Neuron 51, 651–659, September 7, 2006 ©2006 Elsevier Inc. DOI 10.1016/j.neuron.2006.08.006

Motor-Sensory Recalibration Leads to an Illusory Reversal of Action and Sensation

Chess Stetson,[1,7] Xu Cui,[2,7] P. Read Montague,[3,4] and David M. Eagleman[1,5,6,*]
[1] Department of Neurobiology and Anatomy
University of Texas Medical School
6431 Fannin Street, Suite 7046
Houston, Texas 77030
[2] Program in Structural and Computational Biology and Molecular Biophysics
[3] Department of Neuroscience
[4] Menninger Department of Psychiatry and Behavioral Sciences
Baylor College of Medicine
1 Baylor Plaza
Houston, Texas 77030
[5] Departments of Biomedical Engineering and Psychology
University of Texas at Austin
Austin, Texas 77003
[6] Department of Psychology
Rice University
Houston, Texas 77005

Summary

To judge causality, organisms must determine the temporal order of their actions and sensations. However, this judgment may be confounded by changing delays in sensory pathways, suggesting the need for dynamic temporal recalibration. To test for such a mechanism, we artificially injected a fixed delay between participants' actions (keypresses) and subsequent sensations (flashes). After participants adapted to this delay, flashes at unexpectedly short delays after the keypress were often perceived as occurring before the keypress, demonstrating a recalibration of motor-sensory temporal order judgments. When participants experienced illusory reversals, fMRI BOLD signals increased in anterior cingulate cortex/medial frontal cortex (ACC/MFC), a brain region previously implicated in conflict monitoring. This illusion-specific activation suggests that the brain maintains not only a recalibrated representation of timing, but also a less-plastic representation against which to compare it.

Introduction

Correctly judging the order of action and sensation is essential for determining causality, an assessment fundamental for learning and survival. Imagine you are walking through a forest and hear a twig crack. Did it happen when your foot fell or just before? If it happened just before, the sound may alert you to a nearby predator. If the sound happened coincident with your step, then it was a normal occurrence consistent with the sensory feedback expected during walking. However, the ability to correctly judge temporal order may be confounded by changing delays in sensory pathways—e.g., due to retinal response times in different lighting conditions (Purpura et al., 1990) or, on longer time scales, due to limb growth (Campbell et al., 1981). Accordingly, the nervous system must be able to recalibrate its expectations about the normal temporal relationship between actions and sensations to overcome changing latencies and correctly determine causality.

Previous work has reported that a delayed sensory effect is judged to appear slightly earlier in time if it follows a voluntary action (Haggard et al., 2002; Eagleman and Holcombe, 2002). This effect could result from a compression of time between actions and subsequent sensory events, such that all sensations following actions appear to draw closer to the actions (intentional binding). An alternative hypothesis, which we explore here, suggests that the perceived timing of sensory events *shifts* with respect to the perceived timing of actions (Figure 1A). In this alternative hypothesis, sensory events appearing at a consistent delay after motor actions are interpreted as consequences of those actions, and the brain recalibrates timing judgments to make them consistent with a prior expectation that sensory feedback will follow motor actions without delay. This alternative hypothesis predicts a novel illusion: participants will perceive sensory events as occurring before their actions when in reality the events occurred afterward (Figure 1A). A similar effect of order reversal has once been reported anecdotally (Cunningham et al., 2001), but was never verified or quantified. To look for evidence of motor-sensory time recalibration, we designed a psychophysical experiment in which we probed participants' perceptions after they adapted to an injected delay between their actions and sensory feedback.

While much effort has recently focused on finding brain areas involved in time *estimation* using fMRI (Bushara et al., 2001; Coull et al., 2004; Pouthas et al., 2005; Jech et al., 2005; Lewis and Miall, 2006), no experiment has yet looked for neural signature of temporal *recalibration*. There are at least two hypotheses about what might be found using imaging. First, when exposed to a consistent delay, the brain may simply adjust a setpoint in such a way that a contrast between its activity in the adapted and unadapted cases would show no difference. On the other hand, if different activity is found before and after calibration, this could indicate the existence of latent (baseline) representations of time which do not adapt to the brief exposure and which conflict with the quickly adapting representations. That is, if multiple representations of temporal order exist in the brain, we would expect to see signatures of neural conflict during those times when the representations disagree (i.e., one representation concludes that A came before B, the other that B came before A). To distinguish these hypotheses of single versus multiple representations of temporal order, we seek the neural signatures of motor-sensory temporal recalibration using fMRI.

*Correspondence: david@eaglemanlab.net
[7] These authors contributed equally to this work.

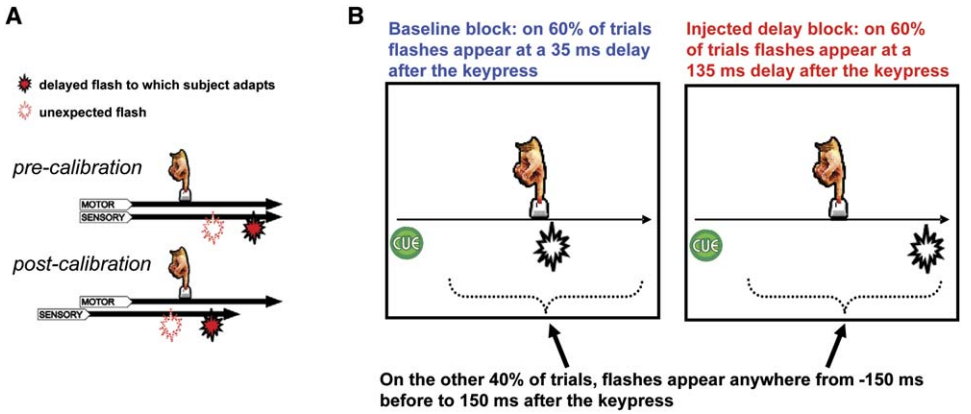

Figure 1. Does Adaptation to a Delay Induce a Shift in Motor-Sensory Temporal Order Judgments?

(A) Schematic of hypothesis: participants exposed to delayed sensory feedback (filled flash) calibrate temporal order judgments to reduce the delay between motor output and sensory feedback. After recalibration, the delayed feedback appears closer in time to the keypress, erroneously causing an unexpected flash appearing at a sooner time (hollow flash) to be perceived as occurring before the motor act. Only one flash appears on any given trial.

(B) Task design: participants are cued to press a key. A flash appears on the screen after the cue and sometime before or after the keypress. On 60% of trials, the flash appears at a fixed time with respect to the keypress (35 ms afterward in the baseline block and 135 ms afterward in the injected delay block). On the other 40% of trials, the flash appears at an unexpected time somewhere before or after the keypress. At the end of every trial, participants report whether the flash appeared before or after their keypress.

