Journal of Experimental Psychology:
Human Perception and Performance
2016, Vol. 42, No. 5, 671–682

0096-1523/16/$12.00 http://dx.doi.org/10.1037/xhp0000179

Computations Underlying Confidence in Visual Perception

Morgan L. Spence, Paul E. Dux, and Derek H. Arnold
The University of Queensland

Humans intuitively evaluate their decisions by forming different levels of confidence. Despite being highly correlated, decisional confidence and sensitivity can be differentiated. The computational processes underlying this remain unknown. Here we find that, for visual judgments concerning global direction, signal range has a greater impact on confidence than it does sensitivity. We equated sensitivity for stimuli containing different degrees of directional variability. This failed, however, to equate confidence—participants were less confident when judging more variable signals despite constant sensitivity. When stimuli were instead calibrated to equate confidence, participants were more sensitive when judging more variable signals. Directional range had no impact on an unrelated judgment of brightness, helping to establish that these results cannot be attributed to a simple decisional confound. Our complementary results show that directional sensitivity and decisional confidence rely on independent transformations of sensory input. We propose that confidence will generally be shaped by the range of differently tuned neural mechanisms responsive to input during evidence accumulation, with this having a lesser impact on sensitivity.

Keywords: confidence, population codes, decision making, Bayesian inference, metacognition

Supplemental materials: http://dx.doi.org/10.1037/xhp0000179.supp

Humans intuitively evaluate their decisions and are often aware of mistakes before receiving explicit feedback concerning task performance. These feelings of confidence are relatively accurate, as they correlate well with objective sensitivity (Henmon, 1911; Peirce & Jastrow, 1885; Volkmann, 1934; Yeung & Summerfield, 2012). Estimates of confidence necessitate that the brain generate a reportable code concerning the precision of its own decisional processes (Fleming, Dolan, & Frith, 2012; Kepecs & Mainen, 2012). Despite considerable interest, the nature of the computations underlying decisional confidence remain unclear (Fleming et al., 2012; Yeung & Summerfield, 2012).

With regard to perceptual decisions, it has long been known that confidence can correlate with objective sensitivity measures (Henmon, 1911; Peirce & Jastrow, 1885; Volkmann, 1934). This would be explicable if confidence reflects the strength of accumulated sensory evidence at the time of decision making (Vickers, 1979). However, more recent evidence has shown that confidence and sensitivity are also separable (De Martino, Fleming, Garrett & Dolan, 2013; Fleming & Dolan, 2012; Fleming, Ryu, Golfinos, & Blackmon, 2014; Li, Hill, & He, 2014). This suggests that encoded signal strength is insufficient to account for computations underlying decision confidence (Yeung & Summerfield, 2012). An alternate possibility is that, in addition to the magnitude of sensory evidence, the brain also estimates the reliability of encoded information, and that this informs feelings of confidence.

We reasoned that a degree of independence between perceptual confidence and sensitivity would be explicable if perceptual confidence were disproportionately governed by the dispersion of activity across a population of neurons tuned to different values of a common stimulus attribute. Sensitivity, by contrast, could be determined by a weighted averaging of such responses (de Gardelle & Summerfield, 2011; Jazayeri & Movshon, 2006; Pouget, Dayan, & Zemel, 2000; Ma & Jazayeri, 2014; Yang & Shadlen, 2007). For example, in a global motion direction judgment the range of differently tuned direction selective cells could be adopted as a proxy for the reliability of the encoded signal, whereas the precision of perception could be governed more by the ability to extract an estimate of the average direction signaled by active neurons (see Figure 1).

Our reasoning was inspired by Bayesian accounts of sensory cue combination (Alais & Burr, 2004; Beck et al., 2008; Ernst & Banks, 2002; Ma, Beck, Latham, & Pouget, 2006), which propose that when the brain combines initially independent sensory signals it weighs these signals in proportion to their associated reliability. Population coding could provide the necessary information for Bayesian accounts of sensory cue combination, in that it offers both a sensory estimate (the maximum likelihood/central tendency of the fitted response function) and an estimate of reliability—the range of differently tuned neurons responsive to an input (the width of the fitted response function). Such estimates of reliability could inform confidence judgments. This would allow the brain to form relatively independent perceptual decisions and estimates of

---

Morgan L. Spence, Paul E. Dux, and Derek H. Arnold, School of Psychology, The University of Queensland.

