Title

The foundations of introspective access: how the relative precision of target encoding influences metacognitive performance

Authors & Affiliations

Sanne Kellij¹

Johannes Fahrenfort^{2,3,4}

Hakwan Lau^{5,6,7,8}

Megan A. K. Peters^{9,10,11}

Brian Odegaard⁵

¹ Sociology department, Rijksuniversiteit Groningen, Netherlands

² Department of Experimental and Applied Psychology, Institute Brain and Behavior Amsterdam (iBBA), VU University Amsterdam, Amsterdam, Netherlands

³ Department of Psychology, University of Amsterdam, Amsterdam, Netherlands

⁴ Amsterdam Brain and Cognition (ABC), University of Amsterdam, Amsterdam, Netherlands

⁵ Department of Psychology, University of California Los Angeles, Los Angeles, CA 90095- 1563

⁶ Brain Research Institute, University of California Los Angeles, Los Angeles, CA 90095-1563

⁷ Department of Psychology, University of Hong Kong, Pokfulam Road, Hong Kong

⁸ State Key Laboratory of Brain and Cognitive Sciences, University of Hong Kong, Pokfulam Road, Hong Kong

⁹ Department of Bioengineering, University of California Riverside, Riverside, CA 92521

¹⁰ Department of Psychology, University of California Riverside, Riverside CA 92521

¹¹ Neuroscience Interdepartmental Graduate Program, University of California Riverside, Riverside CA 92521

Keywords

Metacognition, Signal Detection Theory, Attentional Blink, Phase Scrambling

Abstract

In metacognition research, objective performance is often matched between conditions, so that differences in type II performance (metacognition) cannot be attributed to differences in type I (task-related) performance. A recent study by Meuwese et al (2014) showed a difference in metacognitive ability when comparing metacognition in a detection and a categorization task. Interestingly, this difference only occurred in correct rejection trials, despite the fact that type I performance was matched between tasks. In the present study, we use a different manipulation to investigate the origin of these metacognitive asymmetries. We measured metacognition while matching performance in two visibility manipulations: phase-scrambling and the attentional blink. As in previous work, metacognitive asymmetries emerged: despite matched type 1 performance, metacognitive ability (measured by area under the ROC curve) for reporting stimulus absence was higher in the attentional blink condition. We performed Signal Detection Theoretic (SDT) modeling of the results, showing that differences in metacognition under equal type I performance can be explained when the variance of the signal and noise distributions are unequal. Specifically, the present study suggests that phase scrambling signal trials have a wider distribution (more variability) than attentional blink signal trials, leading to a larger area under the ROC curve for attentional blink trials where subjects reported stimulus absence. These results provide a theoretical basis for the origin of metacognitive differences on trials where subject report stimulus absence, and may also explain previous findings where the absence of evidence during detection tasks results in lower metacognitive performance when compared to categorization.

Introduction

At any given moment in our daily lives, we have varied amounts of confidence in the things we see. For example, as we drive, we may primarily be concerned with *detecting* objects in the scene which might interfere with the path of our vehicle. Potentially, there may be different reasons why we fail to detect something: sometimes we may not detect an object because our perception is degraded (due to fog, low light, or other factors), and other times we may not detect an object due to lapses in our attention. These different scenarios involving judgments of object presence or absence open an interesting question: is our subjective sense of confidence in making judgments in these two scenarios roughly equivalent, or does it differ? And if it does differ, why might this be the case?

In the present study, we aimed to investigate individual subjects' metacognitive abilities in scenarios which manipulated stimulus visibility in different ways. In general, metacognition is defined as the ability to think or reflect about one's perceptions, thoughts and actions (Metcalf, 1994; Grimaldi, Lau & Basso, 2015) and is often measured by means of confidence ratings (Sinanaj, Cojan & Vuilleumier, 2015). In this investigation, we explored the nature of confidence ratings using a detection task with two conditions that reduced subjects' awareness of the stimuli using different methods to determine how the encoding of both signal (i.e., target presence) and noise (i.e., target absence) information influences these judgments.