Results

Effect of Motor-Sensory Adaptation on Temporal Order Judgments

Experiment 1: Illusory Reversal of Temporal Order

To look for evidence of motor-sensory time recalibration, we designed a temporal order experiment in which we manipulated the relationships between motor acts and their sensory consequences. On most trials, we injected a fixed 135 ms delay between participants' keypress and a subsequent flash. On the other trials, the flash appeared at a variable delay either before or after the participant's keypress. After every trial, we asked participants to report their perceptions of the temporal order of flash and keypress. We refer to this block of trails as the "injected delay block." We ran a second, similar test in which the most frequent trials appeared at a fixed delay of 35 ms (the minimum delay possible on our computer system), which we call the "baseline block" (Figure 1B; further details, including our method for presenting the flash before the keypress, can be found in Experimental Procedures).

For each participant, the point of subjective simultaneity (PSS) was taken as the keypress-flash time difference at which he reported "keypress before flash" at 50% probability (Figure 2). We compared each participant's PSS from the baseline and injected delay blocks. Twenty-four of twenty-five participants' psychometric functions shifted in the positive direction (average PSS shift 44 ± 7 ms, standard error of the mean, $p < 10^{-6}$; Figure 3A). Thus, a flash appearing in a 44 ms window after the keypress would usually elicit an "after" report in the baseline block but elicit a "before" report in the injected delay block. This is the first quantification of an illusory reversal of temporal order of action and sensation. The slopes of the curves (reflecting the precision of the temporal order judgment) did not differ significantly in the two blocks (two-tailed paired t test, $p < 0.43$; Figure 3A, first column), suggesting no change in task difficulty.

Experiment 2: Test for Motor-Sensory versus Cross-Sensory Recalibration

To address the possibility that the effect we have shown reflects a cross-sensory, rather than a motor-sensory recalibration, we repeated the experiment with the following variation: the key automatically moved up to tap the participant's finger instead of the other way around, and participants judged the order of the tap and a flash. In this cross-sensory condition, we found a small shift in the PSS between the baseline and injected

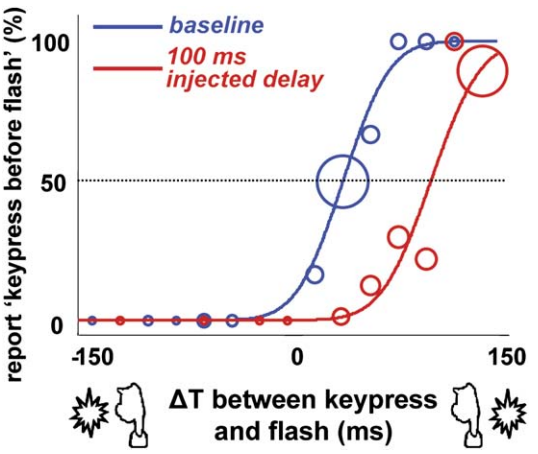

Figure 2. Reversal of Perceived Order of Actions and Subsequent Events

Temporal order experiment data from a representative participant in the baseline (blue) and injected delay (red) blocks. Dot size reflects the number of trials at each delay between key and flash times. Curves are logistic distribution functions. The dashed line intersects the 50% point of each curve, referred to as the point of subjective simultaneity (PSS).

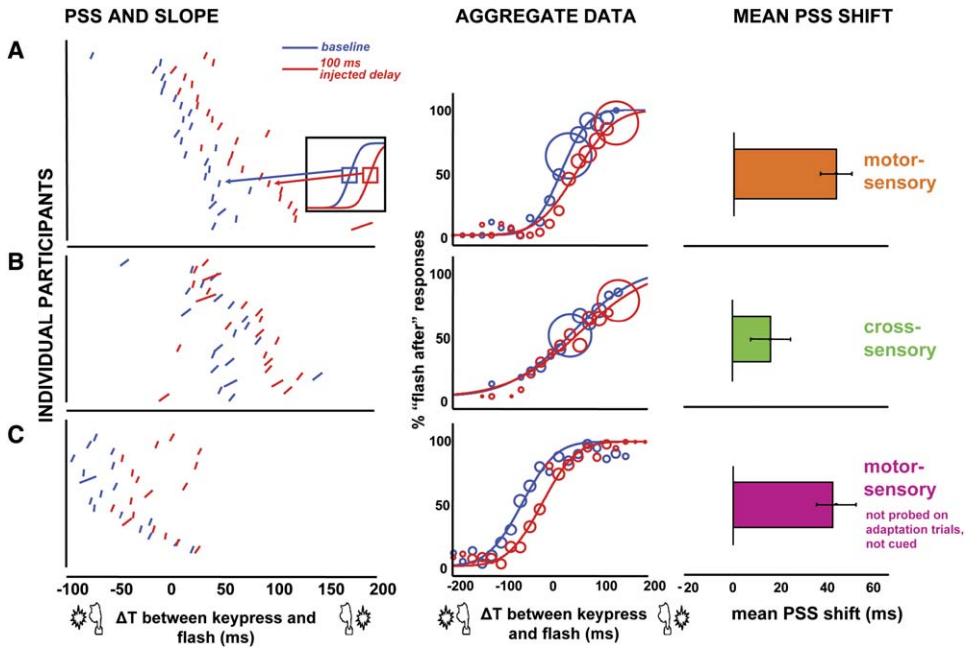

Figure 3. Motor-Sensory, Cross-Sensory, and Uncued Motor-Sensory Data

(A) Motor-sensory experiment from Figures 1 and 2. For each subject, the PSS for the baseline (blue) and 100 ms injected delay (red) blocks are represented by the ticks. The angle of a tick corresponds to the slope of the psychometric curve, and the position of the tick on the x axis corresponds to the mean of the curve. The slopes in the injected delay block and in the baseline block were not significantly different (two-tailed t test, $p < 0.43$). The second column shows the collected data from all participants in each test, plotted as a single psychometric function. As in Figure 2, the area inside the circles reflects the number of trials at each offset. The third column shows the mean PSS shift with SEM.

(B) Cross-sensory condition, in which the key taps the participant's finger rather than the other way around.

