Metacognitive blindness in temporal selection during the deployment of spatial attention

Samuel Recht^{1,2 *}, Vincent de Gardelle^{3 +}, Pascal Mamassian^{2 +}

¹ Department of Experimental Psychology, University of Oxford, Oxford, UK

² Laboratoire des systèmes perceptifs, Département d'études cognitives, École normale supérieure, PSL University, CNRS, Paris, France

³ CNRS and Paris School of Economics, Paris, France

* Corresponding author: samuel.recht@gmail.com

+ these authors share senior authorship

TO CITE THIS PAPER

Recht, S., de Gardelle, V., and Mamassian, P. (2021). "Metacognitive blindness in temporal selection during the deployment of spatial attention." *Cognition* 216:104864, DOI: [10.1016/j.cognition.2021.104864](https://doi.org/10.1016/j.cognition.2021.104864).

ABSTRACT

How does orienting attention in space affect the quality of our confidence judgments? Orienting attention to a particular location is known to boost visual performance, but the deployment of attention is far from being instantaneous. Whether observers are able to monitor the time needed for attention to deploy remains largely unknown. To address this question, we adapted a “Wundt clocks” paradigm, asking observers (N=140) to reproduce the phase of a rotating clock at the time of an attentional cue, and to evaluate their confidence in their responses. Attention affected the latency between objective and perceived events: the average reported phase was delayed in accordance with the known latencies of voluntary and involuntary attention. Yet, we found that confidence remains oblivious to these attention-induced perceptual delays, like a ‘metacognitive blind spot’. In addition, we observed weaker metacognition specifically during the deployment of voluntary attention, suggesting a tight relationship between the attentional and metacognitive systems. While previous work has considered how visual confidence adjusts to fully attended versus unattended locations, our study demonstrates that the very process of orienting attention in space can alter metacognition.

1. INTRODUCTION

Visual confidence is the subjective reliability of a preceding decision in the visual domain; broadly, it can be defined as the self-evaluation of perceptual performance (Mamassian, 2016). Confidence is an important second-order judgement that allows to objectively evaluate the quality of a first-order judgment (e.g. “the traffic light is green”): a confidence estimate can be directly matched to first-order accuracy. One’s ability to reliably track self-performance through confidence judgments has been coined ‘metacognitive ability’, or simply metacognition. Metacognition plays a critical role in task prioritization (Aguilar-Lleyda, Lemarchand, & de Gardelle, 2020), adaptive learning (Guggenmos, Wilbertz, Hebart, & Sterzer, 2016; Hainguerlot, Vergnaud, & De Gardelle, 2018), evidence accumulation (Balsdon, Wyart, & Mamassian, 2020), information seeking (Desender, Boldt, & Yeung, 2018) and the integration of multiple decision stages (van den Berg, Zylberberg, Kiani, Shadlen, & Wolpert, 2016). It can also be used to compare performance between different tasks and modalities (de Gardelle, Le Corre, & Mamassian, 2016; de Gardelle & Mamassian, 2014).

Performance variability is usually well tracked by confidence judgments, despite multiple dissociations being documented in the literature (for a review, see Fleming & Daw, 2017). For instance, in a recent study we found visual confidence to neglect that a single attentional episode can benefit two consecutive targets. When two targets are presented in close succession within a stream of distractors, observers’ discrimination of the second target is degraded when it appears around 300ms after the first target (the well-known ‘Attentional Blink’, e.g. Raymond, Shapiro, & Arnell, 1992). Yet, if the second target is presented immediately after the first, both are on average correctly reported: paying attention to the first item also benefits the second, a phenomenon referred to as ‘lag-1 sparing’ in the literature. However, confidence judgments remain blind to the initial sparing of performance, as if some aspects of attentional processing were concealed from metacognition (Recht, Mamassian, & de Gardelle, 2019).

Given the limitations of cognitive resources and the vast amount of peripheral sensory information, visual attention is thought to play the critical role of a filter: it selects, prioritizes and amplifies some specific sensory information for further processing (Buschman & Kastner, 2015; Carrasco, 2011). Such a selective mechanism provides an organism with an important tool to optimize resources allocation in space and in time. Changes in performance induced by attention should, ideally, be taken into account during decision-making. Intuitively, knowing whether one was paying attention at a given moment should help to estimate the reliability of one’s perception, a subjective reliability that might be used later when trying to optimize one’s behaviour. This reasoning can be extended to a situation where attention is on the verge of orienting, but is not yet fully deployed: such a situation should be considered as less reliable than a condition where attention had enough time to orient before target onset.

Of particular relevance to this question is the classical taxonomy in the literature that differentiates exogenous attention from endogenous attention through their respective temporal dynamics. Exogenous stands for an involuntary, early and short-lasting orienting of attention, while endogenous relates to a voluntary, late and long-lasting allocation (Carrasco, 2011). The nature of an attentional episode is therefore defined primarily by the time it takes to emerge, with exogenous attention taking roughly 100ms to be effective while approximately 300ms are necessary for endogenous attention to be allocated. Whereas the links between metacognition and the temporal aspects of attention have so far been little studied, there are

several reports on its spatial aspects. Some studies have found a dissociation between the two (Rahnev et al., 2011; Wilimzig, Tsuchiya, Fahle, Einhäuser, & Koch, 2008), while others show a tight positive association between confidence and the effects of spatial attention (Denison, Adler, Carrasco, & Ma, 2018; Zizlsperger, Sauvigny, & Haarmeier, 2012; Zizlsperger, Sauvigny, Händel, & Haarmeier, 2014). The distinction between exogenous and endogenous attention regarding confidence judgments is also far from being settled. One study found that only voluntary, but not involuntary attention is reflected in confidence judgments (Kurtz, Shapcott, Kaiser, Schmiedt, & Schmid, 2017). Yet, the experimental design used in this study cannot fully guarantee the strict exogenous nature of the cueing procedure. Moreover, we note that most of the aforementioned studies on spatial attention considered how perceptual and metacognitive judgments vary between attended and unattended locations, often using a typical cueing design known as the ‘Posner paradigm’ (Carrasco, 2011; Posner, 1980). In these studies, the authors measured accuracy and confidence after a valid or invalid pre-cue, and evaluated whether cue-induced changes in accuracy are reflected in confidence judgments. There is, however, an aspect of attentional deployment that has been overlooked so far: do observers incorporate the timing of spatial attention in their confidence judgments?

From the moment the visual system ‘decides’ to orient spatial attention, to the moment attention is actually deployed, some period of time should inevitably pass. In the present work, we investigated whether metacognition considers such delay when estimating perceptual performance. On one hand, confidence could be sensitive to the average delay: observers could be more confident for shorter than longer attentional latencies. On the other hand, confidence could also be sensitive to the variability in temporal selection, observers discerning that on a given trial, selection was triggered in an unusual way (either too early or too late) compared to other trials. In other words, confidence could be implementing both the average timing of attention and the precision of temporal selection to various - potentially distinct - degrees, or neither of these two dimensions.

Interestingly, in a recent work we found confidence to be oblivious to the average latency of temporal attention, despite reflecting some of its variability (Recht et al., 2019), but whether such results hold for the deployment latency of spatial attention remains an open question. Moreover, depending on the voluntary or involuntary nature of attention orienting, the sensitivity of confidence to the average delay in timing and selection variability can differ. Intuitively, it can be argued that both endogenous attention and metacognition involving a central stage processing, one might expect a better metacognitive access to fluctuation in performance induced by the timing of endogenous - compared to exogenous - attention. Hence, to which extent metacognition includes the timing of spatial voluntary or involuntary attention in confidence judgments is the subject of the present work.

Here, in two experiments, we adapted a ‘Wundt clocks’ paradigm where participants have to estimate when an event occurred by reproducing the time of a clock at the onset of a cue. Critically, the cue triggers the deployment of spatial attention towards one of multiple clocks presented in the periphery of the visual field. Such continuous report has been proposed to be a proxy for attentional timing (Carlson, Hogendoorn, & Verstraten, 2006; Chakravarthi & VanRullen, 2011; Hogendoorn, Carlson, VanRullen, & Verstraten, 2010). The ‘Wundt clocks’ design, contrary to the canonical Posner paradigms discussed earlier, also permits to decouple attention from performance: the collected responses both involve an average delay (as a function of attention) and a variability (as a function of temporal selection variability, or

'performance'). In the present study, we combine such a paradigm with a confidence forced-choice procedure (Barthelmé & Mamassian, 2009, 2010; de Gardelle & Mamassian, 2014, 2015) to assess metacognition in different attentional conditions (pre-cue, exogenous, and endogenous). Specifically, after completing two trials, participants were requested to indicate which of the two previous responses they felt more confident about being correct.

To pre-empt our results, our study revealed three main findings. First, visual confidence was oblivious to the average latency of both exogenous and endogenous spatial attention. Secondly, metacognition was sensitive to the precision of temporal selection, but was transiently altered during endogenous/voluntary orienting of attention to a particular location. Finally, to investigate the extent to which confidence in our task is related to a generic ability to evaluate the precision of temporal cognitive processes, we also measured confidence a simple detection task in which participants had to estimate their own response times (e.g. Gorea, Mamassian, & Cardoso-Leite, 2010). We found that metacognitive ability in the clock task was not correlated to metacognition of response times, suggesting that metacognition of temporal selection during attentional orienting is at least partially distinct from the introspection of internal sensorimotor variability.

In Experiment 1, we tested 20 participants in a within-subject design with eye-tracking in a controlled laboratory setting. In Experiment 2, we slightly changed the paradigm to account for potential features-related confounds, and replicated the results from Experiment 1 in an online, between-subjects design (120 participants).