Recent related work by Meuwese, van Loon, Lamme, & Fahrenfort (2014) investigated the subjective experience of object recognition by comparing metacognition for two tasks: object detection and object categorization. On each trial in the detection task, subjects answered a question about whether an animal was present or not in the image shown; in the categorization task, subjects answered a question about whether the animal was from a specific target category; in both tasks, subjects also gave confidence ratings from 1 (not at all confident) to 6 (very confident) about the correctness of their response. Metacognitive ability was assessed by plotting the receiver operating characteristic curves and computing the area under the curve (AUC) independently for each task-stimulus combination. It was shown that overall, metacognitive ability was roughly equivalent for masked and degraded stimuli, but markedly higher for responses in the categorization task compared to the detection task. This difference in metacognitive ability between the two tasks was driven primarily by trials in which subjects reported *target absence*; specifically, on correct rejection trials. Metacognitive ability on trials where subjects reported seeing the target stimulus was approximately equal across tasks, and was comparable for all other stimulus-response contingencies (misses, FAs, and hits). Thus, the authors concluded that metacognitive ability is better in tasks that rely on "positive evidence" to perform the perceptual judgment (e.g., categorization), compared to tasks that rely on assessment of the absence of evidence (e.g., detection) (Luu & Stocker, 2018; Maniscalco, Peters, & Lau, 2016; Zylberberg, Barttfeld, & Sigman, 2012).

We hypothesized that the results of Meuwese and colleagues (2014) may possibly be explained by a violation of the assumption of equal variance in Signal Detection Theory (SDT). Specifically, in SDT it is assumed that on any given trial the subject's nervous system represents the presence of either the nontarget (e.g., noise) stimulus or target stimulus as a random sample from a noisy representation of the information, and compares this sample to an internal criterion to make a

perceptual judgment about stimulus presence or absence. The most common analysis in this framework assumes that the variances of the distributions representing the nontarget and target stimuli are equal (which is an assumption that is sometimes impossible to evaluate, as studies which do not record confidence reports do not have multiple points on the zROC function to assess variance asymmetry). When the variances are equal, standard measures of performance such as the area under the ROC curve (AUC) are unbiased. However, if the distributions representing the nontarget and target stimuli are *unequal*, as is often the case in detection tasks (Green & Swets, 1966; Macmillan & Creelman, 2005), then response-specific AUC measures are strongly influenced by interactions between the variance of the signal distribution and the placement of the response criterion, which can greatly influence results.

To investigate the possibility that unequal variances may be driving differences in this measure of metacognitive ability (i.e., the AUC), we performed simulations and an experiment in which we aimed to assess the metacognitive ability of participants using a detection task with two conditions to manipulate awareness of stimuli: (1) an attentional blink (AB) paradigm and (2) a phase scrambling (PS) paradigm. These two conditions were chosen based on research by Kanai, Walsh, & Tseng (2010), which showed that these manipulations cause the same discrepancy in metacognitive ability for 'No' responses (specifically, correct rejections), but similar metacognitive ability for 'Yes' responses. In line with the results of Kanai et al. (2010), we hypothesized that there would be no differences in metacognitive ability (response-conditional type 1 AUC values) for the 'Yes' response trials and no difference on average over all trials, but that there *would* be a difference for the 'No' response trials. Specifically, we hypothesized that during AB trials, participants would be more aware of when they missed a stimulus and therefore be better at categorizing metacognitive judgements, which would not happen as often during PS trials. Finally, we hypothesized that this difference in metacognition would also be evident in a violation of the equal variance assumption of SDT in the experimental data. To anticipate, all of these hypotheses were supported by both our simulations and experimental results.

Method

Participants

In this study, 38 students at the University of California-Los Angeles (*Age range* = 18-29, *Mean age* = 20.29, *SD* = 1.99, 11 men) participated for partial fulfillment of course requirements. All participants had normal or corrected-to-normal vision, no history of neurological conditions, and gave written informed consent before participation. The experiment was approved by the UCLA Institutional Review Board. For the analyses, three participants (subject #3, #8, and #18) were excluded because they did not use the entire confidence rating scale as was instructed (e.g., if subjects use only one confidence rating for all trials, we cannot assess how confidence differs across our two conditions). In total, 35 students were included in the final analyses (*Age range*: 18 to 29, *M* = 20.31, *SD* = 2.05, 9 men).

Stimuli

The stimuli used in this experiment were from Meuwese et al. (2014), which

included 600 grayscale animal photos from three categories: birds (200), cats (200) and fish (200). These stimulus categories were matched for low-level image statistics. For each image containing an animal, we created a non-animal-counterpart by fully randomizing the phase information in the image, effectively removing all object information (coherence) from the image (e.g., compare the leftmost and rightmost image in Figure 1 to see the difference between an animal and its non-animal counterpart).

Phase-scrambling was also used to parametrically manipulate visibility in the PS task. The phase scrambled versions of these stimuli had similar low-level image statistics (e.g., overall contrast), but object information was gradually removed with increased levels of phase scrambling (see Fig. 1, left to right).