Morgan L. Spence, Paul E. Dux, and Derek H. Arnold designed the experiments. Morgan L. Spence collected and analyzed the data, and prepared the figures. Morgan L. Spence, Paul E. Dux, and Derek H. Arnold wrote the manuscript. Australian Research Council Future Fellowships awarded to Derek H. Arnold (FT130100605) and Paul E. Dux (FT120100033) supported this research.

Correspondence concerning this article should be addressed to Morgan L. Spence, School of Psychology, The University of Queensland, St. Lucia, Queensland 4072, Australia. E-mail: morgan.spence@uqconnect.edu.au

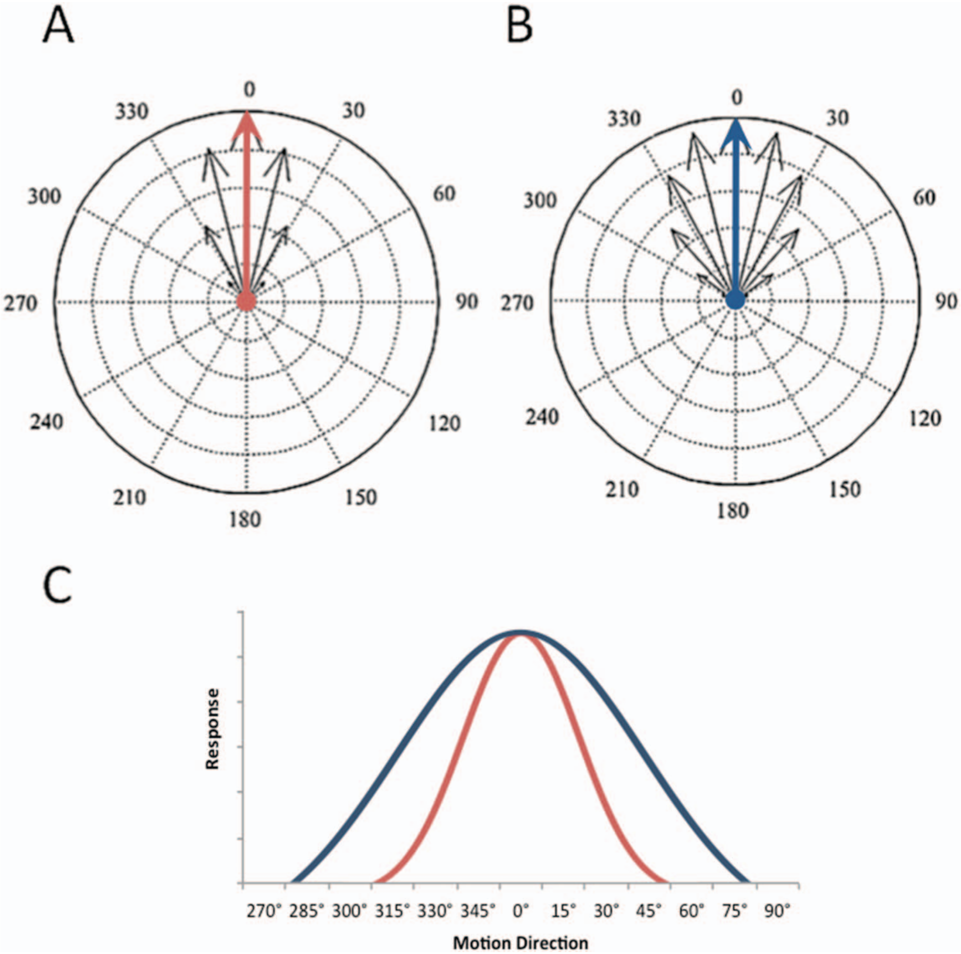

Figure 1. Depiction of a neurally plausible code for motion direction, and of how responsive direction-tuned mechanisms within it would be to stimuli containing different ranges of element directions, uniformly distributed about vertical. (A) Vectors showing the number of model filters responsive to a stimulus containing a uniform distribution of element directions ±5° from vertical. Vectors are scaled in proportion to the maximal response, and the red (light gray) vector coincides with the direction signaled by the vector sum (B). As for (A), but for a stimulus containing a uniform distribution of direction signals ±40° from vertical, indicated by the blue (dark gray) vector. In both cases, represented in (C), the dispersion of activity across the population of differently tuned direction filters is signified by the number of visible vectors, which could be adopted as a proxy for the reliability of the encoded signal. See the online article for the color version of this figure.

decisional confidence (Fleming et al., 2010; Kepecs & Mainen, 2012).