(C) Uncued motor-sensory condition in which participants' perceptions were not probed on trials during which they adapted to the injected delay (60% of trials); instead, participants only made temporal order judgments on the remaining 40% of trials. This procedure equalized the before/after distribution of trial times in which participants were asked to make perceptual reports in the baseline and injected-delay trial blocks. The motor-sensory shifts in (A) and (C) are significantly larger than the cross-sensory shift ($p < 0.02$, $p < 0.04$, respectively, two-tailed t test). Parenthetically, an interesting aspect of the uncued data is the −45 ms offset of the baseline curve, in contrast to the +20 ms offset of the baseline curve in the cued motor-sensory experiment (compare [A] to [C]). This difference may be attributable to the rhythm adopted by participants in the uncued design (Paillard 1948).

delay blocks (16 ± 8 ms, Figure 3B). This cross-sensory shift, though not quite significant in our study (two-tailed t test, $p = 0.06$), is consistent with other reports of small crossmodal recalibration (Fujisaki et al., 2004; Vroomen et al., 2004). The magnitude of these shifts is less than half of the motor-sensory shift of 44 ms, suggesting that active interaction with the world is a powerful mechanism for calibrating timing judgments.

Experiment 3: Control for Cognitive Bias Toward Central Tendency

The cross-sensory data in the previous experiment provides a control against a cognitive bias such as central tendency as an explanation for our motor-sensory data. A central tendency argument would suggest that since 60% of trials appeared at a long delay in the injected-delay block (making the "after" answer more frequent), participants might bias uncertain answers toward "before" to equalize the number of each kind of answer. However, since we would expect this bias to operate just as much in the cross-sensory case as in the motor-sensory case, the significantly larger motor-sensory shift is not suggestive of central tendency.

As a further control against central tendency, we designed a new motor-sensory experiment, in which participants only made temporal order judgments on a fraction of the trials. They were exposed to a fixed delay on 60% of the trials, but they were only probed to make judgments on the remaining 40%. In these trials, the before/after probability was evenly distributed. Thus, the distribution of times between keypress and flash on trials in which subjects reported their perceptions was identical between the baseline and injected-delay blocks.

To additionally ensure that the motor-sensory shift had nothing to do with the cued design, participants were not cued to make their actions but instead made keypresses at a time of their own choosing (see Experimental Procedures for details). The data from this uncued motor-sensory experiment can be seen in Figure 3C. The PSS shift in this version of the experiment is almost identical to the original motor-sensory shift (41 ms ± 8) and significantly larger than the cross-sensory shift ($p < 0.04$; Figure 3B). Given these data, it is unlikely that central tendency can account for the perceptual shift. Additionally, because participants were exposed to delays using red, yellow, or white flashes and then asked to make a temporal judgment about a blue flash (see Experimental Procedures), the result in Figure 3C demonstrates that the motor-visual recalibration is not sensitive to color.

As a framework for understanding these results, we suggest that sensory events appearing at a consistent

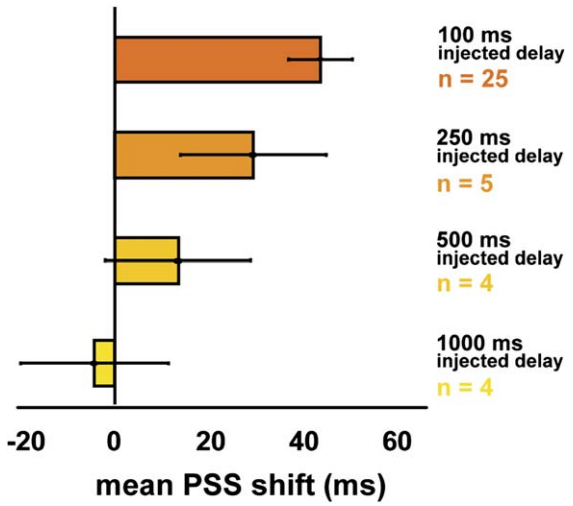

Figure 4. Average PSS Shift for Motor-Sensory Experiments with Injected Delays of 100, 250, 500 and 1000 ms

Bars are significant to $p < 1E\text{-}6$, $p < 0.13$, $p < 0.45$, $p < 0.80$, respectively, two-tailed t test. Error bars are standard error of the mean.

delay after motor actions are interpreted as consequences of those actions, and the brain recalibrates timing judgments to make them consistent with a prior expectation that sensory feedback will follow motor actions without delay. Having recalibrated, the brain can interpret sensory signals appearing at an earlier time than the expected feedback as preceding the motor action.

Experiment 4: Recalibration to Larger Delays

We hypothesize that if the recalibration depends on a neural interpretation of causality, the size of the motor-sensory recalibration should be a function of the size of the injected delay. Specifically, judgments of causality should be modulated by an expectation that very long delays do not represent the consequences of a voluntary action (Eagleman and Holcombe, 2002). To test this hypothesis, we repeated the original experiment using injected delays of 250, 500, and 1000 ms. Participants' PSS shifted by an average of 30 ± 16, 13 ± 16, and −4 ± 16 ms from the baseline to the 250, 500, and 1000 ms injected delay blocks, respectively (Figure 4), demonstrating that the effect diminishes as the motor act and sensory feedback are separated by greater delays.

Experiment 5: Neural Correlates of the Motor-Sensory Illusion of Temporal Order

There has been an extensive search for a single, central clock in the brain, which times events as they come in through the sensory organs and constructs a single temporal representation of the outside world (Nobre and O'Reilly 2004; Mangels et al., 1998; Ivry and Spencer, 2004; Buhusi and Meck, 2005). The recalibration of temporal order we have discovered gives us an opportunity to address the assumption of a single internal clock (Karmarkar and Buonomano, 2003; Nobre and O'Reilly 2004). While participants' responses indicate an internal representation of time which has recalibrated, it may be that there are other internal representations of time which have not recalibrated. If this is the case, we might expect the mismatch to elicit brain activity. To address this hypothesis, we repeated our baseline and injected delay tests while subjects were scanned using functional magnetic resonance imaging (fMRI). While these tests are identical perceptually, the fMRI data might reveal a difference between them neurally. If any area of the brain shows a different activation on trials featuring illusory reversal than it does on other types of trials, it suggests the existence of multiple timing representations in the brain, arguing against the single-clock theory.

ACC/MFC Is Activated by the Illusion

As a way of identifying a neural region-of-interest which might be maximally activated by the illusion, we produced two contrast images. The first contrast was made between "illusion trials" (trials in the injected delay block where the flash came after the baseline PSS but the participant reported that it happened "before") and veridical "before" trials (trials in the baseline block in which the flash appeared before the baseline PSS and participants reported the flash to come "before"). In these two trial types, the stimuli were different, but the perception was the same. The second image contrasted the illusion trials with veridical "after" trials (trials in the baseline block in which the flash appeared after the baseline PSS and participants perceived "after"). Between the latter two trial types, the stimuli were the same, but the perception was different. We reasoned that the voxels common to these two contrasts would give us a good candidate for a region maximally activated by the illusions. The voxels common to both images lie at the interface of the anterior cingulate cortex and medial frontal cortex (ACC/MFC, Figure 5A, peak at MNI, 0, 16, 42; false detection rate [FDR] corrected at $q < 0.05$ with a minimum of five contiguous voxels for each contrast).