Figure 1. Examples of phase scrambling levels. This figure shows the spectrum of the amounts of phase scrambling used in the experimental task, from 0-100%. The rightmost image is the image that was used as the non-animal counterpart in all experimental tasks.

In both tasks, all distractor stimuli were gray and all target stimuli had a blue filter over the photo. Stimuli contained either animals or noise (i.e., a 100% phase scrambled photo) and were shown for 50 ms (see Fig. 2). The blanks in between stimuli lasted for 100ms.

Figure 2. Examples of the targets and distractors for the attentional blink and phase scrambled stimuli. All blue stimuli are targets and all grayscale stimuli are distractors. AB stimuli are shown in the top row, and include examples of the first target (Target 1), examples of the second target (Target 2), and distractor stimuli. PS stimuli are shown in the second row, and include examples of the first target, examples of the second target (57.89% PS), and distractor stimuli.

Procedure

When participants came to the lab, they filled out an eligibility checklist, gave informed, written consent to participate, and then received specific instructions for the main task. The task consisted of a rapid serial visual presentation (RSVP) paradigm and included two trial types: a short lag trial with 300 ms between targets (attentional blink, or “AB”), and a long lag trial with 900 ms between targets (phase scrambled, or “PS”). Each trial started with a fixation dot for 500 ms, after which the RSVP started (Figure 3). First, three to five gray distractors were presented (the number from 3-5 was pseudorandomly selected every trial). Then, the first target (“T1”) was shown in blue. In both conditions, T1 was either fully intact if it contained an animal (non-scrambled) or 100% phase-scrambled if it did not contain an animal (i.e., if it was noise). Therefore, T1 did not differ between trial types. Next, a variable number of gray distractors was shown. After these distractors, the second blue target (“T2”) appeared in blue. For AB trials there was only one gray distractor between T1 and T2 (short lag), while for PS trials there were five gray distractors between T1 and T2 (long lag). For AB trials, T2 was either a blue non-PS picture (animal) or a blue, 100% phase-scrambled picture (no animal). For PS trials, the second target was either a blue partially-PS picture (animal), or blue 100%-scrambled picture (no animal). After the second target, five to seven gray distractors (pseudorandomly assigned using the `np.random.choice` function from the SciPy library (Jones, Oliphant, & Peterson, 2001)) appeared. Finally, the questions for the targets appeared.

Figure 3. Representative trials for the Attentional Blink (AB) and Phase Scrambled (PS) trial types. In both conditions, following 500ms of fixation, 3-5 distractors were shown, followed by the first target, which was either noise or a clear picture of an animal. After the first target, in the AB condition, one distractor was shown, immediately followed by the second target. In the PS condition, 5 distractors were shown before the second target was presented. The trial concluded with 5-7 additional distractors being shown, followed by the question prompts asking whether the first target (T1) was an animal, and confidence regarding whether the second target (T2) contained an animal. All stimuli were shown for 50 ms and all blank screens for 100 ms. The response screens were shown until a key was pressed.

Participants first had to decide whether the first target (T1) was an animal or not and press the corresponding key. Then participants had to indicate whether the second target (T2) contained an animal or not and while simultaneously indicating their confidence about that decision with 4 levels of certainty (8-point scale: 8 entirely certain not an animal, 7 most likely not an animal, 6 probably not an animal, 5 unsure

not an animal, 4 unsure an animal, 3 probably an animal, 2 most likely an animal, 1 entirely certain an animal). This response method combines the classical “Type 1” judgment (assessing whether or not the target contained an animal) with a “Type 2” judgment (evaluating the confidence one has in the judgment of whether the target contained an animal). Participants were explicitly instructed that the first response should only reflect their decision about the first target (T1) and that the second response should only reflect their decision about the second target (T2) that appeared in each sequence.

Participants began the experimental task by completing practice blocks, aimed to target equal performance for the AB and PS trial types. The first block consisted of 54 trials, including 36 AB trials and 18 PS trials, to identify an average percentage correct for the AB trials (PS trials were included so participants would not learn to expect the exact time the stimulus would appear, but were excluded from further analysis). If the performance on the AB trials was not high enough (i.e., was lower than 65% correct), the participant had to redo this block until their performance was higher than 65%. The second block (108 trials; 72 PS trials and 36 AB trials) used a staircasing procedure (QUEST) (Watson & Pelli, 1983) to find a degree of PS such that the percentage correct between the AB and PS trial types was approximately equivalent.