We can evaluate our proposals by examining performance and confidence in perceptual decisions regarding the global direction of a visual motion stimulus. If our hypotheses are supported, confidence should scale negatively with the range of direction signals in a stimulus, even if task difficulty is held constant. This can be achieved by manipulating the magnitude by which global test directions are offset from a decision boundary (here vertical). Conversely, if task difficulty is instead manipulated to maintain a constant level of confidence, performance should improve for stimuli containing an increasing range of different direction signals, as the task will need to be made easier to compensate for deteriorating confidence as a broader range of direction signals is encountered. Here, we establish the validity of both proposals across four core psychophysical experiments, and we rule out a potential confound-based explanation of these data in an additional control experiment.

Experiment 1

Method

Eight volunteers participated in Experiment 1. All were experienced psychophysical observers. Six were naïve as to the purpose of the study (authors M. L. S. and D. H. A. also participated). The number of participants in all experiments reported here was determined a priori as a convenience sample. All participants had normal or corrected-to-normal visual acuity and were seated in a darkened room viewing stimuli from a distance of 70 cm, with their head restrained via a chinrest. The University of Queensland ethical committee approved the study protocol. The observers' task was to discern the direction of global motion in a dot kinematogram as being to the left or right relative to the vertical axis. The first phase of the experiment was a calibration procedure, used to equate performance in the subsequent method of constant stimuli (MOCS) task.

Stimuli. The dot kinematograms were generated using a Cambridge Research Systems ViSaGe stimulus generator driven by Matlab R2013b (MathWorks, Natick, MA) software and was presented on a gamma-corrected 19 in. Dell P1130 monitor (resolution: 1,600 × 1,200 pixels; refresh rate: 85 Hz). Kinematograms were presented for 500 ms within a circular aperture with a diameter subtending 1.79 degrees of visual angle (dva) at the retina. Dot density was 100 dots/3.94 dva. Individual limited lifetime dots (100 ms) were white (luminance = 102 cd/m^2), subtended 0.016 dva in diameter and moved in a linear direction at a speed of 0.72 dva/s. Stimuli were centered on fixation. The vertical axis was signaled via three static red disks, each subtending 0.07 dva in diameter, positioned at fixation and 2.77 dva directly above and below fixation. The global direction of test motion, upward or downward, alternated on successive trials to mitigate the build-up of motion aftereffect signals. Individual dots moved in one of 10 directions, uniformly distributed about a range from the mean test direction (see Figure 2A–D). Four test ranges were sampled, ±5, 10, 20, and 30° from the mean test direction. We adopted this range manipulation as the range of element direction signals within a stimulus determines the bandwidth of neural responses across a population of direction-tuned cells (de Gardelle & Summerfield, 2011; Jazayeri & Movshon, 2006; Pouget et al., 2000; Ma & Jazayeri, 2014; Yang & Shadlen, 2007; Britten, Shadlen, Newsome, & Movshon, 1992).

Procedure. In the calibration phase, we determined magnitudes of global direction offset, left or right from vertical, resulting in ~70% correct direction discrimination performance for each range using one-up two-down staircase procedures (Levitt, 1971). This enabled us to equate performance for stimuli containing different ranges of directional signals by adjusting the magnitudes