1. EXPERIMENT 1

2.1. MATERIAL & METHODS

2.1.1. Participants

20 adult volunteers were recruited from the French RISC pool of participants (age $M \pm SD = 25.85 \pm 2.30$, 14 females). They all provided informed written consent prior to the experiment. Participants were compensated for their time at a rate of 10€ per hour. The experiment consisted in two 1-hour sessions. The experimental procedure was approved by the ethics review board of the Paris School of Economics (PSE).

2.1.2. Overview

Participants were engaged in two identical sessions. In both sessions, there were 3 blocks of the main task, and 2 blocks of the response time (RT) task. The order of the two tasks was counterbalanced across sessions and participants. The full experiment consisted in 432 trials for the main task and 288 trials for the RT task.

For both tasks, participants were instructed to fixate the centre of the screen during the whole trial period and their gaze was monitored online using an eye-tracker (EyeLink 1000 Hz, SR Research). Any trial during which participants blinked or moved their gaze away from the fixation dot were automatically aborted, and a new sample of the trials pair was added at the end of the block. Fixation was enforced from 200ms before cue onset.

Participants sat 60cm from the monitor (11 in front of a 1280x1024 pixels CRT monitor, 85Hz refresh rate, and 9 in front of a 1920x1080 pixels monitor, 60Hz refresh rate), their head

maintained with a chin-rest. The tasks were coded using the Python programming language and the PsychoPy library (Peirce, 2007), on a Mac OS computer.

2.1.3. Stimuli

Each trial started with a central fixation dot presented on a grey background for 1000ms (fig. 1). Then six clocks (black outline, 4° diameter, 0.2° thick) were presented on the right and left side of the screen, with 4 clocks on the diagonals (4° eccentricity from fixation), and 2 on the horizontal midline (6° from fixation). This small difference in eccentricity did not produce observable effects on attention or confidence (see Supplementary Material). Each clock had a black dot (diameter 0.2°) at the centre. The hand was a line starting 0.2° from the clock's centre (length: 1.4°). We call phase the clock time indicated by the hand orientation. Each clock rotated for 4s, at a fixed speed of 1 revolution per second, starting from a random onset phase.

At a random time uniformly drawn between 1000ms and 2000ms after clocks' onset, the probe was presented for about 20ms (1 frame for the 60Hz monitor and 2 frames for the 85 Hz monitor). In the pre-cue and exogenous conditions, this probe was a red annulus surrounding the clock (diameter: 4.1°), whereas in the endogenous condition, the cue was a central black line (length: 1°) pointing towards the clock. In addition, in the pre-cue condition, a green line (length: 1°) was displayed centrally for the whole trial duration, indicating with 100% validity the clock about to be probed. Each clock was probed with probability 1/6.

2.1.4. Main task

In each session, the 3 attention conditions were presented in 3 separate blocks of 72 trials each, in a randomized order. Each block started with 10 trials of practice. Participants were instructed to fixate the centre of the screen, to monitor all clocks (exogenous/endogenous conditions) or only the pre-cued clock (pre-cue condition), and to register the phase of the relevant clock at probe onset. They had to report the phase at the end of the trial using the mouse cursor (fig. 1). The display of the clock's hand during the report period was initiated by the participant's mouse click. Every two trials, participants were asked to select which of the two previous responses they felt more confident about being correct, by clicking on one of two squares (2°x2°) displayed at 6° eccentricity on each side below the fixation cross, flanked by '1' and '2' (for first or second trial in the pair). Based on this report, for each pair of trials, one trial was labelled as 'higher confidence' and the other one as 'lower confidence'. Participants were not instructed to make speeded responses. At the end of each ~15min block, participants had the opportunity to take a break. Our design did not involve between-conditions comparisons (i.e., one exogenous and one endogenous trial in the same pair) because it would have concerned a distinct experimental question, namely, the relative effect of cue type on confidence. Moreover, it would have added a significant cognitive switching cost for the participant. Yet, testing how observers are able to compare the differential effect of exogenous and endogenous attention is a question that could be of interest for further studies.

Figure 1: Experimental protocol (Exp 1). Main task: On each trial, the stimuli consisted in 6 clocks rotating for 4 seconds, at 1 revolution per second but with random starting phases. After a variable delay, one of the clocks was probed, either peripherally ('exogenous' condition) or centrally ('endogenous' condition). A third ('pre-cue') condition included a central pre-cue through the whole trial, indicating with 100% validity which clock will be probed. At the end of the trial, participants had to reproduce the phase of the clock at probe onset. Every two trials, participants had to indicate the best of their two preceding responses.

2.1.5. Reaction times (RT) task

Each session consisted of one block of 72 trials for the pre-cue condition and one block of 72 trials for the exogenous condition, in a random order. We did not test the endogenous condition in the RT task. The stimuli were identical to the main task, and participants had to make a speeded response by pressing a key at probe onset. Every two trials, participants were asked to select which of the two previous responses had the shortest reaction time (i.e. confidence forced-choice).

2.2. ANALYSES

2.2.1. Latency and overall precision

For our analyses, we focused on the two parameters of the von Mises distribution - typically used for circular data in place of the normal distribution - to probe the average latency of attention and the overall precision of temporal selection, together with confidence:

$$f(x | \mu, k) = \frac{e^{k \cos(x-\mu)}}{2\pi I_0(k)} \text{ for } \mu, x \text{ in } (-\pi, \pi] \text{ and } k > 0,$$

where I_0 is the modified Bessel function of the first kind of order 0, μ is the location parameter, k is the concentration parameter and x the angular error in a given trial.

For each trial, the angular error between the phase of the clock at cue onset and the phase reported by the participant was calculated. We fitted von Mises distributions on angular values using maximum likelihood estimation (MLE) separately for each participant and condition (see Supplementary Material). The location parameter of the distribution (equivalent to the mean of a normal distribution) is an estimate for the latency of response errors, since it relates to the average time difference between the objective event and the perceived event. The concentration parameter (or ‘kappa’) is an estimate for the overall precision of the responses, and its inverse is analogous to the variance of a normal distribution. We confirmed that the simple von Mises model was not significantly worse than two other alternatives: a skewed von Mises, to account for potential asymmetry due to the direction of rotation, and a mixture of two von Mises, with distinct location and precision parameters, to account for potential systematic differences in early/late orienting of attention leading to bimodality (see Supplementary Material).

It is important to distinguish the experiment-wise, overall precision of responses from the true underlying, trial-wise precision of memory encoding. While the two are fundamentally related, the exact nature of encoding precision in working memory has been subject to debate, and has led to a myriad of ‘resource models’ (for a review, see Ma, Husain, & Bays, 2014). Yet, they often differ in their understanding of how resource allocation is actually implemented. We therefore decided to consider a standard, non-specific statistics of error dispersion (i.e. the concentration parameter kappa of von Mises distributions), rather than a more specific mechanistic implementation of the cognitive process giving rise to such errors. We use the term ‘overall precision’ to refer to the concentration parameter of the von Mises fits, to make it distinct from the trial-wise, internal precision level of the participant (see Supplementary Material and Section 4 for a discussion).

2.2.2. Within-pair difference in error magnitude

In a second type of analyses, we looked at the difference in error magnitude within a pair of trials. To predict confidence from the difference in error between the two trials of the pair, we used logistic regression models. One potential challenge to such an approach comes from the circularity of our decision space: projected on a $[-180^\circ, 180^\circ)$ axis, a few, large-magnitude errors might be categorised as negative values, despite being generated by a positive delay. To minimise the occurrence of such cases, we centred the error space on the

circular mean (i.e. the average latency) of the distribution per participant, condition and confidence level. The space was then cut at the exact opposite of the circular mean, and the axis was then moved back to the original range. This approach allowed us to recalibrate the sign of extreme errors by moving to a near-linear space, on the realistic assumption that the underlying generative process obeys a von Mises distribution (see Supplementary Material).

Simple model

$$\Delta\varepsilon = |\varepsilon_A| - |\varepsilon_B| \text{ (eq. 1)}$$

where ε_A and ε_B are the error in first and second trial in the confidence pair. A negative value of $\Delta\varepsilon$ would indicate a greater error for trial B, and a positive value a greater error for trial A.

Centred model

$$\Delta\varepsilon = |\varepsilon_A - \mu| - |\varepsilon_B - \mu| \text{ (eq. 2)}$$

where μ is the average circular error (or latency) for the considered participant/condition.

Scaling model

$$\Delta\varepsilon = \frac{|\varepsilon_A - \mu| - |\varepsilon_B - \mu|}{|\varepsilon_A - \mu| + |\varepsilon_B - \mu|} \text{ (eq. 3)}$$

For all models, we used the relative error ($\Delta\varepsilon$) as a predictor of confidence in a logistic regression model (logit), estimated per participant and condition separately:

$$\text{logit}(p) = \log\left(\frac{p}{1-p}\right) = \alpha + \beta\Delta\varepsilon,$$

where

$$p = p(\text{Higher Confidence Trial} = \text{Trial}_B)$$

Given the symmetry of the confidence forced-choice paradigm, there should not be an advantage to choose the first or second decision of a confidence pair, and therefore the parameter α should theoretically be zero. Non-zero positive values of this parameter represent biases to choose the second decision as more confident.

All the analyses were carried out using the R programming language (R Core Team, 2013). When necessary, ANOVAs were corrected using the Greenhouse-Geisser adjustment and t-tests were corrected using the Welch-Satterthwaite adjustment. We report Wilcoxon signed rank test using uppercase T when the Shapiro-Wilk normality test failed, and Student test using lowercase t otherwise. For the Student-tests and Pearson correlations, we also systematically indicated the Bayes Factor (BF_{10}), computed using the 'ttestBF' and the 'correlationBF' functions from the 'BayesFactor' R package (Morey & Rouder, 2018).

2.3. Results

2.3.1. Main task

2.3.1.1. Spatial orienting of attention and temporal selection

We compared our 3 conditions of attention (pre-cue, exogenous, endogenous) in terms of the latency and overall precision of participants' responses (fig. 2).