During the main experiment, participants completed eight blocks of 72 trials, for a total of 576 trials. Within each block, AB and PS trials were randomly interleaved. To ensure performance was matched in all blocks for T2 responses for the AB and PS trials, whenever the percentage correct on PS trials became more than five percentage points higher than the percentage correct on AB trials, the PS level was raised by about five percentage points (1/19, making 20 possible PS levels as 0% was also included). If the percentage correct on long lag trials got worse than ten percentage points below the short lag, the PS level was lowered by approximately five (1/19) percentage points.

Results

To evaluate whether the objective difficulty between trial types was successfully matched, we performed a paired-samples t-test on the percentage correct for T2 responses for AB and PS trials, analyzing the subset of trials which included a correct T1 response. (This is because in AB trials, it is easier to have the T2 response correct if T1 was unattended or ignored; thus, we aimed for consistency across trial types.) Results revealed that performance for T2 (when T1 was correct) was similar for the two trial types (AB trials: 70.00% correct [range = 54.39-81.07%, $SD = .061$], PS trials: 69.28% correct [range = 57.75-79.21%, $SD = .053$], $t(34) = 1.40$, $p = .17$). All but three participants were above 60% performance; these three participants were excluded in the following analyses.

To visualize the behavioral responses, we first plotted pseudo Type 1-ROC curves, which were formed by using each confidence rating as a decision criterion for performing the detection task and plotting the corresponding false alarm and hit percentages for each of these points (Fig. 4A). A Wilcoxon signed-rank test of the AUC values (computed using MATLAB’s “perfcure” function) between trial types revealed that the values were not significantly different from one another ($z = -1.16$, $p = .24$, Fig. 4B), although the shapes of the curves were quite different from one another. Overall, this provided further support for matched type 1 performance between the PS and AB trial types.

Figure 4. Pseudo Type-1 ROC curves and AUC values for the attentional blink (AB) and phase scrambled (PS) trials. (A) The ROC curves for each condition. Each curve plots the average data across all subjects. Each dot on the graph represents data from a different confidence rating; specifically, each dot represents a different decision criterion (for a given confidence rating) and the corresponding percentage of hits and false alarms. The black dotted line provides a reference for chance-level performance. Error bars represent SEM across subjects. (B) The average AUC values for the PS and AUC conditions. These AUC bars plot the averaged AUC values for each individual subjects' curves, with error bars denoting SEM across subjects.

Next, we analyzed whether there was a general difference for metacognitive ability between trial types using the same methods as Meuwese et al. (2014). Figure 5A shows the ability of confidence ratings to distinguish between correct and incorrect responses, using the data from all subjects. There was a significant difference for average confidence ratings to better distinguish correct and incorrect trials in AB trials than in PS trials, as shown by a Wilcoxon signed-rank test of the area under the curve (AUC) values ($z = -2.09$, $p = .036$). To explore the difference between trial types in more detail, type 2 ROC curves were made for only the 'Yes, animal present' responses (false alarm & hit trials). These curves (see Fig. 5C,5D) showed that the performance on these trials was quite similar between the PS and AB conditions ($z = 1.15$, $p = .25$). However, the type 2 ROC curves of the 'No, animal not present' responses (misses & correct rejections (CR)) (see Fig. 5E,5F) did reveal a significant difference in AUC value between the AB and PS trial types ($z = -2.00$, $p = .046$). Thus, subjects were better at using metacognitive judgments in AB trials to distinguish between correct and incorrect trials for 'No' responses compared to when they were using metacognitive judgments in PS trials.

Figure 5. Response-conditional ROC curves using across-subject averages. (A-B) The probability for each confidence judgment given either a correct or incorrect response. Correct responses were composed of the summed counts for correct rejections and hits for each confidence level, and incorrect responses were composed of the summed counts for misses and false alarms across each confidence level. The ROC curve plots the average results across all subjects, and

the AUC bars plot the averaged AUC values for each individual subjects' curves. Error bars represent the standard error of the mean across subjects. (C-D) The probability of each confidence judgment given a Yes response. The y-axis shows this probability for Yes responses that were correct (hits), and the x-axis shows this probability for the Yes responses that were incorrect (false alarms). Conventions are the same as panels A-B. (E-F) The probability of each confidence judgment given a No response. The y-axis shows this probability for no responses that were correct (correct rejections), and the x-axis shows this probability for no responses that were incorrect (misses). Conventions are the same as panels A-B. * $p < .05$.