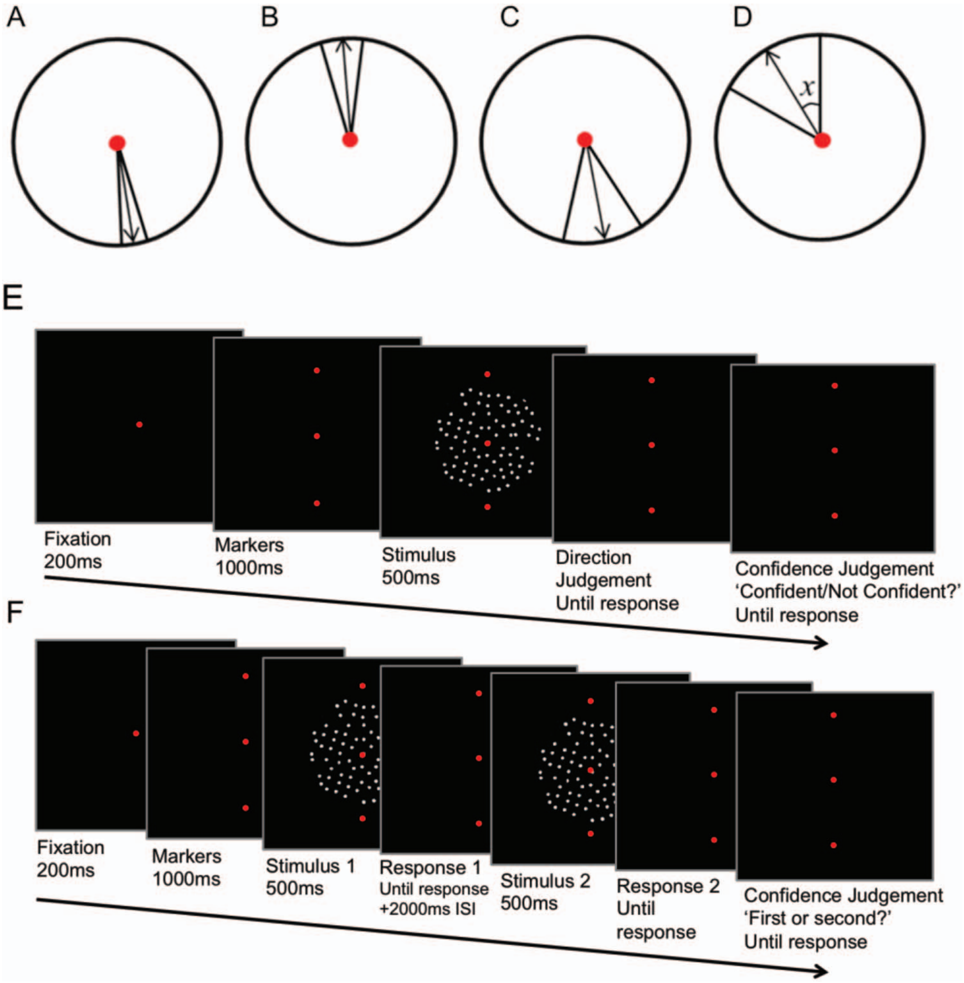

Figure 2. Graphic depiction of stimuli (A–D) and experimental protocols for Experiment 1 (E), Experiments 2 and 3 (F). (A–D) Here angles depict the range of directions, about a mean test direction, that individual dots moved in. There were four conditions: (A) ±5°, (B) 10°, (C) 20°, and (D) 30°. See the online article for the color version of this figure.

by which global test directions were offset from the decisional boundary (vertical; see Figure 2A–D). Two staircases were conducted for each test range. One was instigated at an angular offset of ±12° (starting high) and another at an angular offset of ±1° (starting low). These values were the maximal and minimal possible test values.

On each staircase trial, participants first reported the direction of global motion (left or right from the vertical). The first three trials for each test range were used as practice trials, presented at the maximum angular offset. Thereafter, staircase procedures were implemented with angular offsets manipulated in 1° steps, up after incorrect responses, and down after two successive correct responses within each staircase. Each staircase was sampled for 30 individual trials, so 60 individual trials in total were conducted for each test range. A block of trials for this preliminary task involved 240 individual trials, with presentations of the different test ranges randomly interleaved. Data were collated across the two staircase procedures conducted for each test range condition, and a logistic function fitted to proportion correct data as a function of angular offset magnitude, in order to estimate angular offsets resulting in ~70% correct task performance.

In the second primary MOCS phase of the experiment, angular offsets were set to individual thresholds obtained using the previous calibration procedure for each test range condition. As in the calibration phase, participants first indicated the perceived direction of motion relative to vertical (left/right), then whether they had felt a high or low level of confidence in their perceptual decision (see Maniscalco & Lau, 2012). Both responses were made using right and left mouse buttons. During a block of trials each of the four test range conditions was sampled 100 times in random order, yielding a total of 400 individual trials. Each participant completed two blocks of trials, providing proportion correct and proportion confident scores across 200 individual trials for each test range.