The latency profile (fig. 2B) was consistent with results from the spatial attention literature showing a faster orienting for exogenous/peripheral cues compared to endogenous/central cues (see Carlson et al., 2006 for a study using the same paradigm and Carrasco, 2011 for a review). A repeated-measures ANOVA revealed a main effect of condition ($F(1.51, 28.76) = 194.10$, $MSE = 1698\text{ms}$, $p < 0.001$). Bonferroni corrected t-tests ($\alpha=0.017$, corrected for 3 tests) confirmed that latency was lower for the pre-cue condition than for both the exogenous condition ($t(19) = -6.28$, $p < 0.001$; $M = -53.70\text{ms}$, 95% CI $[-71.59 -35.81]$; $BF_{10} = 4.16 \times 10^3$) and endogenous condition ($t(19) = -15.28$, $p < 0.001$; $M = -214.63\text{ms}$, 95% CI $[-244.02 -185.22]$; $BF_{10} = 1.90 \times 10^9$), and lower for the exogenous condition than for the endogenous condition ($t(19) = -14.98$, $p < 0.001$; $M = -160.93\text{ms}$, 95% CI $[-183.41 -138.45]$; $BF_{10} = 1.37 \times 10^9$).

Figure 2: Latency of attention and temporal selection (Exp 1). (A) Distributions of errors for a representative participant. The distribution of responses represents the angular error (reported phase minus objective phase) converted in milliseconds. Von Mises distributions were fitted to estimate the latency (location parameter) and the overall precision (concentration parameter) of temporal selection. The pre-cue, exogenous and endogenous conditions are represented in green, red and blue, respectively. (B) The average latency for each condition, reproducing the delays generally observed in the literature. (C) Overall precision for each condition, showing no difference between pre-cue and endogenous condition. Coloured dots correspond to individual participants in the given condition. Large dots with a black outline represent the mean across participants. Error bars represent across participants ± 1 SEM.

Regarding overall precision, a second ANOVA with concentration parameter as a dependent variable showed a main effect of condition ($F(1.91, 36.23) = 4.04$, $MSE = 0.43$, $p = 0.03$). However, Bonferroni-corrected t-tests (with $\alpha=0.017$, corrected for 3 tests) showed no convincing evidence for a difference between pre-cue and exogenous condition ($t(19) = -2.19$, $p = 0.041$; $M = -0.46$, 95% CI $[-0.89 -0.02]$, $BF_{10} = 1.62$), between pre-cue and endogenous condition ($t(19) = 0.40$, $p = 0.690$; $M = 0.07$, 95% CI $[-0.30 0.45]$, $BF_{10} = 0.25$) or between exogenous and endogenous condition ($t(19) = 2.44$, $p = 0.024$; $M = 0.53$, 95% CI $[0.08 0.98]$, $BF_{10} = 2.45$). Importantly, the profile of the concentration parameter suggested that pre-cue and endogenous conditions led to roughly similar performance, and only some evidence was found for a gain during exogenous orienting of attention (fig. 2C).

2.3.1.2. Metacognition of attentional effects

To evaluate metacognitive ability, we estimated location and concentration parameters separately for the higher confidence trials and the lower confidence trials (fig. 3A), according to the confidence forced-choice response within each condition. We then evaluated in two distinct ANOVAs how location and concentration depended on confidence, the attentional condition, and the interaction between these two factors (fig. 3B and 3C).

Figure 3: Metacognition of latency and overall precision (Exp 1). (A) Latency and concentration parameters were estimated by fitting von Mises distributions to the ‘higher (purple) and ‘lower’ (grey) confidence trials for each participant/condition. The parameter differences between ‘higher’ and ‘lower’ confidence give an estimate of metacognitive access to attentional latency and overall precision. The figure plots the two distributions for a representative participant for one of the three conditions. (B) The average difference in latency between ‘lower’ and ‘higher’ confidence trials for each condition. The absence of a significant difference suggests that confidence is oblivious to attentional latency. (C) The average difference in overall precision between ‘higher’ and ‘lower’ confidence trials. Positive values indicate that confidence has access to the variability of responses for all conditions, even if there was a lower difference for endogenous orienting compared to pre-cue. Coloured dots correspond to individual participants in the given condition. Large dots with a black outline represent the mean across participants. Error bars represent across participants ± 1 SEM.

For latency (i.e. the location parameter, fig. 3B), the ANOVA indicated a main effect of condition ($F(1.51, 28.64) = 194.75$, $MSE = 3391ms$, $p < 0.001$) as expected given the large effect of latency across conditions when the data were not split between higher and lower confidence. However, we did not find a significant effect of confidence ($F(1, 19) = 0.99$, $MSE = 341.00ms$, $p = 0.33$), nor a confidence $\times$ condition interaction ($F(1.73, 32.92) = 1.00$, $MSE = 248.66ms$, $p = 0.37$). The absence of a main effect of confidence suggests that within each attentional condition, confidence was oblivious to the delays induced by spatial orienting of attention. Regarding precision (i.e. the concentration parameter, fig. 3C), as expected, we found significant main effects of condition ($F(1.95, 37.02) = 3.73$, $MSE = 1.13$, $p = 0.03$) and confidence

($F(1,19) = 48.25$, $MSE = 1.0$, $p < 0.001$). Interestingly, we also observed a condition $\times$ confidence interaction ($F(1.87,35.54)=7.05$, $MSE=0.38$, $p=0.003$). These results demonstrate that trials labelled with higher confidence were associated with greater overall precision compared to lower confidence trials. Bonferroni-corrected t-tests ($\alpha=0.017$, corrected for 3 tests) confirmed a greater precision for higher confidence trials for all conditions, namely pre-cue ($T(19) = 210$, $p < 0.001$), exogenous ($T(19) = 209$, $p < 0.001$) and endogenous ($t(19) = 6.70$, $p < 0.001$; $M = 0.76$, 95% CI [0.52 0.99] ; $BF_{10} = 9.1 \times 10^3$) conditions. Corrected paired tests ($\alpha=0.017$) showed a significant difference only between the pre-cue and the endogenous condition ($T(19) = 179$, $p = 0.004$). We found no strong evidence for a difference between pre-cue and exogenous ($t(19) = 1.85$, $p = 0.080$; $M = 0.46$, 95% CI [-0.06 0.99]; $BF_{10} = 0.97$) or between exogenous and endogenous condition ($t(19) = 2.17$, $p = 0.043$; $M = 0.54$, 95% CI [0.02 1.05]; $BF_{10} = 1.57$). Together these results indicate that confidence computation has access to the overall precision of responses. Furthermore, the greater metacognitive ability for the pre-cue compared to the endogenous condition points to a potential interaction between the deployment of voluntary attention and metacognitive ability.

2.3.1.3. Trial-by-trial variability of precision

To investigate if overall precision was distinct from trial-by-trial encoding precision, we tested different alternative models to the simple von Mises, with parameters fixed or not between conditions (see Supplementary Material). No alternative model was significantly better, but we nonetheless tested parameter estimations using the best performing alternative, a model with a fluctuating precision level following a higher order distribution. This variable-precision model (Van Den Berg, Shin, Chou, George, & Ma, 2012) did not significantly change the pattern of estimated precision (see Supplementary Material for a BIC and AIC model comparison).

2.3.1.4. How well does confidence predict pair-wise error difference?

The previous analyses suggested that participants were oblivious to latency, but were accurately monitoring the precision of their response across trials. Furthermore, orienting endogenous attention appeared to induce a decrease in metacognition of overall precision. These analyses however give us a broad picture of attention orientation across conditions, leaving open the question of the relative error estimation between two trials. A straightforward approach to this question is to assume that confidence has access to some form of evidence for each trial, which can be - from the experimenter point of view - related to the difference in absolute error between trial A and trial B in each pair of trials (that is, the relative error magnitude). To investigate how metacognition tracked trial-by-trial errors, we used a logistic regression model to predict confidence judgment from the difference in error magnitude between the two trials.

Three models were compared. In the 'simple model', confidence choices were predicted from the comparison of absolute error magnitude between the two trials (eq. 1 in Section 2.2.2.). In the 'centred model', given that confidence was not able to access the overall attentional delay within each condition (Section 2.3.1.2. and fig. 3B), the deviation from the mean error within the condition was used, instead of the absolute error (eq. 2). Finally, in the 'scaling model', the comparison of errors between two trials was divided by the magnitude of both errors (eq. 3), a form of scaling observed in perception (Fechner, 1964; Shepard, 1987).

The total log-likelihood across participants and conditions favoured the centred model (simple model: -2697; centred model: -2655; scaling model: -2703). We compared the goodness-of-fit of each model using a Log-Likelihood Ratio (LLR), all models sharing the same number of parameters. Model comparison provided evidence for the 'centred' model compared to both the 'simple' model (LLR = 84.18) and the 'scaling' model (LLR = 96.89). We also found greater likelihood for the 'simple' model compared to the 'scaling' model (LLR = 12.71). We therefore selected the 'centred' model for all subsequent analyses. The average classification accuracy for this model, estimated as the Area Under the Curve (AUC) was 62% (chance performance is 50%) across participants and conditions (overall AUC $M = 0.62$, 95% CI [0.60 0.64]; pre-cue AUC $M = 0.65$, 95% CI [0.62 0.68]; exogenous AUC $M = 0.61$, 95% CI [0.58 0.64]; endogenous AUC $M = 0.60$, 95% CI [0.57 0.63]). The pseudo- R^2 (McFadden, 1973) for this model was on average, across participants and conditions, equal to 0.05 (95% CI [0.04 0.06]). It should be noted that the pseudo- R^2 is not directly comparable to R^2 , the former having considerably lower values than the latter. Pseudo- R^2 values of 0.2 - 0.4 are said to represent excellent fits (chapter 15 in Hensher & Stopher, 2021). However, given the relatively low pseudo- R^2 value, the following analyses have to be considered in light of our previous overall analysis of metacognition (Section 2.3.1.2).