We then sorted our data to determine if differences in confidence between the AB and PS conditions were driven by a specific Type 1 stimulus-response combination. Figure 6 displays ROC curves for each Type 1 response, including (A) hits, (B) misses, (C) false alarms, and (D) correct rejections. As can be seen in the results, a clear pattern emerges: 'yes' responses (left column) lead to a stronger confidence profile in AB compared to PS, regardless of whether the yes response is correct, whereas 'no' responses (right column) lead to under-confidence in the AB compared to PS, regardless of correctness of the response.

Figure 6. ROC curves for the PS and AB conditions for all Type-1 responses using across-subject averages. Each point on the graph represents the criterion for a given confidence rating. (A) Hits, (B) Misses, (C) False Alarms (FA), (D) Correct Rejections (CR).

Interestingly, our results conceptually replicate several findings by Meuwese and colleagues (2014). In their study, they reported significant differences due to the

type of target detection task that participants were asked to perform (i.e., categorization and detection). They showed that metacognitive ability was overall higher for categorization compared to detection (based on AUC values plotting correct vs. incorrect responses) across all trials, and that this difference was primarily driven by differences in metacognitive ability in the ‘No’ response trials (their Fig. 2). Additionally, they did not show a significant difference in metacognitive ability between tasks based solely on the ‘Yes’ response trials. The authors conclude that metacognitive ability is better in tasks that rely on positive evidence (categorization) than in tasks that rely more strongly on absence of evidence (detection).

However, the quantitative framework used by Meuwese and colleagues to analyze their behavioral results, SDT, provides a possible alternative explanation of why differences in metacognitive ability may have emerged across tasks. SDT is typically used to analyze perceptual decisions that are made in uncertain or ambiguous situations (Wickens, 2002). In this framework (Fig. 7), whenever a stimulus is presented, it is assumed that an internal response is generated which encodes the stimulus (e.g., either signal or noise), and this response represents a sample from a distribution with some amount of variability. On each trial, the observer compares the strength of the evidence they obtained to an internal criterion, and makes a judgment as to whether this stimulation originated from a signal present in the environment or random noise. For a detection task (e.g., “is ‘A’ present or not?”) the decision-making process can be represented by two curves, one for when ‘A’ is absent (e.g., a noise trial) and one for when ‘A’ is present (e.g., a signal trial). This gives four different possibilities for type 1 performance: hits, misses, false alarms, and correct rejections.

Figure 7. Signal Detection Theoretical (SDT) Framework. In SDT, it is assumed that on each trial of a detection experiment, an observer’s internal representation can be modeled as a sample drawn from either a “stimulus present” (blue) or “stimulus absent” (green) distribution. To make a decision about stimulus presence or absence, this sample is then compared to an internal criterion; if the sample is above the criterion, the participant responds “signal present,” and if the sample is below the criterion, the participant responds “signal absent.” This results in four different trial types: trials where the stimulus was present and the subject correctly responds that it

was presented (hits, light blue); trials where the stimulus was present and the subject says that it was not (misses, dark blue); trials where the stimulus was not present and the subject says that it was not (correct rejections, light green), and trials where the stimulus was not present, but the subject says that it was (false alarms, dark green).

In many studies employing SDT, it is assumed that the signal and noise distributions have equal variances. With equal variances for these two distributions, type 1 response-conditional ROC curves have a characteristic shape, as shown in our simulation in Figure 8. For example, ROC curves that plot the probability of confidence given a hit vs. the probability of confidence given a false alarm, and ROC curves that plot the probability of confidence given a correct rejection vs. the probability of confidence given a miss, will have similar shapes.

Figure 8. Simulation of a detection task with equal variances for the signal and noise distribution, using standard signal detection theoretic tools. (A) Graphical representation for stimulus encoding and perceptual decisions in this framework. Shown here are two distributions, representing the encoding of a sensory signal (blue) and the encoding of noise (black). On any given trial, either a stimulus or noise is presented, and the observer compares a sample drawn from one of these two distributions to the internal criterion (green) to make their perceptual judgment about stimulus presence/absence. (B) ROC curves for trials where subjects responded “yes, stimulus present” (left) or “no, stimulus absent” (right). The different

confidence ratings in this simulated experiment (20 total) are shown in blue. The closer the point on the blue curve is to the origin, the higher the given confidence rating (blue dots). Whenever the curve is above the black diagonal, it indicates that the metacognitive responses of the observer can distinguish correct from incorrect trials.