Task Difficulty Does Not Explain ACC/MFC Activation

To address whether the ACC/MFC was maximally activated by the illusory trials (as compared to other trial types), we needed to correct for trial difficulty. To ensure that the trials analyzed were all of approximately equal difficulty, we chose only those trials in which the flash appeared close in time to the baseline PSS (within ±35 ms, see Difficulty Matching in Experimental Procedures). We grouped these trials by the physical order of the flash with respect to the baseline PSS and by the participants' perceptions on each trial.

The average hemodynamic response on illusory trials was significantly higher than on any other type of trial (Figure 5B; $p < 0.05$, one-tailed t test). Note that the ACC/MFC is more active during illusions than on trials in which participants veridically perceived the flash to occur before their keypress, so the increased activation cannot be attributed solely to the perception of being "beaten by the flash" (Figure 5C, top panel). Further, the ACC/MFC is activated more by illusion trials than by trials with the same delays in which participants correctly perceive the flash afterward, so we cannot attribute the high activation to something about the physical timing of the flash (Figure 5C, second panel). The increased ACC/MFC activity is unlikely to represent increased task difficulty, because the trials chosen for this plot are all from the same range of maximal difficulty

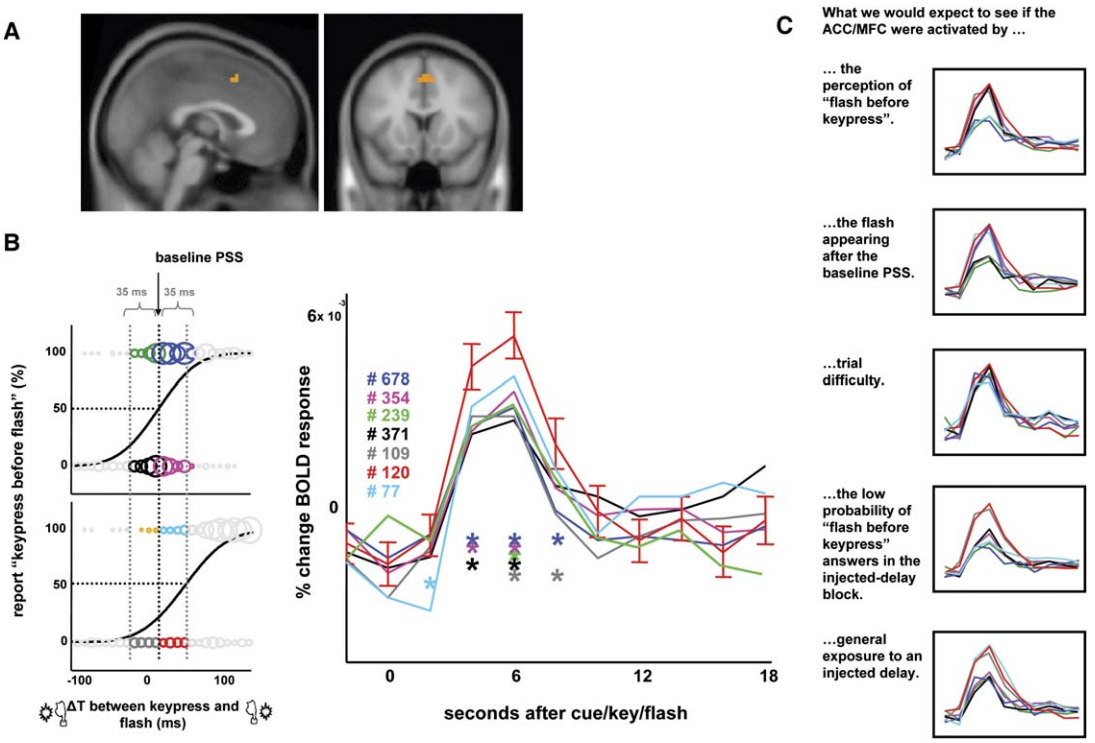

Figure 5. Illusion Maximally Activates ACC/MFC

(A) Region of interest formed by the voxels shared by two contrasts between illusory trials and veridical trials (see text).

(B) Averaged time courses of the BOLD signals in the voxels in (A). Only trials within 35 ms of the baseline PSS were analyzed in order to match for difficulty between the two blocks. Each curve corresponds to trials indicated by the same color in the inset (left). For example, illusory trials (red) are those in the injected delay block in which the flash follows the keypress but participants report "flash before keypress." The number of trials included appears above the traces, and error bars are standard error of the mean. *Significance value of $p < 0.05$ (two-tailed t test) between the illusion data (red) and the data with the color of the star. There were too few orange trials for inclusion.

(C) Schematic: what would be expected if the ACC/MFC activation represented the perception of the flash before the keypress, the flash appearance after the baseline PSS, the difficulty of the trial, the flash appearance as a low probability event ("oddball") in the injected delay block, or a general effect of being exposed to an injected delay.

in the participants' behavioral curves (Figure 5C, third panel). The ACC is known to be activated by unexpected stimuli ("oddballs," Linden et al., 1999), thus we might expect that it is activated by the "before" answer in the injected delay block simply because of that answer's infrequency. However, if this were true, one would also expect veridical "before" answers to elicit high activation (gray trials, Figure 5C, fourth panel), which does not happen. Finally, if the ACC/MFC activation resulted from the recalibration itself or reflected nothing more than a new learned relationship between motor timing and sensory feedback, one might expect it to persist throughout the entire injected delay block (Figure 5C, fifth panel). Yet the ACC/MFC signal is not elevated on trials when the perception in the injected delay block agrees with what it would have been in the baseline block (e.g., compare the gray and black curves in Figure 5B). In summary, if the ACC/MFC were maximally activated by anything other than the illusory trials, the data in Figure 5B would look different.

Neural Signal Reflects Amount of Recalibration

We have established that the ACC/MFC is maximally activated during trials when participants experienced an order-reversal illusion. Since some participants recalibrated more than others, we reasoned that ACC/MFC activity might reflect the range of the illusory reports. To explore this, we separately analyzed participants whose PSS shift was less than 33 ms (the "small adaptors"), and those whose shift of more than 33 ms (the "big adaptors"; Figure 6).