The slope (β) of the model gives an estimate of metacognitive ability (fig. 4). After Bonferroni correction ($\alpha = 0.017$), the slope was significantly positive for the three conditions, pre-cue ($t(19) = 7.58$, $p < 0.001$; $M = 0.007$, 95% CI [0.005 0.009]; $BF_{10} = 4.43 \times 10^4$), exogenous ($t(19) = 6.40$, $p < 0.001$; $M = 0.005$, 95% CI [0.003 0.006]; $BF_{10} = 5.24 \times 10^3$) and endogenous condition ($t(19) = 5.78$, $p < 0.001$; $M = 0.003$, 95% CI [0.002 0.004]; $BF_{10} = 1.60 \times 10^3$), showing above chance metacognitive ability for all conditions at the group level. Importantly, a comparison of the three conditions using a repeated-measures ANOVA showed a significant effect of condition ($F(1.98, 37.66) = 7.57$, $MSE = 0.00$, $p = 0.002$). After Bonferroni-correction ($\alpha = 0.05/3$), we found a strong evidence for a difference between the pre-cue and the endogenous conditions ($t(19) = 4.07$, $p < 0.001$; $M = 0.004$, 95% CI [0.002 0.006]; $BF_{10} = 53.56$). However, only little evidence was found for a difference between pre-cue and exogenous conditions ($t(19) = 2.15$, $p = 0.045$; $M = 0.002$, 95% CI [0 0.004]; $BF_{10} = 1.52$) and no evidence for a difference between exogenous and endogenous conditions ($t(19) = 1.64$, $p = 0.12$; $M = 0.002$, 95% CI [0 0.004]; $BF_{10} = 0.73$).

For completeness, we also analysed the intercept of the regression, which provides an estimate of metacognitive bias in favour of the second trial in the pair (independently of actual performance). We found moderate evidence for a bias in the pre-cue condition ($t(19) = -2.72$, $p = 0.014$; $M = -0.49$, 95% CI [-0.86 -0.11]; $BF_{10} = 3.98$). We found no evidence for a bias in the exogenous ($t(19) = -1$, $p = 0.330$; $M = -0.13$, 95% CI [-0.41 0.14]; $BF_{10} = 0.36$) or endogenous conditions ($t(19) = -0.99$, $p = 0.334$; $M = -0.13$, 95% CI [-0.42 0.15]; $BF_{10} = 0.36$). The ANOVA showed a significant effect of condition ($F(1.98, 37.57) = 8.89$, $MSE = 0.10$, $p < 0.001$). Bonferroni-corrected t-tests confirmed a significantly greater bias for the pre-cue condition compared to

Figure 4: Confidence judgment and the difference in error magnitudes (Exp 1). (A) The probability of higher confidence for trial B in the pair, as a function of the error difference between the two trials ($\Delta\epsilon$) at the group level. For illustrative purposes, errors have been grouped by quantiles at the individual level, and the group average is represented with error bars for each quantile. The significant decrease in slope between the pre-cue (green) and endogenous condition (blue) suggests a metacognitive cost during voluntary orienting of spatial attention. (B) The average correlation coefficient (or slope) for each condition, used as a measure of metacognitive ability. Coloured dots correspond to individual participants in the given condition. Large dots with a black outline represent the mean across participants. Error bars represent across participants ± 1 SEM.

the endogenous ($t(19) = -3.84$, $p = 0.001$; $M = -0.36$, 95% CI $[-0.55 -0.16]$; $BF_{10} = 33.84$) and exogenous conditions ($t(19) = -3.64$, $p = 0.002$; $M = -0.36$, 95% CI $[-0.56 -0.15]$; $BF_{10} = 21.91$), but no significant difference between exogenous and endogenous conditions ($t(19) = 0.009$, $p = 0.99$; $M = 0.0008$, 95% CI $[-0.21 0.21]$; $BF_{10} = 0.23$). Therefore, there was evidence for a bias in the pre-cue condition only. The pre-cue average bias log-odds was -0.49, which corresponds to a 12% increase in the probability of selecting the first trial compared to the second trial in the pair. To rule out a potential effect of the bias on the estimated metacognitive ability (or slope) found in the pre-cue condition, we tested the correlation between the value of the bias and the slope across participants. We found no evidence for a correlation ($t(18) = -0.09$, $p = 0.93$; $r = -0.01$, 95% CI $[-0.26 0.24]$; $BF_{10} = 0.47$). We extended such analysis across conditions and participants, with a similar outcome ($t(58) = -0.09$, $p = 0.93$; $r = -0.01$, 95% CI $[-0.26 0.24]$; $BF_{10} = 0.29$).

2.3.2. RT task

The RT task was designed to probe the metacognitive ability regarding the timing of a distinct, sensorimotor process, using the same stimuli as in the clock task. This task included only the pre-cue and exogenous conditions, which we did not expect to differ. Two participants were excluded from the analyses due to a technical error ($N = 18$). A repeated-

measures ANOVA with median response time as a dependent variable and condition/confidence as independent variables showed an effect of confidence ($F(1,17) = 39.93$, $MSE = 0$, $p < 0.001$) but no effect of condition ($F(1,17) = 0.95$, $MSE=0$, $p = 0.34$) and no interaction ($F(1,17) = 1.41$, $MSE = 0.0$, $p=0.25$). Participants were therefore able to discriminate between fast and slow response times, and, as expected, condition did not significantly affect this ability. For all subsequent analyses, we therefore combined both conditions together.

We tested whether our results for metacognition in the main attention task could be a general ability that extends to the timing of another cognitive and sensorimotor process (i.e., metacognition in the RT task), by evaluating across participants the correlation between metacognitive abilities in the attention and sensorimotor domain. A lack of correlation here would suggest that metacognitive ability in our main task does not necessarily reduce to sensorimotor uncertainty, and might thus be related to another process, probably linked to the state of spatial attention.

We found no evidence for a correlation between the metacognition in the RT task (i.e. the difference in median RT between lower and higher confidence trials) and the metacognition in the clock task (i.e. the difference in kappa between higher and lower confidence trials), for any of the conditions (correlation with the pre-cue condition: $t(16) = 0.19$, $p = 0.85$; $r = 0.05$, 95% CI [-0.43 0.50]; $BF_{10} = 0.54$; exogenous: $t(16) = 0.16$, $p = 0.88$; $r = 0.04$, 95% CI [-0.43 0.50]; $BF_{10} = 0.5$; endogenous: $t(16) = -0.47$, $p = 0.64$; $r = -0.12$, 95% CI [-0.55 0.37]; $BF_{10} = 0.33$).

In a similar vein to the first task, we evaluated how pair-wise, second-order comparison judgments (here, RT comparisons) could be predicted from the difference in first-order performance between the two trials (here, the difference in RTs), using a logistic regression. We then compared the simple model (ΔRT , log-likelihood: -1458.66) to the scaling model ($\frac{\Delta RT}{\Sigma RT}$, log-likelihood: -1454.09) using the likelihood ratio. Contrary to the first task, we found evidence in favour of the scaling model over the simple model ($LLR = 9.14$). We did not use the centred model for this analysis, given the nature of the task. Indeed, participants had to select the trial in which they were the fastest, not the trial in which the RT was closest to the mean RT in the condition. This result suggests a potential difference in the nature of the metacognitive evidence used in each task. We therefore used the scaling model in all subsequent analyses (group-level average pseudo- $R^2 = 0.10$; 95% CI [0.07 0.013]). Metacognitive ability (the β in the model) was significantly positive at the group level ($t(17) = 6.83$, $p < 0.001$; $M = 5.47$, 95% CI [3.78 7.16]; $BF_{10} = 6.79 \times 10^3$), and we found no significant evidence for a metacognitive bias (the intercept in the model, $t(17) = -1.72$, $p = 0.10$; $M = -0.38$, 95% CI [-0.84 0.09]; $BF_{10} = 0.83$).

We then tested the correlation of the confidence slope (beta) between the clock task and the RT task. We found no strong evidence for a correlation, for any of the conditions (correlation with the pre-cue condition: $t(16) = 2.34$, $p = 0.033$; $r = 0.50$, 95% CI [0.05 0.79]; $BF_{10} = 2.85$; exogenous: $t(16) = 1.62$, $p = 0.12$; $r = 0.38$, 95% CI [-0.11 0.72]; $BF_{10} = 1.23$; endogenous: ($t(16) = 0.21$, $p = 0.84$; $r = 0.05$, 95% CI [-0.42 0.51]; $BF_{10} = 0.50$) using Bonferroni correction ($\alpha=0.05/3$).

In sum, given the lack of evidence for a strong correlation between sensorimotor and temporal selection metacognition, and given the different models for these two domains (with the scaling model best fitting the RT task but not the main task data), our data suggest that

metacognition of errors in the first task and metacognition of RT in the second task are relying on – at least partially – distinct sources of second-order evidence.

3. Experiment 2

In Experiment 1, we did not find any significant difference in precision between the pre-cue and endogenous condition, but we found significantly lower metacognition for the endogenous condition. However, this effect could be in principle mediated by the difference in visual features between conditions rather than attention per se. A similar argument can be made about the apparent latency blind spot. In particular, the ‘pre-cue’ condition in Exp 1 is a hybrid of both exogenous and endogenous conditions, and might benefit from their respective features. To rule out this possibility, we conducted a second, online, between-subjects experiment where we equated as much as possible the visual characteristics of each condition. In particular, we decoupled the temporal cue from the spatial cue. The moment to be remembered was indicated in the exact same way for all conditions, by flashing all the clocks at once. The clock to be reported was indicated by a red spatial cue. Hence, the pre-cue condition was virtually the same as the endogenous condition at flash onset. Moreover, only the colour of the temporal cue was differing in the exogenous condition, to allow for the peripheral, automatic capture of attention. We also equated fixation dot and central cue features across conditions.