However, the assumption of equal variance between the signal and noise distributions can be violated. In circumstances where the variances of these two distributions differ, distinct patterns emerge for the response-conditional ROC curves for stimulus absence, which can directly influence AUC magnitude (Figure 9). Critically, violation of this assumption strongly influences the ROC curves for “signal absent” trials, illustrating the importance of evaluating this assumption before applying SDT analyses to data.

Figure 9. Theoretical simulation of a detection task with unequal variances for the signal and noise distributions. (A) Graphical representation of the type 1 decision space. Shown here is an example where the variability of encoding of the sensory signal is either lower (blue) or higher (red) than the precision of encoding the noise signal (black). The distributions differ in standard deviation (.8 & 1.43 for red and blue, respectively), but have the same mean (2.44). Perceptual decisions are made by comparing the sample drawn to an internal criterion (shown in green). (B) Response-conditional ROC Curves for simulated data for trials where the subject responded “yes, stimulus present” (hits & false alarms) or “no, stimulus absent” (misses & correct rejections). As can be seen in this example, the ROC curves are quite similar for both the low-variance and high-variance cases for trials where the

subject responses “yes, stimulus present,” indicating variance of the signal distribution has minimal impact on this measure. However, the variance of the signal distribution has a much greater influence on the trials where the subject responded “no, stimulus absent,” illustrating the importance of testing the assumption of equal variances when computing AUC values for this measure. (C) Area under the curve for the “yes” response trials (hits and false alarms) and area under the curve for the “no” response trials (correct rejections and misses).

To explicitly evaluate whether the equal variance assumption was violated for the AB and the PS trials in our experiment, Type 1 zROC analyses were performed (Figure 10). Specifically, the hit and false alarm rates of the type 1 ROC analysis were z-scored for each participant. A linear line was then fit to each participant’s data through the points in zROC space, with an estimated slope and intercept.

Figure 10. zROC space for one representative subject. Shown in this figure are the z-scored false alarm rates and hit rates for the AB condition for a representative subject, with the linear model fit to these points.

If the slopes of these lines in zROC space are near 1, this supports the idea that the equal variance assumption has been met; if the slopes of these lines are different from 1, it is likely that the equal variance assumption has been violated. However, as noted by (Wickens, 2002), “any theoretical model, once its parameters are estimated, implies a comparable table of predicted frequencies. If the observed frequencies are similar to the predicted frequencies, then the model fits, and if they differ, the model fails (p.213).” Therefore, we also performed chi-square goodness of fit tests, comparing our observed frequencies to the expected distributions under the equal-variance model and unequal variance model, and conducted a hierarchical model comparison to see if one model was significantly better than the other.

Results confirmed the hypothesis that for many subjects, both AB and PS trial types were marked by unequal variances for the signal and noise distributions (see Fig. 11). For PS trials, the slopes on average were *below* 1 (mean = 0.87, $p < .0001$), and 17 of the 35 subjects were better explained by a model that assumed unequal variances (chi-square value > 3.84 for these subjects). For AB trials, the average

slope of the zROC lines was *above* 1 (mean = 1.11, $p < .01$), and 12 of the 35 subjects were better explained by a model assuming unequal variances (chi-square value > 3.84 for these subjects).

Figure 11. Slopes of zROC lines. (A) Phase Scrambled (PS) trials. Shown in this panel are the slopes of the zROC lines for each individual subject. In general, lines that deviate from a slope of 1 (the black dotted line) indicate that the equal variance assumption has been violated. The black dot on the right hand side indicates the average slope of the zROC lines, with SEM across subjects. (B) Attentional blink (AB) trials, with the same conventions as panel A.

In a final analysis, using the estimated slopes and intercepts for the AB and PS trials, we computed the average slope and average intercept across all subjects to determine if our data supported the hypothesis of unequal variances for the signal and noise distributions for the two conditions. As can be seen in Fig. 12, the estimated variance for the distributions of both the AB and PS conditions appear to violate the assumption of equal variance, and using these values, it is clear that this scenario can produce matched ROC curves and AUC values for “yes” responses (Fig. 12B, left, Fig. 12C, left) and mismatched ROC curves and AUC responses for “no” responses (Fig. 12B, right, Fig. 12C, right). This analysis supports the hypothesis that unequal variances are driving the observed effects, and is further verified by its congruence with results from Figure 5.