The peak ACC/MFC BOLD response from each trial were binned by the time between keypress and flash (ΔT) of the trial (50 ms bins, peak times were taken 6 s after the trial start). For each bin, we tested for a significant difference between the peaks of the activity in the injected delay block and in the baseline block. For the big adaptors, the difference between the ACC/MFC activation in the injected delay blocks was significant ($p < 0.05$, two-sample t test) when ΔT was between 16 and 66 ms, which is exactly where we see the maximum behavioral difference. For small adaptors, we did not find a significant difference in any range, nor any bin width.

Discussion

We have found that events which consistently lag actions can lead to a recalibration of temporal order judgments. This suggests that temporal order judgments between motor acts and sensory events dynamically change in order to keep causality assessments

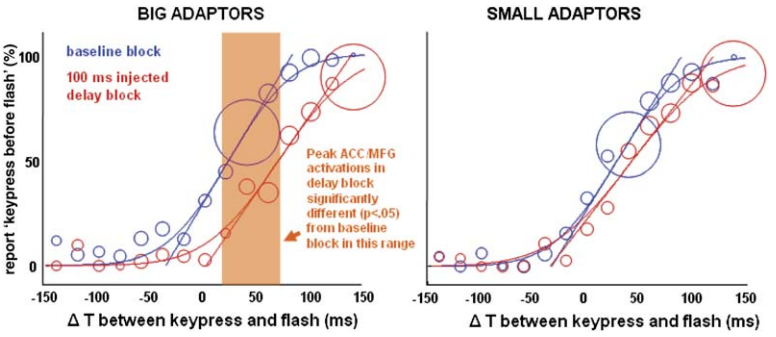

Figure 6. Maximal ACC/MFC Activation Overlaps Behavioral Shift of Big and Small Adaptors

When comparing baseline to injected blocks, the ACC/MFC activation difference is most significant when the behavioral difference is maximal. We separated participants into those whose PSS shift was larger than 33 ms (the big adaptors) and those with a positive shift smaller than 33 ms (the small adaptors). For each ΔT between keypress and flash, we collected the peak ACC/MFC activations for each trial (in 50 ms bins centered around the baseline PSS). For big adaptors, the ACC/MFC activation in trials where ΔT was between 16 and 66 ms was significantly greater in the 100 ms injected delay block than in the adapted block. This is exactly the range of the maximal behavioral difference for the big adaptors, whose mean PSS shifted from 16 to 57 ± 6 ms. We failed to find a significant difference in ACC/MFC activation for range of ΔT for the small adaptors. This provides further evidence that the ACC/MFC activation correlates with illusions of temporal order resulting from recalibration.

appropriately calibrated. This recalibration causes sensory inputs which occurred just after motor actions to appear as though they came beforehand—an illusory reversal of temporal order (Figure 1A). We suggest that the calibration of motor-sensory timing results from a prior expectation of little or no delay between outgoing actions and resulting sensory effects. The illusory reversals coincide with maximal ACC/MFC activation, suggesting the existence of at least two representations of temporal order in the brain—one which rapidly adapts to the injected delay and one which retains its baseline timing judgments.

With regard to the temporal order illusion, an open question is why only a few participants showed a PSS shift equal to the size of the injected delay. We attempted to increase the size of the shift by exposing participants to 75 trials at the adaptation delay before the test began, but that had little effect. By pooling all of the participants' behavioral data and examining a running average of the PSS shift, we noticed that the recalibration appears to reach its full magnitude within ~20 trials (data not shown). It may be that motor-sensory timing shifts of 100 ms are beyond the hardware limitations of the calibration mechanisms. Another possibility is that adaptation to injected delays battles the existing calibration cemented in by years of experience, such that a much longer exposure to the injected delay would be required to increase the effect.

Differential Neural Responses Suggest the Possibility of Multiple Timing Representations

We have found that a small brain region at the border of the ACC and MFC is maximally activated when participants, having adapted to an injected delay, experience an illusion that a flash preceded their keypress. We suggest that for a region of the brain to be maximally activated exclusively by illusion trials, the brain must contain representations of both the recalibrated and the baseline timing. Our reasoning is as follows: the PSS shift in the behavioral data shows that a timing representation exists (namely, the one which determines the perceptual report) which adapts to the injected delay. In order for a region of the brain to become maximally activated only on those trials when the recalibrated representation disagrees with what the baseline representation would have reported, the brain must have a latent or stored version of the baseline representation. We find such an activation in a region that has been implicated in conflict monitoring (Botvinick et al., 2001; Holroyd et al., 2004).

The ACC is activated in other fMRI experiments which have employed the Stroop task (Pardo et al., 1990; Leung et al., 2000) and Flanker task (Botvinick et al., 1999—tasks in which sensory input is thought to activate conflicting neural representations (Figure 7). Neighboring areas in the medial frontal cortex have also been implicated in conflict (Nachev et al., 2005). Our results may be consistent with the above studies, in that ACC/MFC activation in our experiment seems to result from a mismatch between answers resulting from separate timing representations. Many theories of ACC function suggest that it selects behavioral actions (Paus, 2001), monitors conflict in support of behavioral adjustment (Kerns et al., 2004), assesses the consequences of a decision (Walton et al., 2004), predicts error likelihood (Brown and Braver, 2005; Carter et al., 1998), or signals

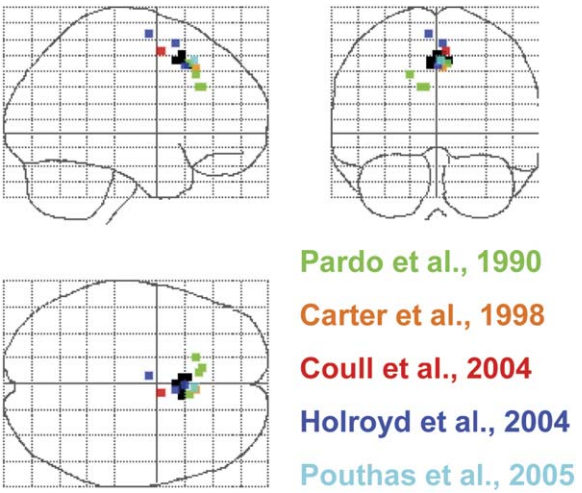

Figure 7. ACC/MFC in the Conflict and Timing Literatures

Coordinates of peak activation found in five papers overlaid onto the ACC/MFC activation shown in Figure 5 (black).

the difference between predicted and actual consequences of an action (Ito et al., 2003). These theories stress that the ACC is involved in the pursuit of reward. Our task involves no explicit reward or error signals, so it is difficult to say whether our data supports any of these theories. Note that if we expect that the ACC is activated by conflict, we would only expect it to be maximally activated at times when two timing representations do not agree (that is, during the illusory trials), rather than during the entire trial block.