3.1. MATERIAL & METHODS

3.1.1. Participants

123 adult volunteers were recruited via the Prolific online platform (age $M \pm SD = 30.5 \pm 8$ year-old, 57 females). They all provided informed electronic consent prior to the experiment. Participants were compensated 10 € for their time, plus a bonus of 1.75 € for the 50% best performing participants. An a priori <80% accuracy cut-off in catch trials was set to ensure covert orienting of attention (see below). Following this rule, one participant (60% accuracy) was excluded. Two other participants were discarded based on their written feedbacks at the end of the experiment: one admitting to try different screen orientations during the course of the experiment and another having repetitive display problems. The experiment consisted in one session (duration $M \pm SD = 57.1 \pm 13.6$ min). The experimental procedure was approved by the ethics review board of the Paris School of Economics (PSE).

3.1.2. Overview

Participants were randomly assigned to one of three treatment groups, corresponding to the three conditions from Experiment 1. The experiment consisted in 220 trials in total, including 20 catch trial pairs, where a catch target appeared on first, second, or both trials (see below). These 20 pairs were discarded from analyses. In 8 participants, a technical error led to the loss of 2 trials (out of the 180 trials kept for further analyses).

Participants were instructed to fixate the centre of the screen during the whole trial period. To ensure participants were compliant, 9% of trials were ‘catch trials’, where a letter

target was briefly presented at the fixation point. The target was nearly impossible to discriminate without foveal fixation, but was very easy to recognize with foveal fixation, as assessed beforehand in a small pilot study. Participants were informed that the target letter was meant to check their fixation compliance during the task, and that failure to accurately report it would disqualify them from getting a bonus. All but one participant performed very well during these catch trials, confirming that they were fixating the centre of the screen during the experiment (average accuracy $M \pm SD = 0.97 \pm 0.04$). Outliers with a circular IQR on clock reproduction error above/below 1.5 times the treatment group IQR were excluded: 8, 4 and 2 in the pre-cue, exogenous and endogenous group, respectively, leaving 107 participants for further analysis. For these outliers, the distribution of the angular errors was resembling a uniform distribution. Individual raw error plots for all participants, including these outliers, are provided in the Supplementary Material.

The tasks were coded in javascript using the jsPsych library (de Leeuw, 2015). The visual angle values indicated in the following paragraphs assume an average 60cm distance from screen.

3.1.3. stimuli

The experiment was presented in fullscreen mode, and exiting the mode during course of the experiment was systematically reversed in the next trial. Each trial started with a central fixation dot (0.6°) presented on a light grey background for 500ms. Six black arrows (length: 0.6°) were systematically presented around the fixation dot, each pointing to six possible clock locations (fig. 5). Then six, equally-spaced clocks (black outline, 4° diameter, 0.2° thick) were presented at 5° -eccentricity around the fixation point. Each clock had a black dot (diameter 0.2°) at the centre. The hand was a line starting 0.2° from the clock's centre (length: 1.6°). Each clock rotated for 5s, at a fixed speed of 1 revolution per second, starting from a random onset phase. At a random time uniformly drawn between 2s and 3s after clocks' onset, all clocks were briefly flashed in mid grey for 80ms. This temporal cue was complemented with a spatial red cue specific to the treatment group. In the pre-cue group, the cue was one of the central arrows, indicating during the whole trial and with 100% validity the clock to be remembered at flash onset. In the exogenous group, the clock of interest was flashed in red instead of mid grey, making it distinguishable from the other clocks. In the endogenous group, the central arrow pointing to the clock of interest briefly turned red for 80ms. Each clock was probed with probability $1/6$.

3.1.4. Procedure

Following a consent form, participants were presented with a standard visual adjustment procedure, in order to adapt the experiment to the characteristics of their screen. They had to rescale a rectangle using the mouse cursor to reproduce the size of a credit card held up on the screen. The estimated scaling factor was then used for the whole experiment, and allowed to present stimuli in standardised units. Following the procedure, participants were presented with instructions, followed by 8 practise trials, 4 trials in slow motion (half the speed), and 4 trials at full speed. Participants were instructed to fixate the centre of the screen, to monitor all clocks (exogenous/endogenous group) or only the pre-cued clock (pre-cue group), and to register the phase of the probed clock at flash onset (fig. 5). The mouse cursor was systematically hidden during the stimuli phase. Participants had to report the phase at the end of the trial using the mouse cursor. The display of the clock's hand during the report period was initiated by the participant's mouse click. To validate their response, participants

were required to click on the fixation point, which turned blue following participant's response. This approach was used to get the cursor centred before each trial. Every two trials, participants were asked to select which of the two previous responses they felt more confident about being correct, by clicking on one of two squares ($2^\circ \times 2^\circ$) displayed at 1.4° eccentricity on each side of the fixation point, flanked by '1' and '2' (for first or second trial in the pair). Once the trial was selected, the fixation dot turned blue, allowing participants to click on it. Based on this report, for each pair of trials, one trial was labelled as 'higher confidence' and the other one as 'lower confidence'. Participants were not instructed to make speeded responses. They were informed that some of the trials were catch trials, during which a target letter ('A' vs. 'X', 0.4°) was presented for 80ms within the fixation point, at the same time as the cue. At the end of these catch trials, participants were requested first to report the phase of the clock, and then the target letter using the mouse cursor. A feedback on the letter discrimination task was given immediately after their response. Catch targets were randomly interleaved within the course of the experiment, occurring at first, second or on both trials in a given pair. At the end of each block that lasted about 4min, participants had the opportunity to take a 20s break. Participants were informed that accuracy in both clock reproduction, letter report and confidence judgements will be considered for extra rewards calculation, and that failure to accurately report the target letter was eliminatory. At the end of the experiment, participants had the opportunity to provide a written feedback on the task.

Figure 5: Experimental protocol for Exp 2. On each trial, the stimuli consisted in 6 clocks rotating for 5 seconds, at 1 revolution per second but with random starting phases. After a variable delay, all clocks were briefly flashed, and one of the clocks was cued, either

peripherally ('exogenous' group) or centrally ('endogenous' group). A third ('pre-cue') group was presented with a central pre-cue through the whole trial, indicating with 100% validity which clock had to be monitored. At the end of the trial, participants had to reproduce the phase of the clock at flash onset. Every two trials, participants had to indicate the best of their two preceding responses. In 9% of trials, participants had to also report the identity of a letter ('A' vs. 'X') presented at fixation during flash onset.

3.2. ANALYSES

We applied the same approach in Exp 2 than in Exp 1. Please refer to Section 2.2. for details.

3.3. RESULTS

3.3.1. Spatial orienting of attention and temporal selection

We compared our 3 groups (pre-cue, exogenous, endogenous) in terms of the latency and overall precision of participants' distribution of responses (fig. 6).

The latency profile (fig. 6B) was strikingly similar to Exp 1, despite the experiment being ported online, and the timing characteristics / visual features being adapted. A repeated-measures ANOVA revealed a main effect of condition ($F(2,103) = 203.68$, $MSE = 2415.26\text{ms}$, $p < 0.001$). Bonferroni corrected t-tests ($\alpha=0.017$, corrected for 3 tests) gave very strong evidence for lower latency in the pre-cue condition compared to both the exogenous condition ($t(63.89) = -6.53$, $p < 0.001$; $M = -70.42\text{ms}$, 95% CI $[-91.95 -48.89]$; $BF_{10} = 5.66 \times 10^5$) and endogenous condition ($t(65.13) = -20.5$, $p < 0.001$; $M = -229.41\text{ms}$, 95% CI $[-251.76 -207.06]$; $BF_{10} = 3.00 \times 10^{26}$), as well as lower latency for the exogenous condition than for the endogenous condition ($t(71.93) = -12.84$, $p < 0.001$; $M = -159.00\text{ms}$, 95% CI $[-183.68 -134.30]$; $BF_{10} = 9.52 \times 10^{16}$).

Similar to Exp 1, overall precision was assessed using a second ANOVA with concentration as a dependent variable, but this time showed no significant effect of condition ($F(2,103) = 1.07$, $MSE = 1.54$, $p = 0.345$). Therefore, no evidence was found for a difference in overall precision between the three groups (fig. 6C).

Figure 6: Latency of attention and temporal selection (Exp 2). (A) Distributions of errors for three representative participants. The distribution of responses represents the angular error (reported phase

minus objective phase) converted in milliseconds. The pre-cue, exogenous and endogenous conditions are represented in green, red and blue, respectively. (B) The average latency for each condition. (C) The average concentration for each condition, with no evidence for an effect of condition. Coloured dots correspond to individual participants in the given condition. Large dots with a black outline represent the mean across participants. Error bars represent across participants ± 1 SEM.

3.3.2. Metacognition of attentional effects

We then evaluated metacognitive ability, analysing the location and concentration parameters separately for the ‘higher’ and the ‘lower’ confidence trials (fig. 7A), according to the confidence forced choice response within each condition. Two distinct ANOVAs tested how location and concentration depended on confidence, the attentional condition, and the interaction between these two factors (fig. 7B and 7C).

For latency (fig. 7B), the ANOVA indicated a main effect of condition ($F(2,103)=200.54$, $MSE=4900.51ms$, $p<0.001$), a significant effect of confidence ($F(1,103) = 19.61$, $MSE = 255.46ms$, $p = 0.009$), and no confidence $\times$ condition interaction ($F(2,103) = 0.14$, $MSE = 255.46ms$, $p = 0.87$). In the same vein as for Exp 1, confidence seems oblivious to the large delays induced by the spatial orienting of attention (i.e. pre-cue vs. endogenous/exogenous conditions). The significant effect of confidence suggests that in Exp 2, participants were on average able to introspect a small portion of their latency. However, and importantly, we found no evidence for an interaction.