Figure 12. Detection task results based on experimental data, using the estimated slopes, intercepts, and criteria, averaged across all participants. (A) Graphical representation of the decision space. Shown here are the estimated variances for the signal distributions for the PS (blue) and AB (red) conditions using the average values from the slopes and intercepts of the zROC lines for each participant. We also include the average criteria for the AB (dark green) and PS (light green) conditions (B) ROC Curves using the average zROC slope, zROC intercept, and criterion across subjects for the PS (blue) and AB (red) “yes, stimulus present” (hits & false alarms) or “no, stimulus absent” (misses & correct rejections). As can be seen in this example, the ROC curves are quite similar for both the low-variance and high-

variance cases for trials where the subject responses “yes, stimulus present,” indicating variance of the signal distribution has minimal impact on this measure. However, the variance of the signal distribution has a much greater influence on the trials where the subject responded “no, stimulus absent,” illustrating the importance of testing the assumption of equal variances when computing AUC values for this measure. (C) Area under the curve for the “yes” response trials (hits and false alarms) and area under the curve for the “no” response trials (correct rejections and misses) for the PS (blue) and AB (red) conditions.

Discussion

Previously, it has been shown that metacognitive ability is higher for responses in a categorization task compared to a detection task, and that this difference is primarily driven by performance on trials in which subjects reported that a target stimulus was not presented on a given trial (Meuwese et al., 2014). We hypothesized that this previous finding may be accounted for by violations of the assumption of equal variance in signal detection theory, and performed simulations and an experiment to replicate these results and demonstrate possible theoretical underpinnings of this reported effect. Through our modeling simulations, we find that indeed, when the variance of “signal present” distribution is different from the “signal absent” distribution, this has a minimal impact on the area under the ROC curve for trials where participants report that a target stimulus was presented, but a significant impact on the ROC curve for trials when participants report that a stimulus was *not* presented (Figure 9). In our experiment, we had subjects perform a detection task using two conditions shown to produce matched type 1 performance and different levels of metacognition on trials with reported stimulus absence (Figures 4 & 5) and find that these conditions violate the assumption of equal variance in signal detection theory, and therefore can at least partially explain previous results (Figures 10-12).

Kanai and colleagues (2010) offered one possible explanation of what may be driving these effects: differences in metacognitive judgements may be due to an “awareness availability.” That is, the awareness in AB or short lag trials might be in an all or none fashion (as suggested by Sergent & Dehaene, 2004), while the PS or long lag trials may be more gradual. We find some evidence to support these claims: as shown in Figure 6, ‘yes’ responses (left column) lead to a stronger confidence profile in the AB compared to phase scrambling (regardless of whether the yes response is correct) whereas ‘no’ responses (right column) lead to under-confidence in the AB compared to PS. As ‘signal’ trials are degraded in PS, it seems reasonable that subjects have more confidence in hits during the AB than during PS, and interestingly, this effect carries over to false alarms. Conversely, the performance drop for “no” responses during the AB is caused by attentional lapses, and this makes participants less confident that their response is correct when they miss a target or report absence of a trial, compared to PS. Thus, at some level subjects seem to ‘know’ when they miss a trial in the AB when compared to a PS trial, as a miss is caused at a higher level in the decision making hierarchy.

While our results provide strong evidence for our conclusion, there were some limitations in this study. One important limitation is the inability to staircase the AB trials to a set percentage. Staircasing would involve changing the stimulus duration in general or contrast of the AB (second) target. Both would be theoretically difficult.

Changing stimulus duration would also have to be paralleled in the PS trials, but then it would be difficult to staircase the long lag trials since the staircased variable (level PS) difficulty would depend on the (slightly) changing stimulus duration. Changing the contrast of the AB second target would also be theoretically difficult, since then the AB trials would include both perceptual and attentional blindness manipulations (Kanai et al., 2010) and the rationale was to include one perceptual (PS) and one attentional blindness manipulation (AB). Nevertheless, the analyses showed that the average percentages were around seventy percent correct across both conditions, which is near threshold.

Another important limitation was the lack of a fitting control trial type. We considered including trials in which the second target would be outside the AB timespan (as with the long lag) and without phase scrambling. However, such trials would be substantially less challenging or difficult and could have influenced how the confidence ratings would be distributed across trials of the same difficulty. The control trials would be more clearly visible, and thus also would have received higher confidence ratings while the experimental trials (AB and PS) would have gotten the lower confidence ratings. For this reason, we did not include separate control trials.

Nonetheless, our results remain intriguing. As we demonstrate in our simulations and experimental data, it is important to assess whether the equal variance assumption has been violated when analyzing detection-task data. A suggestion for future research is to combine this paradigm with EEG, to see if there is a difference in how well a decoder can categorize animal targets from no-animal targets for both trial types (short lag and long lag) and if this correlates with the metacognitive judgments. Finally, we also suggest that researchers compare conditions using manipulations other than attentional blink and phase scrambling, to see if these same trends would emerge.