Results

Directional offset magnitudes determined via our preliminary calibration task were 5.7 ± 3.3, 5.7 ± 1.7, 5.0 ± 1.1, and 7.2 ± 2.0 angular degrees, respectively, for stimuli with direction ranges of ±5°, 10°, 20°, and 30° from the test mean. One-way repeated measures ANOVAs revealed that, as intended, there was no difference in task performance across the four direction range conditions, $F(1, 7) = 0.610$, $p = .555$, $\eta_p^2 = 0.08$ (see Figure 3A). There was, however, a discernible impact of direction signal range on confidence, $F(1, 7) = 10.106$, $p = .001$, $\eta_p^2 = 0.59$ (see online

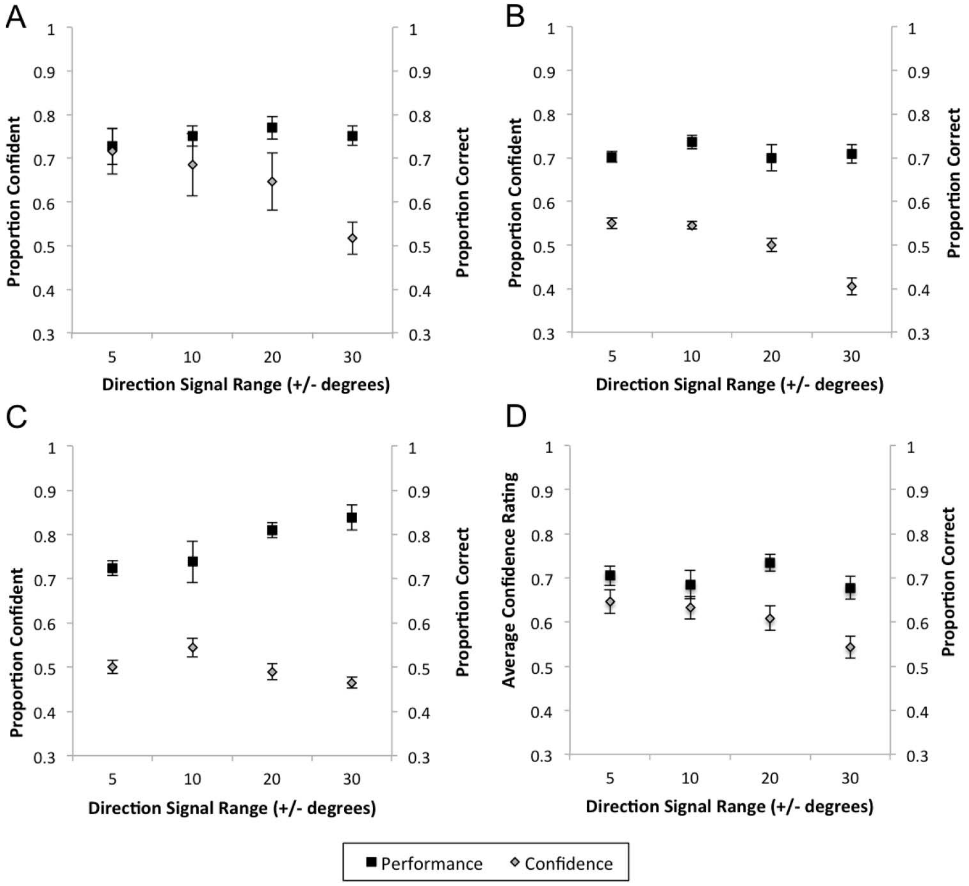

Figure 3. Graphs depicting results of Experiments 1–4. (A) Results of Experiment 1. As the range of test direction signals increased, task performance was invariant, but reported confidence declined. (B) Results of Experiment 2. (C) Results of Experiment 3. (D) Results of Experiment 4. In all cases, error bars depict ±1 *SEM*.

supplementary material Tables S1–2 individual participant data). Confidence declined as the range of stimulus element directions broadened: linear trend analysis, $F(1, 7) = 36.33$, $p = .001$, $\eta_p^2 = 0.84$.

In addition to the one-way repeated-measures ANOVAs, we performed one-way repeated-measures Bayes analysis of variance (BANOVA; R statistical computing software, R Core Team, 2014). This allowed us to examine the likelihood of the observed performance and confidence results against both the null and alternate hypotheses. For proportion correct data, this yielded a raw Bayes factor BF10 (the likelihood of the data occurring under assumptions of the alternative hypothesis over the null) of 0.262 (±0.46%), providing evidence in favor of the null hypothesis, that there is no difference in performance across variability conditions (Kass & Raftery, 1995). For proportion confident data this analysis yielded a Bayes factor BF10 of 92.47 (±0.88%), providing strong support for the alternative hypothesis that confidence varies across the four direction range conditions.