Regarding overall precision (fig. 7C), we found no significant main effects of condition ($F(2,103) = 1.40$, $MSE = 4.11$, $p = 0.251$), but an effect of confidence ($F(2,103) = 164.96$, $MSE = 0.75$, $p < 0.001$). As for Exp 1, we observed a condition $\times$ confidence interaction ($F(2,103)=7.83$, $MSE=0.75$, $p < 0.001$). These results also replicated Exp 1 findings that trials labelled with higher confidence were associated with higher overall precision compared to lower confidence

Figure 7: Metacognition of latency and overall precision (Exp 2). (A) Plot of the two distributions of ‘higher’ (purple) and ‘lower’ (grey) confidence for a representative participant in one of the three conditions. (B) The average difference in latency between ‘lower’ and ‘higher’ confidence trials for each condition. The absence of a significant difference between conditions confirms the latency blindness observed in Exp 1. (C) The average difference in overall precision between ‘higher’ and ‘lower’ confidence trials. Positive values confirm that confidence has access to the variability of response

precision for all conditions, despite the significant metacognitive cost for endogenous compared to pre-cue and exogenous conditions. Coloured dots correspond to individual participants in the given condition. Large dots with a black outline represent the mean across participants. Error bars represent across participants ± 1 SEM.

trials. Bonferroni-corrected t-tests ($\alpha=0.017$, corrected for 3 tests) confirmed a greater precision for higher confidence trials for both pre-cue ($T(31) = 871$, $p < 0.001$), exogenous ($T(35) = 660$, $p < 0.001$) and endogenous ($T(37) = 735$, $p < 0.001$) conditions. Corrected paired tests on the difference in concentration between higher and lower confidence trials ($\alpha=0.017$) also showed a significant difference between the pre-cue and the endogenous condition ($T(68) = 871$, $p = 0.002$), and, in contrast to Exp 1, the difference between the exogenous and the endogenous condition also reached significance level ($T(72) = 909$, $p = 0.015$). As for Exp 1, we found no evidence for a difference between pre-cue and exogenous condition ($T(66) = 700$, $p = 0.13$). Together these results show that confidence computation has access to the overall precision of responses. Furthermore, they show that metacognitive ability is significantly greater for the pre-cue and exogenous conditions compared to endogenous condition, confirming the interaction between the deployment of voluntary attention and overall metacognitive ability found in Exp 1.

3.3.3. Trial-by-trial variability of precision

As for Exp 1, we performed an additional analysis to probe the potential variability in precision from trial to trial, which might not be detectable from the overall empirical distribution of errors. We found evidence for a between-trials variability in encoding precision in the pre-cue group, and to a less extent in the exogenous group (see Supplementary Material). We did not find any evidence for such variability in the endogenous group. This suggests that despite overall response variability being seemingly matched between conditions, the trial-by-trial fluctuation of internal precision level might be differentially distributed between the pre-cue/exogenous and the endogenous conditions (see Section 4. below for a discussion).

3.3.4. How well does confidence predict pair-wise error difference?

Using the same approach as for Exp 1, we employed a logistic regression model to estimate how the difference in error magnitude within a pair of trials was predictive of confidence (fig. 8). We compared the three model candidates used previously (see Section 2.2.2 and 2.3.1.4).

Figure 8: Confidence judgment and the difference in error magnitudes (Exp 2). (A) The probability of higher confidence for trial B in the pair, as a function of the error difference between the two trials ($\Delta\epsilon$) at the group level. Errors have been grouped by quantiles at the individual level, and the group average is represented with error bars for each quantile. (B) The average correlation coefficient (or slope) for each condition, used as a measure of metacognitive ability. Coloured dots correspond to individual participants in the given condition. Large dots with a black outline represent the mean across participants. Error bars represent across participants ± 1 SEM.

The total log-likelihood across participants and conditions favoured the centred model (simple model: -5951; centred model: -5828; scaling model: -5964). Once again, a model comparison provided evidence for the 'centred' model compared to both the 'simple' model (LLR = 246.45) and the 'scaling' model (LLR = 271.18). We also found greater likelihood for the 'simple' model compared to the 'scaling' model (LLR = 24.73). The 'centred' model was kept for all subsequent analyses. The average classification accuracy for this model was nearly the same as for Exp 1, namely 63% across participants and conditions (overall AUC M = 0.63, 95% CI [0.62 0.64]; pre-cue AUC M = 0.65, 95% CI [0.63 0.67]; exogenous AUC M = 0.64, 95% CI [0.62 0.66]; endogenous AUC M = 0.60, 95% CI [0.58 0.62]). The pseudo- R^2 (McFadden, 1973) for this model was on average, across participants and conditions, equal to 0.06 (95% CI [0.051 0.069]).

Following our approach in Exp 1, we considered the slope of the model (β) as a measure of metacognitive ability. After Bonferroni correction ($\alpha = 0.017$), the slope was significantly positive for the three conditions, pre-cue ($t(31) = 10.33$, $p < 0.001$; $M = 0.006$, 95% CI [0.005 0.007]; $BF_{10} = 6.30 \times 10^8$), exogenous ($t(35) = 10.59$, $p < 0.001$; $M = 0.005$, 95% CI [0.004 0.006]; $BF_{10} = 4.60 \times 10^9$) and endogenous condition ($T(37) = 735$, $p < 0.001$), showing above chance metacognitive ability for all conditions at the group level. A repeated-measures ANOVA showed a significant effect of condition ($F(2,103) = 3.71$, $MSE = 0.00$, $p = 0.03$). After Bonferroni-correction ($\alpha = 0.05/3$), we once again found evidence for a difference between the pre-cue and the endogenous conditions ($T(68) = 851$, $p < 0.001$). However, we found no evidence for a difference between pre-cue and exogenous conditions ($t(59.1) = 1.63$, $p = 0.10$; $M_{\text{Pre-cue}} = 0.006$, $M_{\text{Exo}} = 0.005$, 95% CI [0 0.01]; $BF_{10} = 0.78$) and no evidence for a difference

between exogenous and endogenous conditions ($T(72) = 840$, $p = 0.09$). Therefore, contrary to the significant difference found in metacognition of overall precision (Section 3.3.2), here the difference between exogenous and endogenous conditions did not reach significance.

We finally tested the intercept of the model, a proxy for the metacognitive bias in favour of the second trial in the pair. Contrary to Exp 1, we found no evidence for a bias in neither the pre-cue condition ($t(31) = -1.37$, $p = 0.18$; $M = -0.15$, 95% CI $[-0.36 -0.07]$; $BF_{10} = 0.44$), nor the exogenous condition ($T(35) = 226$, $p = 0.09$), but a significant bias in the endogenous conditions ($t(37) = -3.04$, $p = 0.004$; $M = -0.31$, 95% CI $[-0.52 -0.10]$; $BF_{10} = 8.52$). However the ANOVA showed no significant effect of condition ($F(2,103) = 1.06$, $MSE = 7.09$, $p = 0.35$). Interestingly, the prevalence of a bias in the pre-cue condition in Exp 1 was not found in Exp 2, where a bias was, on the contrary, detectable for the endogenous condition. This difference in bias (pre-cue only for Exp 1, and endogenous only for Exp 2), despite the effect on the beta being similar in the two experiments, rules out the role of bias in solely driving the observed difference in metacognitive ability (i.e. in β parameter).

4. DISCUSSION

Our results shed light on three aspects of attention orientation that we will discuss in turn. The first is the time taken by the orienting process, which appeared to be different across conditions but relatively constant within each condition. The second element is the seemingly oversight of delay: metacognition appeared to be strongly bound to attention, to the point of making confidence blind to the temporal cost of attentional deployment. Yet, such a confidence-attention bound still allowed confidence to discriminate between different levels of error magnitude. Finally, confidence judgments appeared to be degraded *during* (or immediately after) the deployment of voluntary attention, revealing a novel constraint on the metacognitive system.

4.1. TIMING VOLUNTARY AND INVOLUNTARY ATTENTION, AND THEIR EFFECT ON TEMPORAL SELECTION

Our data fit well with the results from both the attention and time perception literature. First, our results replicated previous studies that have found exogenous and endogenous attention to affect the perceived phase of moving clocks (Carlson et al., 2006; Chakravarthi & VanRullen, 2011; Hogendoorn et al., 2010). These results are also consistent with the observation that the reported time of visual events is directly affected by their relative distance from the attentional locus (Jovanovic & Mamassian, 2019, 2020).

A canonical experimental design in the spatial attention literature involves a location-specific pre-cue followed by a target, with a cue-to-target delay known to maximise attentional effects and target discrimination (e.g., 300ms for endogenous attention). These paradigms however can overlook the variability of the orienting process from trial to trial: sometimes attention is allocated earlier, sometimes later (Hogendoorn et al., 2010; Zivony & Eimer, 2020). In our paradigm, the orientation of attention is expected to occur either at the beginning of the trial (pre-cue condition) or at the very moment the observer needs to register the phase of the clock (exogenous and endogenous condition).

The current experimental design allowed us to measure the time cost of attentional deployment for the exogenous and endogenous condition compared to the pre-cue (or

‘sustained attention’) condition: we found that on average the reported phase was strongly delayed, in accordance with the known average latency for each attention type. Importantly, even if our paradigm led to very reasonable delay estimations for exogenous (101ms in Exp 1 and 135ms in Exp 2) and endogenous (262ms and 294ms) attention, the absolute value of these mean errors in milliseconds is not necessary directly interpretable, as the use of temporally autocorrelated stimuli is known to increase perceived lag compared to decorrelated ones (Callahan-Flintoft, Holcombe, & Wyble, 2020; Sheth, Nijhawan, & Shimojo, 2000). In other words, temporally autocorrelated features widen the latency of information sampling, leading to less moments being sampled in a given time range. It is mainly for this reason that they should be interpreted relatively to the pre-cue condition, where attention is pre-allocated at the correct location.