Author Contributions

S.K. and J.F. developed the study design, with input from H.L. S.K. coded the experiment in PsychoPy2 (version 1.84.2) and collected all of the data. M.A.K.P. coded the SDT simulation file, which was modified by S.K. and B.O. S.K. created figures 1, 2, 3, and 6. B.O. coded the files for remaining figures and analyses, which were proofread and independently verified in analyses coded by S.K. All authors contributed to the writing of the manuscript.

Acknowledgements

The authors would like to thank Musen Li for his helpful feedback and proofreading of the code for the Chi-Square test.

References

- Bowers, J.S., & Jones, K.W. (2008). Detecting objects is easier than categorizing them. *Quarterly Journal of Experimental Psychology*, 61, 552-557.
- Fahrenfort, J.J., Van Leeuwen, J., Olivers, C.N.L., & Hogendoorn, H. (2017). Perceptual integration without conscious access. *Proceedings of the National Academy of Sciences of the United States of America*, 114(14), 3744-3749.
- Geisler, W.S. (2011). Contributions of ideal observer theory to vision research. *Vision Research*, 51(7), 771-781.
- Green, D. M., & Swets, J. A. (1966). Signal detection theory and psychophysics. 1966. New York.
- Grimaldi, P., Lau, H., & Basso, M.A. (2015). There are things that we know, and there are things that we do not know; Confidence in decision-making. *Neuroscience and Biobehavioral Reviews*, 55, 88-97.
- Jones, E., Oliphant, T., & Peterson, P. (2001--). {SciPy}: Open source scientific tools for {Python}. Retrieved from <http://www.citeulike.org/group/19049/article/13344001>
- Kanai, R., Walsh, V., Tseng, C.H. (2010). Subjective discriminability of invisibility: A framework for distinguishing perceptual and attentional failures of awareness. *Consciousness and Cognition*, 19, 1045-1057.
- Luu, L., & Stocker, A. A. (2018). Post-decision biases reveal a self-consistency principle in perceptual inference. *eLife*, 7. <https://doi.org/10.7554/eLife.33334>
- Mack, M.L., & Palmeri, T.J. (2010). Decoupling object detection and categorization. *Journal of Experimental Psychology: Human Perception and Performance*, 36, 1067-1079.
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide*, 2nd ed.

- Lawrence Erlbaum Associates Publishers. Retrieved from
<http://psycnet.apa.org/psycinfo/2004-19022-000>
- Maniscalco, B., & Lau, H. (2014). Signal Detection Theory Analysis of Type 1 and Type 2 Data: Meta-d', Response-Specific Meta-d', and the Unequal Variance SDT Model. In *The Cognitive Neuroscience of Metacognition* (pp. 25–66). Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-45190-4_3
- Maniscalco, B., Peters, M. A. K., & Lau, H. (2016). Heuristic use of perceptual evidence leads to dissociation between performance and metacognitive sensitivity. *Attention, Perception & Psychophysics*, 78(3), 923–937.
<https://doi.org/10.3758/s13414-016-1059-x>
- Meuwese, J.D.I., Van Loon, A.M., Lamme, V.A.F., & Fahrenfort, J.J. (2014). The subjective experience of objection recognition: Comparing metacognition for object detection and object categorization. *Attention, Perception, & Psychophysics*, 76, 1057-1068.
- Metcalfe, J. (1994). *Metacognition: Knowing about knowing*. Cambridge, MA: Bradford Books.
- Rossion, B., & Caharel, S. (2011). ERP evidence for the speed of face categorization in the human brain: Disentangling the contribution of low-level visual cues from face perception. *Vision Research*, 51, 1297-1311.
- Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, 8(10), 1391-1400.
- Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science*, 15, 721-728.

- Sinanaj, I., Cojan, Y., Vuilleumier, P. (2015). Inter-individual variability in metacognitive ability for visuomotor performance and underlying brain structures. *Consciousness and Cognition*, 36, 327-337.
- Watson, A.B., & Pelli, D.G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception and Psychophysics*, 33(2), 113-120.
- Wickens, T.D. (2002). *Elementary signal detection theory*. New York, NY: Oxford University Press.
- Wixted, J. T., (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, 114, 152-176.
- Zylberberg, A., Barttfeld, P., & Sigman, M. (2012). The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience*, 6, 79.
<https://doi.org/10.3389/fnint.2012.00079>