Since ACC is involved in signaling conflict, we might expect it to become activated during other kinds of perceptual illusions, and indeed it does—for example, in binocular rivalry (Lumer et al., 1998; Cosmelli et al., 2004). Note that not all perceptual illusions will be expected to activate the ACC, but only those in which multiple representations of the stimuli compete in the brain.

Some experiments have found activation of similar regions as ours in time-estimation tasks (Coull et al., 2004; Lau et al., 2004; Eagleman, 2004; Pouthas et al., 2005; Rao et al., 2001), while others have not (Nenadic et al., 2003). Looking at all seven hemodynamic response functions in Figure 5B, it appears that while the ACC/MFC is preferentially activated by illusory trials, it is somewhat activated by all trials in the task. It is difficult to know whether the baseline ACC/MFC activation on all trials in this study represents time estimation, as other studies might suggest, a basal level of conflict, or some other kind of activity. It will be interesting to explore how ACC/MFC activity relates to Insular activity, which has been shown to encode the perception of agency (Farrer and Frith, 2002).

Some theories of the ACC suggest that it is involved in conflict resolution (Botvinick et al., 2001). If this occurred in our experiment, we might expect the activation to go away after a sufficiently long period of adaptation; further experiments will be necessary to know whether this happens. Moreover, if the conflict resolution acted to resolve the discrepancy between timing representations by recalibrating the injected-delay PSS to its baseline value, we might expect the behavioral evidence of adaptation to disappear after the subject sees enough illusions. This will be explored in further studies.

Experimental Procedures

Behavioral Methods

Each trial block consisted of 100 trials and was run once per participant. On each trial, the participant pressed the button as quickly as he could in response to a cue. The software kept a running average of each participant's reaction time to the cue, making it possible to probabilistically place flashes just before the keypress. There were two randomly ordered trial blocks—the "baseline" and "injected delay" blocks. Flashes were distributed throughout a trial block such that 60% appeared at a consistent delay (35 ms in baseline block, 135 ms in injected delay block) and the rest were distributed in a Gaussian centered 60 ms after the keypress with a width of 80 ms—this distribution attempted to maximize the number of trials at the steep part of the psychometric functions. The distribution of these 40 trials was the approximately the same in both the baseline and injected delay blocks. For a sense of the distribution of flashes, refer to the sizes of the dots in Figure 2 and Figure 3, middle column. In those figures, the radius of the dots is proportional to the square root of the number of trials. Trials of different delays were randomly interleaved within a block, and the blocks were randomly ordered. On average, in the fMRI baseline block, 47% of the reports ended up being "flash before keypress" as opposed to 32% in the injected delay block (the probabilities were 38% and 26% outside the scanner).

In this manuscript, the PSS at a 35 ms delay (that is, during the baseline block) is called the baseline PSS. This is the marker against which the adapted PSS is compared, and by which participants are combined with each other in the fMRI analysis. We classify as "illusions" those trials in the injected delay block in which the flash appears after the baseline PSS, but the participant reports it to occur "before." Since we do not know how participants choose the marker against which to compare the time of the flash (e.g., the time at the first activity in motor cortex, the time when his finger touches the key, or the time when the key is fully depressed), the baseline PSS is the best measure of simultaneity we have for a given subject. It may be that what we call the baseline PSS (on average 20 ± 5 ms after the time that the computer registered a keypress) is actually representative of some adaptation to a 35 ms delay typical of computers. If true, our measure of the PSS shift may be diluted—i.e., a motor-sensory shift from a 0 ms delay condition would be even larger than what we have reported here.

Before running the experiment, participants were required to pass a training version of the task, in which flashes appeared randomly distributed before or after their keypress on all trials (i.e., there was no consistently presented delay). For feedback during this training, flashes were defined as "before" the keypress if they came anywhere up to 35 ms (the delay due to our graphics card) after the keypress and "after" if they came later than that. Participants gained one point for correctly identifying before/after (as defined above) and lost five points for incorrect answers. If they could accrue 25 points, we allowed them to participate in the experiment.

For the cross-sensory test (Figure 3B), the button lightly tapped the participant (the button was attached to a driver which provided an ~10 ms tap to the bottom of the participant's finger). This cross-sensory (or "involuntary") condition was designed such that the statistics of timing between key events and flashes was the same as in the motor-sensory experiment.

In the uncued motor-sensory experiment shown in Figure 3C, a participant's finger tap caused a red, white, or yellow light-emitting diode (led) to flash at a fixed delay (0 ms in the baseline block, 100 ms in the injected delay block). After 5–12 such finger taps, a blue flash appeared sometime just before or after the participant's next finger tap (our software attempted to predict the time of the participant's next finger tap from the loose rhythm participants tended to adopt and then presented the flash in a uniform random distribution between ±150 ms of the participant's keypress on each trial). Participants were instructed to recognize the trial with the blue flash as the "probe" trial and answered whether the blue flash had come before or after the time of their last finger tap. These probe trials were evenly distributed between "before" and "after" answers, and this was true in both the baseline and injected delay blocks.

Finally, a note on our methods—Libet's clock face paradigm has been used in the past to test questions of relative timing of action and effect (Haggard et al., 2002). In that paradigm, participants watch a rapidly moving clock hand and report the position of the hand at the moment they perceive some event (the "probe"). We have been concerned that this methodology may introduce a confound in the form of the flash-lag effect (Nijhawan, 1994; Eagleman and Sejnowski, 2000a), in which the position of a moving object is reported to be further along in its trajectory than its physical position at the time of probing. Importantly, the flash lag effect is modified by predictability (Eagleman and Sejnowski, 2000b, Namba and Baldo, 2004): the more predictable the time of the probe, the smaller the positional offset. Thus, it may be that beeps following a voluntary button press are more predictable than surprise beeps. If true, this complicates the interpretation of an earlier report by Haggard et al. (2002), in which the use of a Libet clock paradigm produced the result that beeps following 250 ms after a voluntary keypress were reported at an earlier time on the clock face. To avoid confounds, we designed our experiment so that flashes are equally unpredictable in all conditions.

fMRI Methods

Of 38 right-handed participants scanned, 19 pressed the key in response to the cue with their right hand and answered the "before/after" question with their left hand, while the other 19 performed the task with hands reversed. High-resolution T1-weighted scans

were acquired using an MPRage sequence (Siemens). Six to twelve seconds elapsed between the time of the cue (and flash and initial keypress) and the time at which the participant was asked to report her perceptions. For some participants, we reduced the 6–12 s delay to 2–4 s on some of the adaptation trials (the frequent trials appearing at 35 or 135 ms) in order to decrease the total time the participant spent in the scanner, but these trials were excluded from the time course plot in Figure 5B.