On the other hand, we found no strong evidence for a difference between conditions regarding the overall variability of temporal selection. Notably, average responses precision in the pre-cue and endogenous condition was approximately matched in both experiments, suggesting that the deployment of attention takes a relatively fixed time in this task, and that the precision we measured mostly reflects the average quality of encoding. This allowed for a systematic analysis of confidence with approximately equated average performance across attention conditions.

In the between-subjects design of Exp 2, we found evidence for a difference in the cognitive process involved when comparing the pre-cue/exogenous condition to the endogenous condition. In the latter, errors appeared more homogeneous, leading to a very stable encoding precision between trials. On the contrary, for the pre-cue and exogenous conditions, the overall precision was better explained with a mixture of fluctuating encoding precisions at the trial level, the nature of the mixture being, however, equivocal (see Supplementary Material, and fig. S6 for a schematic diagram). Such a difference might be partly explained by the sustained nature of the pre-cue condition, which could favour spontaneous fluctuations of the overall attentional level, but the difference between exogenous and endogenous attention remains interesting.

4.2. METACOGNITION IGNORES ATTENTIONAL LATENCY BUT TRACKS THE PRECISION OF TEMPORAL SELECTION

Participants appeared fully oblivious to the delay of both exogenous and endogenous attention (fig. 3B & 7B). In Exp 1 there was no difference in average delays regarding confidence. In Exp 2, there was a small but consistent subjective latency difference (10ms on average) but importantly this difference was the same across conditions. In comparison, both experiments showed a strong difference in objective report latency between conditions (fig 2B & 6B). For metacognition to be sensitive to attentional delay, one should have observed at least a proportional increase of both objective and subjective latencies. Therefore, the constant delay found in Exp 2 (fig. 7B)– also present during sustained attention (i.e. in the pre-cue condition) – is more likely to be related to non-attentional processing, like the general arousal level.

The inability to monitor the delay of spatial attention mirrors what has been recently found for temporal attention (Recht et al., 2019), as well as previous results on executive control (Corallo, Sackur, Dehaene, & Sigman, 2008; Marti, Sackur, Sigman, & Dehaene, 2010). These studies, using dual-task paradigms, showed that metacognition ignores the delay in temporal attention induced by the Attentional Blink (Recht et al., 2019) or the delay in response times induced by the Psychological Refractory Period (Corallo et al., 2008; Marti et al., 2010).

The present study provides evidence for a new ‘metacognitive blind spot’, namely, the delay of spatial attention. It also confirms that the introspective blindness about the timing of certain cognitive processes is generalisable to single task paradigms, thought to involve lower cognitive load.

However, the inability to monitor the average timing of cognitive events needs not preclude a fine-grained introspection of other aspects of processing, like the variability of temporal selection (or relative error magnitude). Participants in our study were able to discriminate between error magnitudes, giving higher confidence judgments to more precise trials for all conditions (fig. 3C & 7C, and fig. 4 & 8). They were also able to discriminate between larger and shorter reaction times in our RT task of Exp 1. We note that in a similar vein, observers have been shown to be metacognitively aware of some of the processing stages during visual search and implicit spatial shifts of attention (Reyes & Sackur, 2014, 2017).

4.3. ALLOCATING VOLUNTARY ATTENTION IN SPACE WEAKENS METACOGNITION

Voluntary attention takes time to be allocated, and requires cognitive control to be maintained (Carrasco, 2011). Our results are suggestive of a metacognitive cost of endogenous attention orientation in space. During voluntary orienting, the sensitivity of confidence to overall precision drops compared to pre-allocated, sustained attention (Exp 1 and 2), and compared to transient, exogenous orienting of attention (Exp 2). The weakening was confirmed in Experiment 2, where the visual characteristics of the pre-cue and endogenous conditions were strictly matched, ruling out a feature-based explanation. The cost was also observable when analysing confidence sensitivity to the within-pair difference in error magnitude for both experiments: a comparable change in confidence required a much larger error difference in the endogenous compared to the pre-cue condition (fig. 4 & 8). In both experiments, a metacognitive cost is found despite overall selection variability remaining relatively spared by attentional orienting. Therefore, it is possible that a bifurcation occurs at some stage between the stream of evidence used for perceptual report and the evidence used for metacognitive judgment. The relationship between first-order (here, phase reproduction) and second-order (here, confidence judgment) decisions is the subject of ongoing debates: the account of confidence using only the first-order decision evidence (e.g. Kiani & Shadlen, 2009) is challenged by empirical dissociations between subjective and objective performance (e.g. Rahnev et al., 2011; Recht et al., 2019) and the existence of changes of mind (e.g. Resulaj, Kiani, Wolpert, & Shadlen, 2009). At first glance, the ability to assess the magnitude of self-made errors might seem paradoxical: if the participant knows about the size of the estimation error, why not correct for this error in the first place? Such comment has often been made in the context of error detection, but also in confidence studies (Yeung & Summerfield, 2012): an influential account solves this contradiction by distinguishing between the evidence used for perceptual decision and the evidence used by metacognition (see Fleming & Daw, 2017 for a review). Our results support such distinction, and suggest that the process of deploying voluntary attention has a direct effect on confidence. Interestingly, participants’ ability to introspect on their reaction times in Exp 1 did not particularly correlate with their metacognition of temporal selection, suggesting that each type of metacognitive ability was driven by (partially) distinct sources of second-order evidence.

Yet, the results from our model comparison also raise the interesting possibility that first-order evidence itself might have a different flavour during – or immediately following –

endogenous orienting, despite overall performance being relatively spared. When comparing models for Exp 2 (see Supplementary Material), we found evidence for a trial-to-trial variability (or ‘mixture’) of internal precision levels in the pre-cue and exogenous groups, but no evidence for variable precision in the endogenous group. Regarding the pre-cue group, sustained attention could facilitate error discriminations precisely via spontaneous fluctuations in attentional focus, leading to more marked error types. However, we also found that metacognition was lower in endogenous compared to exogenous group (Exp 2), with evidence for a mixture of precision levels during exogenous orienting as well, and despite of the latter being transient. These elements suggest that it is the voluntary nature of the orienting process that impairs metacognitive ability, rather than sustained attention solely boosting it.

The current finding that the process of allocating endogenous attention elicited a metacognitive impairment might also suggest that the top-down, frontal mechanisms needed for both voluntary orienting of attention and metacognition could share certain central resources. For example, the neuroanatomical and functional bases of visual attention have been located within a large fronto-parietal network involving, amongst other areas, the frontal-eye-field (Buschman & Kastner, 2015), while the neural bases of visual metacognition are proposed to be mostly residing within the dorsolateral and anterior parts of the prefrontal cortex (Fleming & Dolan, 2012; Fleming, Ryu, Golfinos, & Blackmon, 2014; Fleming, Van Der Putten, & Daw, 2018; Shekhar & Rahnev, 2018). All of these regions have a strong implication in top-down cognitive control, biasing incoming signals from early visual cortices and monitoring perceptual selection and decision-making (Gilbert & Li, 2013; Rahnev, 2017). Further work is needed to address how the timing of attention and metacognition interact at the functional level, to better understand the neural underpinnings of the metacognitive cost observed in the present study.

5. CONCLUSION

Metacognition allows individuals to reflect on the quality of their perceptual decisions. Yet, our results demonstrate that metacognition can be oblivious to the latency of spatial attention, an important modulator of perceptual accuracy. Furthermore, our experiments tap into the computational limitations of metacognition: the very process of voluntarily deploying attention in space seems to weaken metacognition of temporal selection. Together, our results provide critical information to our understanding of metacognition and its relationship with spatial attention.

DATA AVAILABILITY

The data for the experiments are freely available via Open Science Framework:
<https://osf.io/ynm3k>

ACKNOWLEDGEMENTS

We would like to thank Monica Barbir for helpful comments during the drafting of this manuscript. This research was supported by the Agence Nationale de la Recherche (ANR-16-CE28-0002 and ANR-16-ASTR-0014 to VdG, ANR-17-EURE-0017 and ANR-18-CE28-0015 to PM) and by ENS-PSL University (doctoral scholarship to SR).