To determine whether declining confidence with increasing signal range was associated with a decline in metacognitive sensitivity we assessed the degree to which confidence ratings could discriminate between correct and incorrect perceptual decisions as a function of test range (Maniscalco & Lau, 2012; Barrett, Dienes, & Seth, 2013; for a review, see Fleming & Lau, 2014). We did so by plotting proportion correct data for trials associated with high and low confidence ratings, generating a conditional accuracy plot (see Figure 4). If the observed decline in confidence as a function of direction signal range was due solely to a reduction in metacognitive sensitivity, we would expect to see a marked interaction between high and low confidence proportion correct, with performance on these two types of trial becoming increasingly similar for greater signal ranges. An alternate possibility is that there is no such interaction, which would suggest that declining reports of confidence were due to participants adopting increasingly conservative criteria for reporting high-levels of confidence as they encounter a greater range of directional signals.

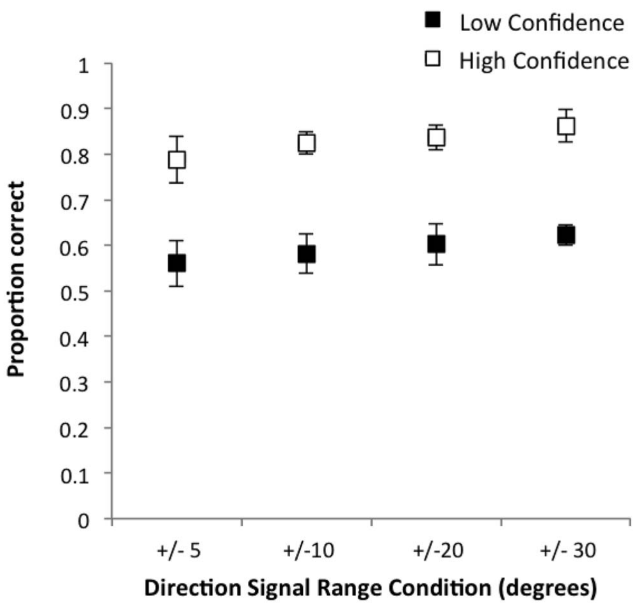

Figure 4. Graph depicting conditional accuracy plot for Experiment 1. Error bars depict ±1 *SEM*.

Results of a two-way repeated measures ANOVA revealed a main effect of confidence level, such that task performance was greater on high- than on low-confidence trials, $F(1, 7) = 118.54$, $p < .001$, $\eta_p^2 = 0.94$. This indicates that participants had good insight into their objective performance. However, there was no main effect of direction signal range, $F(1, 7) = 1.47$, $p = .260$, $\eta_p^2 = 0.17$, and no interaction between confidence level and signal range, $F(1, 7) = .02$, $p = .969$, $\eta_p^2 = 0.003$. Type II sensitivity (the precision with which confidence could predict performance) was therefore similar across range conditions, which nonetheless were marked by declining levels of overall confidence as signal range increased. This suggests our data reflect a shift in the criteria used to demark high from low levels of confidence as increasingly broad ranges of direction signals were encountered.

Discussion

The results of Experiment 1 show that confidence in global direction judgments declines for stimuli containing an increasingly broad range of directional signals, and that this influence is disproportionate relative to direction discrimination performance, which was held constant by experimental design. This suggests a disproportionate weighting of the range of direction signals in computations underlying confidence, relative to sensitivity. In Experiment 2 we attempted to replicate this observation using a potentially more sensitive measure of confidence in order to assess the generality of our results.

Experiment 2

Method

All details for Experiment 2 were as for Experiment 1 ($N = 8$), with the following exceptions. On each trial of the test phase participants were presented two successive motion stimuli. After each, participants made a global motion direction discrimination (left/right of vertical) and then, following the second global direction discrimination, the participant also indicated which judgment (first/second) they had felt most confidence in (Barthelme & Mamassian, 2009, 2010; de Gardelle & Mamassian, 2014).

Global direction offset (left/right of vertical) was counterbalanced across a block of trials, and presentation order was randomized. Pairs of successive stimuli consisted of all combinations of test range (± 5 and 10, 5 and 20, 5 and 30, 10 and 20, 10 and 30, 20 and 30). During a block of trials each combination of test ranges was presented 30 times, for a total of 180 individual trials. Each participant completed two blocks of trials (360 trials in total).