Functional run details were as follows: echo-planar imaging, gradient recalled echo; repetition time (TR) = 2000 ms; echo time (TE) = 40 ms; flip angle = 90°; 64 × 64 matrix, 29 4 mm axial slices, yielding functional 3.4 mm × 3.4 mm × 4.0 mm voxels. Data analysis was performed using SPM2 (http://www.fil.ion.ucl.ac.uk/spm/software/spm2) and visualized using xjView (http://people.hnl.bcm.tmc.edu/cuixu/xjView). Motion correction to the first functional scan was performed using a six parameter rigid-body transformation (Kao et al., 2005). The average of the motion-corrected images was coregistered to each individual's structural MRI using a 12 parameter affine transformation. The images were spatially normalized to the MNI template by applying a 12 parameter affine transformation, followed by a nonlinear warping using basis functions (Kao et al., 2005). Images were then smoothed with an 8 mm isotropic Gaussian kernel and highpass filtered in the temporal domain (filter width of 128s, King-Casas et al., 2005).

Fourteen participants in the behavioral tests also participated in the fMRI tests, and 11 participants from the right-handed set of fMRI tests were repeated in the left-handed version. The mean PSS shift in the fMRI scanner was 35 ± 5 ms. In order to find a region of interest, we performed a general-linear-model regression. For each block, we separated trials in which the flash preceded the keypress from those in which it followed the keypress. We further separated those trials after which the participant would report "flash before key" from those after which the participant would report "key before flash."

Using SPM2, we performed two separate multiple regressions, one for each trial block. To define the regressors, we identified "cue" trials as those in which the participants saw the cue and flash and pressed the key and "response" trials as those in which the participant reported his perception (usually about 6 s after each "cue" trial). We further separated the cue trials by whether the flash came before or after the baseline PSS and by whether the participant perceived the flash as before or after his keypress, leaving us with six total regressors. After performing the regressions, we formed two random-effect contrasts (a t test of differences in β values). We contrasted the illusory trials (trials in the injected-delay block when the flash appeared after the baseline PSS but the perception was "before") to the veridical trials (baseline block, flash before baseline PSS, perception "before"; baseline block, flash after baseline PSS, perception "after"). Both of these contrasts were required to pass a threshold of $p < 0.001$ with at least five continuous voxels. They were further required to pass a false discovery rate (FDR) test of $q < 0.05$ (Genovese et al., 2002). We used the common voxels of these two contrasts as a region of interest (Figure 5A). Note that the GLM contrast is not difficulty matched, as the GLM contrast is intended only as a means of obtaining a region-of-interest to use in the time-course plots. For the reasoning behind differentiating based on the participants baseline PSS (rather than true simultaneity of keypress and flash as measured by the computer) see Behavioral Methods, above. If one repeats the analysis for Figure 5B based on true 0, rather than the baseline PSS, the results are qualitatively similar, but some of the significant differences disappear.

Difficulty Matching

In order to match trials for difficulty (Figure 5) we chose trials that lay between $PSS_{baseline}$ −35 ms and $PSS_{baseline}$ +35 ms. This window of 35 ms is smaller than half of the average standard deviation of the psychometric functions, ensuring that all trials within the window are of approximately equal difficulty. 35 ms is also the average size of the PSS shift, ensuring that the trials in the 35 ms window after the baseline PSS in the injected delay block have approximately equal distribution of answers as those in the 35 ms window before the baseline PSS in the baseline block. We separated trials by their timing with respect to their temporal relationship to the baseline PSS and by the participant's perceptions, resulting in four trial types for each trial block (Figure 5B, inset). We then plotted the time-course of the activation in our region-of-interest (the ACC/MFC) for each type of trial during the 12 s after the participant makes the temporal order judgment (the cues to report one's perceptions occurred 6–12 s after the x axis point labeled 0 in this plot). Difficulty matching in this way is important to the conclusion that the ACC/MFC is most highly activated by the illusion. For example, without difficulty matching around the baseline PSS (the point at which participants are normally maximally uncertain about their answer), one could argue that the ACC/MFC activity on "illusion" trials is higher than the activity in veridical "before" trials because most of the "illusion" trials happen just after the baseline PSS, whereas veridical "before" trials can happen up to 150 ms before the participant's keypress, making these trials easier and polluting the average. However, in the difficulty-matched analysis, there are an equal distribution of trials close to the baseline PSS in the veridical "before" cases (black and gray, Figure 5B, inset) as in the illusory case. Thus, if trial difficulty (as measured by proximity to the baseline PSS) were the cause of the ACC/MFC activation, we would expect to find the black curve just as high, if not higher, than red curve in Figure 5B.

Acknowledgments

We thank Alex Holcombe, Brooks King-Casas, Adnan Lakhani, Jian Li, and Keith Kline for helpful discussions. This work was supported by the University of Texas, Houston Medical School (C.S. and D.M.E.) and NIDA R01 DA11723 and Kane Family Foundation and NINDS grant NS045790 (P.R.M.).

Received: October 13, 2005
Revised: March 27, 2006
Accepted: August 8, 2006
Published: September 6, 2006

References

Botvinick, M., Nystrom, L.E., Fissell, K., Carter, C.S., and Cohen, J.D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. Nature *402*, 179–181.

Botvinick, M., Braver, T.S., Barch, D.M., Carter, C.S., and Cohen, J.D. (2001). Conflict monitoring and cognitive control. Psychol. Rev. *108*, 624–652.

Brown, J.W., and Braver, T.S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. Science *307*, 1118–1121.

Bushara, K.O., Grafman, J., and Hallett, M. (2001). Neural correlates of auditory-visual stimulus onset asynchrony detection. J. Neurosci. *21*, 300–304.

Buhusi, C.V., and Meck, W.H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. Nat. Rev. Neurosci. 755–765.

Campbell, W.W., Ward, L.C., and Swift, T.R. (1981). Nerve conduction velocity varies inversely with height. Muscle Nerve *4*, 520–523.

Carter, C.S., Braver, T.S., Barch, D.M., Botvinick, M.M., Noll, D., and Cohen, J.D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. Science *280*, 747–749.

Coull, J.T., Vidal, F., Nazarian, B., and Macar, F. (2004). Functional anatomy of the attentional modulation of time estimation. Science 1506–1508.

Cosmelli, D., David, O., Lachaux, J.P., Martinerie, J., Garnero, L., Renault, B., and Varela, F. (2004). Waves of consciousness: ongoing cortical patterns during binocular rivalry. Neuroimage 128–140.