CONTRIBUTIONS

SR, VdG and PM designed the experiment. SR conducted the experiment. SR, VdG and PM analyzed the data and wrote the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

REFERENCES

- Aguilar-Lleyda, D., Lemarchand, M., & de Gardelle, V. (2020). Confidence as a Priority Signal. *Psychological Science*, 31(9). <https://doi.org/10.1177/0956797620925039>
- Balsdon, T., Wyart, V., & Mamassian, P. (2020). Confidence controls perceptual evidence accumulation. *Nature Communications*, 11(1), 1753. <https://doi.org/10.1038/s41467-020-15561-w>
- Barthelmé, S., & Mamassian, P. (2009). Evaluation of Objective Uncertainty in the Visual System. *PLoS Computational Biology*, 5(9), e1000504. <https://doi.org/10.1371/journal.pcbi.1000504>
- Barthelmé, S., & Mamassian, P. (2010). Flexible mechanisms underlie the evaluation of visual confidence. *Proceedings of the National Academy of Sciences*, 107(48), 20834–20839. <https://doi.org/10.1073/pnas.1007704107>
- Buschman, T. J., & Kastner, S. (2015). From Behavior to Neural Dynamics: An Integrated Theory of Attention. *Neuron*, 88(1), 127–144. <https://doi.org/10.1016/j.neuron.2015.09.017>
- Callahan-Flintoft, C., Holcombe, A. O., & Wyble, B. (2020). A delay in sampling information from temporally autocorrelated visual stimuli. *Nature Communications*, 11(1), 1852. <https://doi.org/10.1038/s41467-020-15675-1>
- Carlson, T. A., Hogendoorn, H., & Verstraten, F. A. J. (2006). The speed of visual attention: What time is it? *Journal of Vision*, 6, 1406–1411. <https://doi.org/10.1167/6.12.6>
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*. <https://doi.org/10.1016/j.visres.2011.04.012>
- Chakravarthi, R., & VanRullen, R. (2011). Bullet trains and steam engines: Exogenous attention zips but endogenous attention chugs along. *Journal of Vision*. <https://doi.org/10.1167/11.4.1>
- Corallo, G., Sackur, J., Dehaene, S., & Sigman, M. (2008). Limits on Introspection. *Psychological Science*. <https://doi.org/10.1111/j.1467-9280.2008.02211.x>
- de Gardelle, V., Le Corre, F., & Mamassian, P. (2016). Confidence as a Common Currency between Vision and Audition. *PLoS One*, 11(1), e0147901. <https://doi.org/10.1371/journal.pone.0147901>
- de Gardelle, V., & Mamassian, P. (2014). Does Confidence Use a Common Currency Across Two Visual Tasks? *Psychological Science*, 25(6), 1286–1288. <https://doi.org/10.1177/0956797614528956>
- de Gardelle, V., & Mamassian, P. (2015). Weighting Mean and Variability during Confidence Judgments. *Plos One*, 10(3), e0120870. <https://doi.org/10.1371/journal.pone.0120870>
- de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47(1), 1–12. <https://doi.org/10.3758/s13428-015-0600-0>

014-0458-y

- Denison, R. N., Adler, W. T., Carrasco, M., & Ma, W. J. (2018). Humans incorporate attention-dependent uncertainty into perceptual decisions and confidence. *Proceedings of the National Academy of Sciences*, 201717720. <https://doi.org/10.1073/pnas.1717720115>
- Desender, K., Boldt, A., & Yeung, N. (2018). Subjective Confidence Predicts Information Seeking in Decision Making. *Psychological Science*.
<https://doi.org/10.1177/0956797617744771>
- Fechner, G. T. (1964). *Elemente der Psychophysik*. 2 (1964). Bonset.
- Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision performance: A general Bayesian framework for metacognitive computation. *Psychological Review*, 124(1), 1–59.
<https://doi.org/10.1037/rev0000045>
- Fleming, S. M., & Dolan, R. J. (2012). The neural basis of metacognitive ability. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 367(1594), 1338–1349. <https://doi.org/10.1098/rstb.2011.0417>
- Fleming, S. M., Ryu, J., Golfinos, J. G., & Blackmon, K. E. (2014). Domain-specific impairment in metacognitive accuracy following anterior prefrontal lesions. *Brain*.
<https://doi.org/10.1093/brain/awu221>
- Fleming, S. M., Van Der Putten, E. J., & Daw, N. D. (2018). Neural mediators of changes of mind about perceptual decisions. *Nature Neuroscience*. <https://doi.org/10.1038/s41593-018-0104-6>
- Gilbert, C. D., & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews Neuroscience*. <https://doi.org/10.1038/nrn3476>
- Gorea, A., Mamassian, P., & Cardoso-Leite, P. (2010). Introspective duration estimation of reactive and proactive motor responses. *Acta Psychologica*, 134(2), 142–153.
<https://doi.org/10.1016/j.actpsy.2010.01.005>
- Guggenmos, M., Wilbertz, G., Hebart, M. N., & Sterzer, P. (2016). Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *ELife*.
<https://doi.org/10.7554/eLife.13388>
- Hainguerlot, M., Vergnaud, J. C., & De Gardelle, V. (2018). Metacognitive ability predicts learning cue-stimulus associations in the absence of external feedback. *Scientific Reports*.
<https://doi.org/10.1038/s41598-018-23936-9>
- Hensher, D. A., & Stopher, P. R. (2021). Behavioural Travel Modelling. In *Behavioural Travel Modelling* (pp. 11–52). Routledge. <https://doi.org/10.4324/9781003156055-2>
- Hogendoorn, H., Carlson, T. A., VanRullen, R., & Verstraten, F. A. J. (2010). Timing divided attention. *Attention, Perception, and Psychophysics*, 72(8), 2059–2068.
<https://doi.org/10.3758/APP.72.8.2059>
- Jovanovic, L., & Mamassian, P. (2019). When an Event Is Perceived Depends on Where We Attend. *I-Perception*, 10(3), 1–10. <https://doi.org/10.1177/2041669519858096>
- Jovanovic, L., & Mamassian, P. (2020). Temporal context affects the perceived time of visual events. *Psychonomic Bulletin and Review*, 27, 56–61. <https://doi.org/10.3758/s13423-019-01682-x>
- Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science (New York, N.Y.)*, 324(5928), 759–764.
<https://doi.org/10.1126/science.1169405>
- Kurtz, P., Shapcott, K. A., Kaiser, J., Schmiedt, J. T., & Schmid, M. C. (2017). The Influence of

- Endogenous and Exogenous Spatial Attention on Decision Confidence. *Scientific Reports*, 7(1). <https://doi.org/10.1038/s41598-017-06715-w>
- Ma, W. J., Husain, M., & Bays, P. M. (2014, March). Changing concepts of working memory. *Nature Neuroscience*. <https://doi.org/10.1038/nn.3655>
- Mamassian, P. (2016). Visual Confidence. *Annual Review of Vision Science*, 2(1), 459–481. <https://doi.org/10.1146/annurev-vision-111815-114630>
- Marti, S., Sackur, J., Sigman, M., & Dehaene, S. (2010). Mapping introspection's blind spot: Reconstruction of dual-task phenomenology using quantified introspection. *Cognition*. <https://doi.org/10.1016/j.cognition.2010.01.003>
- McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. Retrieved from <http://elsa.berkeley.edu/reprints/mcfadden/zarembka.pdf>
- Morey, R. D., & Rouder, J. N. (2018). BayesFactor: Computation of Bayes Factors for Common Designs. Retrieved from <https://cran.r-project.org/package=BayesFactor>
- Peirce, J. W. (2007). PsychoPy-Psychophysics software in Python. *Journal of Neuroscience Methods*, 162(1–2), 8–13. <https://doi.org/10.1016/j.jneumeth.2006.11.017>
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*. <https://doi.org/10.1080/00335558008248231>
- R Core Team. (2013). R: A Language and Environment for Statistical Computing. Vienna, Austria. Retrieved from <http://www.r-project.org/>
- Rahnev, D., Maniscalco, B., Graves, T., Huang, E., de Lange, F. P., & Lau, H. (2011). Attention induces conservative subjective biases in visual perception. *Nature Neuroscience*, 14(12), 1513–1515. <https://doi.org/10.1038/nn.2948>
- Rahnev, Dobromir. (2017). Top-Down Control of Perceptual Decision Making by the Prefrontal Cortex. *Current Directions in Psychological Science*. <https://doi.org/10.1177/0963721417709807>
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary Suppression of Visual Processing in an RSVP Task: An Attentional Blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18(3), 849–860. <https://doi.org/10.1037/0096-1523.18.3.849>
- Recht, S., Mamassian, P., & de Gardelle, V. (2019). Temporal attention causes systematic biases in visual confidence. *Scientific Reports*, 9(1). <https://doi.org/10.1038/s41598-019-48063-x>
- Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, 461(7261), 263–266. <https://doi.org/10.1038/nature08275>
- Reyes, G., & Sackur, J. (2014). Introspection during visual search. *Consciousness and Cognition*. <https://doi.org/10.1016/j.concog.2014.08.009>
- Reyes, G., & Sackur, J. (2017). Introspective access to implicit shifts of attention. *Consciousness and Cognition*, 48, 11–20. <https://doi.org/10.1016/j.concog.2016.10.003>
- Shekhar, M., & Rahnev, D. (2018). Distinguishing the roles of dorsolateral and anterior PFC in visual metacognition. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.3484-17.2018>
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*. <https://doi.org/10.1126/science.3629243>
- Sheth, B. R., Nijhawan, R., & Shimojo, S. (2000). Changing objects lead briefly flashed ones. *Nature Neuroscience*. <https://doi.org/10.1038/74865>

- Van Den Berg, R., Shin, H., Chou, W. C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America*.
<https://doi.org/10.1073/pnas.1117465109>
- Van den Berg, R., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert, D. M. (2016). Confidence Is the Bridge between Multi-stage Decisions. *Current Biology*.
<https://doi.org/10.1016/j.cub.2016.10.021>
- Wilimzig, C., Tsuchiya, N., Fahle, M., Einhäuser, W., & Koch, C. (2008). Spatial attention increases performance but not subjective confidence in a discrimination task. *Journal of Vision*, 8(5), 7. <https://doi.org/10.1167/8.5.7>
- Yeung, N., & Summerfield, C. (2012). Metacognition in human decision-making: confidence and error monitoring. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 367(1594), 1310–1321. <https://doi.org/10.1098/rstb.2011.0416>
- Zivony, A., & Eimer, M. (2020). Distractor Intrusions Are the Result of Delayed Attentional Engagement: A New Temporal Variability Account of Attentional Selectivity in Dynamic Visual Tasks. *Journal of Experimental Psychology: General*.
<https://doi.org/10.1037/xge0000789>
- Zizlsperger, L., Sauvigny, T., & Haarmeier, T. (2012). Selective attention increases choice certainty in human decision making. *PLoS ONE*.
<https://doi.org/10.1371/journal.pone.0041136>
- Zizlsperger, L., Sauvigny, T., Händel, B., & Haarmeier, T. (2014). Cortical representations of confidence in a visual perceptual decision. *Nature Communications*.
<https://doi.org/10.1038/ncomms4940>