might make it easier to follow a specific conversation in a crowded social setting. However, in real-world settings, these sensory interactions are likely to be modulated by prior experience and expectation (Guski & Troje, 2003; Jackson, 1953; Radeau & Bertelson, 1987; Spence, 2007; Vatakis & Spence, 2007; Welch & Warren, 1980) rather than being driven solely by the physical timings of different sensory events. These suggestions could be tested in experiments using more naturalistic stimuli.

Conclusions

Previous studies of AV timing perception have suggested the existence of an AV simultaneity window—an interval during which AUD and VIS signals are necessarily perceived as synchronous (Dixon & Spitz, 1980; Lewald et al., 2001; Lewald & Guski, 2003). Our data show that the AV simultaneity window is not fixed, but dynamic, shaped by interactions between temporally proximate sensory events. These interactions can make temporally offset AUD and VIS events seem less related when presented in temporal proximity to other events than they would be if presented in isolation—an apparent perceptual *segregation* across time. These interactions can enhance the objective sensitivity of timing judgments and are likely to be common in the cluttered environments in which humans exist.

Acknowledgments

We would like to thank Tom Wallis and Waka Fujisaki for comments and discussions during the course of this project. This research was supported by an Australian Research Council discovery project grant and fellowship awarded to DHA.

Commercial relationships: none.

Corresponding author: Warrick Roseboom.

Email: roseboom@psy.uq.edu.au.

Address: School of Psychology, The University of Queensland, St. Lucia, QLD, 4072, Australia.

References

- Dixon, N. F., & Spitz, L. (1980). The detection of auditory visual desynchrony. *Perception*, *9*, 719–721. [PubMed]
- Fendrich, R., & Corballis, P. M. (2001). The temporal cross-capture of audition and vision. *Perception & Psychophysics*, *63*, 719–725. [PubMed] [Article]

- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Guski, R., & Troje, N. F. (2003). Audiovisual phenomenal causality. *Perception & Psychophysics*, *65*, 789–800. [[PubMed](#)] [[Article](#)]
- Jackson, C. V. (1953). Visual factors in auditory localization. *Quarterly Journal of Experimental Psychology*, *5*, 52–65.
- Keetels, M., Stekelenburg, J., & Vroomen, J. (2007). Auditory grouping occurs prior to intersensory pairing: Evidence from temporal ventriloquism. *Experimental Brain Research*, *180*, 449–456. [[PubMed](#)]
- Lewald, J., Ehrenstein, W. H., & Guski, R. (2001). Spatio-temporal constraints for auditory–visual integration. *Behavioural Brain Research*, *121*, 69–79. [[PubMed](#)]
- Lewald, J., & Guski, R. (2003). Cross-modal perceptual integration of spatially and temporally disparate auditory and visual stimuli. *Cognitive Brain Research*, *16*, 468–478. [[PubMed](#)]
- Macmillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide*. Cambridge: Cambridge University Press.
- Miyazaki, M., Yamamoto, S., Uchida, S., & Kitazawa, S. (2006). Bayesian calibration of simultaneity in tactile temporal order judgement. *Nature Neuroscience*, *9*, 875–877. [[PubMed](#)]
- Morein-Zamir, S., Soto-Faraco, S., & Kingstone, A. (2003). Auditory capture of vision: Examining temporal ventriloquism. *Cognitive Brain Research*, *17*, 154–163. [[PubMed](#)]
- Nishida, S., & Johnston, A. (2002). Marker correspondence, not processing latency, determines temporal binding of visual attributes. *Current Biology*, *5*, 359–368. [[PubMed](#)]
- Radeau, M., & Bertelson, P. (1987). Auditory–visual interaction and the timing of inputs. Thomas (1941) revisited. *Psychological Research*, *49*, 17–22. [[PubMed](#)]
- Spence, C. (2007). Audiovisual multisensory integration. *Acoustical Science and Technology*, *28*, 61–70.
- Spence, C., & Squire, S. B. (2003). Multisensory integration: Maintaining the perception of synchrony. *Current Biology*, *13*, R519–R521. [[PubMed](#)]
- Vatakis, A., & Spence, C. (2007). Crossmodal binding: Evaluating the “unity assumption” using audiovisual speech stimuli. *Perception & Psychophysics*, *69*, 744–756. [[PubMed](#)] [[Article](#)]
- Vroomen, J., & Keetels, M. (2006). The spatial constraint in intersensory pairing: No role in temporal ventriloquism. *Journal of Experimental Psychology: Human Perception and Performance*, *32*, 1063–1071. [[PubMed](#)]
- Welch, R. B., & Warren, D. H. (1980). Immediate perceptual response to intersensory discrepancy. *Psychological Bulletin*, *88*, 638–667. [[PubMed](#)]