Cunningham, D.W., Billock, V.A., and Tsou, B.H. (2001). Sensorimotor adaptation to violation of temporal contiguity. Psychol. Sci. *12*, 532–535.

Eagleman, D.M. (2004). The where and when of intention. Science *303*, 1144–1146.

Eagleman, D.M., and Sejnowski, T.J. (2000a). Motion integration and postdiction in visual awareness. Science *287*, 2036–2038.

Eagleman, D.M., and Sejnowski, T.J. (2000b). The position of moving objects: response. Science *289*, 1107a.

Eagleman, D.M., and Holcombe, A.O. (2002). Causality and the perception of time. Trends Cogn. Sci. *6*, 323–325.

Farrer, C., and Frith, C.D. (2002). Experiencing oneself vs another person as being the cause of an action: the neural correlates of the experience of agency. Neuroimage *15*, 596–603.

Fujisaki, W., Shimojo, S., Kashino, M., and Nishida, S. (2004). Recalibration of audio-visual simultaneity. Nat. Neurosci. *7*, 773–778.

Genovese, C.R., Lazar, N.A., and Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. Neuroimage *15*, 870–878.

Haggard, P., Clark, S., and Kalogeras, J. (2002). Voluntary action and conscious awareness. Nat. Neurosci. *5*, 382.

Holroyd, C., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R.B., Coles, M.G.H., and Cohen, J.D. (2004). Dorsal anterior cingulated cortex shows fMRI response to interal and external error signals. Nat. Neurosci. *7*, 497–498.

Ito, S., Stuphorn, V., Brown, J., and Schall, J.D. (2003). Performance monitoring by the anterior cingulate cortex during saccade countermanding. Science *302*, 120–122.

Ivry, R.B., and Spencer, R.M. (2004). The neural representation of time. Curr. Opin. Neurobiol. *14*, 225–232.

Jech, R., Dušek, P., Wackermann, J., and Vymazal, J. (2005). Cumulative blood oxygenation-level-dependent signal changes support the 'time accumulator' hypothesis. Neuroreport *16*, 1467–1471.

Kao, Y., Davis, E., and Gabrieli, J. (2005). Neural correlates of actual and predicted memory formation. Nat. Neurosci. *8*, 1776–1783.

Karmarkar, U.R., and Buonomano, D.V. (2003). Temporal specificity of perceptual learning in an auditory discrimination task. Learn. Mem. *10*, 141–147.

Kerns, J.G., Cohen, J.D., MacDonald, A.W., III, Cho, R.Y., Stenger, V.A., and Carter, C.S. (2004). Anterior cingulate conflict monitoring and adjustments in control. Science *303*, 1023–1026.

King-Casas, B., Tomlin, D., Anen, C., Camerer, C.F., Quartz, S.R., and Montague, P.R. (2005). Getting to know you: reputation and trust in a two-person economic exchange. Science *308*, 78–83.

Lau, H.C., Rogers, R.D., Haggard, P., and Passingham, R.E. (2004). Attention to intention. Science *303*, 1208–1210.

Leung, H.C., Skudlarski, P., Gatenby, J.C., Peterson, B.S., and Gore, J.C. (2000). An event-related functional MRI study of the stroop color word interference task. Cereb. Cortex *10*, 552–560.

Lewis, P.A., and Miall, R.C. (2006). A right hemispheric prefrontal system for cognitive time measurement. Behav. Processes *71*, 226–234.

Linden, D.E.J., Prvulovic, D., Formisano, E., Voellinger, M., Zanella, F., Goebel, R., and Dierks, T. (1999). The functional neuroanatomy of target detection: an fMRI study of visual and auditory oddball tasks. Cereb. Cortex *9*, 815–823.

Lumer, E.D., Friston, K.J., and Rees, G. (1998). Neural correlates of perceptual rivalry in the human brain. Science *280*, 1930–1934.

Mangels, J.A., Ivry, R.B., and Shimizu, N. (1998). Dissociable contributions of the prefrontal and neocerebellar cortex to time perception. Brain Res. Cogn. Brain Res. *7*, 15–39.

Nachev, P., Rees, G., Parton, A., Kennard, C., and Husain, M. (2005). Volition and conflict in human medial frontal cortex. Curr. Biol. *15*, 122–128.

Namba, J., and Baldo, V.C. (2004). The modulation of the flash-lag effect by voluntary attention. Perception *33*, 621–631.

Nenadic, I., Gaser, C., Volz, H., Rammsayer, T., Haeger, F., and Sauer, H. (2003). Processing of temporal information and the basal ganglia: new evidence from fMRI. Exp. Brain Res. *148*, 238–246.

Nijhawan, R. (1994). Motion extrapolation in catching. Nature *370*, 256–257.

Nobre, A.C., and O'Reilly, J. (2004). Time is of the essence. Trends Cogn. Sci. *8*, 387–389.

Paillard, J. (1948). Recherche sur la simultanéité dans l'exécution des mouvements réactifs. C. R. Seances Soc. Biol. Fil. *142*, 935–937.

Paus, T. (2001). Primate anterior cingulate cortex: where motor control, drive and cognition interface. Nat. Rev. Neurosci. *2*, 417–424.

Pardo, J.V., Pardo, P.J., Janer, K.W., and Raichle, M.E. (1990). The anterior cingulate cortex mediates processing selection in the Stroop attentional conflict paradigm. Proc. Natl. Acad. Sci. USA *87*, 256–259.

Pouthas, V., George, N., Poline, J.B., Pfeuty, M., Vandemoorteele, P.F., Hugueville, L., Ferrandez, A.M., Lehericy, S., Lebihan, D., and Renault, B. (2005). Neural network involved in time perception: an fMRI study comparing long and short interval estimation. Hum. Brain Mapp. *25*, 433–441.

Purpura, K., Tranchina, D., Kaplan, E., and Shapley, R.M. (1990). Light adaptation in the primate retina: analysis of changes in gain and dynamics of monkey retinal ganglion cells. Vis. Neurosci. *4*, 75–93.

Rao, S.M., Mayer, A.R., and Harrington, D.L. (2001). The evolution of brain activation during temporal processing. Nat. Neurosci. *4*, 317–323.

Vroomen, J., Keetelsa, M., de Geldera, B., and Bertelson, P. (2004). Recalibration of temporal order perception by exposure to audio-visual asynchrony. Brain Res. Cogn. Brain Res. *22*, 32–35.

Walton, M.E., Devlin, J.T., and Rushworth, M.F.S. (2004). Interactions between decision making and performance monitoring within prefrontal cortex. Nat. Neurosci. *7*, 1259–1265.