Results

Data from both trial blocks, for each experimental condition, were collated and proportion correct global direction judgments and proportion confident scores were calculated for each participant. Proportion confident scores reflect the proportion of trials in which a stimulus condition was presented, as one of the two conditions presented on that trial, and was then also chosen as the stimulus that had elicited greater decisional confidence. One-way repeated measures ANOVAs revealed that direction discrimination performance was again invariant for tests containing different ranges of element direction signals, $F(1, 7) = 0.711$, $p = .484$, $\eta_p^2 = 0.09$ (see Figure 3B). Means and *SEM* for proportion correct were 0.70 (0.01), 0.74 (0.02), 0.70 (0.03), and 0.71 (0.02) for ± 5°, 10°, 20°, and 30°, respectively, which equated to *d′* values of 1.03 (0.06), 1.31 (0.08), 1.13 (0.18), and 1.13 (0.10). Although performance was constant, confidence varied across conditions, $F(1, 7) = 16.01$, $p < .001$, $\eta_p^2 = 0.70$ (see online supplementary material Tables S3–4 for individual participant data). As in Experiment 1, levels of confidence declined for tests containing progressively broad ranges of direction signals: linear trend analysis, $F(1, 7) = 32.74$, $p = .001$, $\eta_p^2 = 0.82$.

In addition to one-way repeated-measures ANOVAs, we subjected data to BANOVAs. For proportion correct data this yielded a Bayes factor BF10 of 0.305 (±0.51%) providing evidence supporting the null hypothesis that there would be no difference in task performance across our test conditions. For proportion confident data this yielded a Bayes factor BF10 of 57,960.82 (±0.84%), providing support for the alternate hypothesis, that confidence would vary across conditions.

As in Experiment 1, we examined whether declining confidence with increasing signal range was associated with a decline in the degree to which participants correctly distinguished correct from incorrect decisions with their confidence responses. The motion direction judgment that was chosen within each pair of judgments as having been associated with more relative confidence was classified as "high confidence." Conversely, unchosen judgments were classified as "low confidence." A conditional accuracy plot generated with proportion correct for high (chosen) and low (unchosen) confidence judgments is shown in Figure 5.

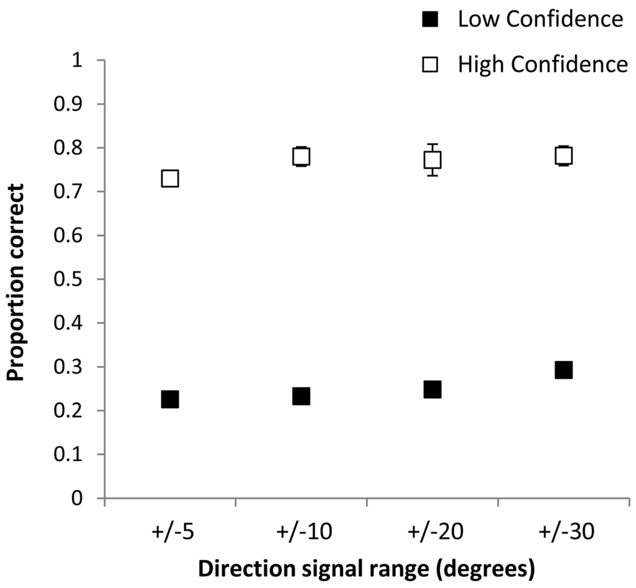

Figure 5. Graph depicting conditional accuracy plot for Experiment 2. Error bars depict ±1 *SEM*. It is important to note that, unlike Experiment 1, participants can be correct or incorrect on both trials in the pair despite choosing only one as having been associated with high confidence. Results are shown for each range condition as it was paired with all other conditions.

Results of a two-way repeated measures ANOVA revealed a main effect of confidence level, such performance was improved for high- relative to low-confidence choices, $F(1, 7) = 1663.35$, $p < .001$, $\eta_p^2 = 0.99$. There was, however, no main effect of direction signal range, $F(1, 7) = 2.41$, $p = .148$, $\eta_p^2 = 0.29$, and no Confidence × Signal Range Condition interaction, $F(1, 7) = .82$, $p = .498$, $\eta_p^2 = 0.12$, indicating that there was no change in metacognitive sensitivity across range conditions.

Discussion

The results of Experiment 2 replicate findings of Experiment 1, showing that confidence in global direction judgments declines for stimuli containing an increasingly broad range of direction signals, and that this influence is disproportionate relative to objective decisional accuracy. In Experiment 3 we examined if we could demonstrate a complementary result. If signal range impacts confidence disproportionately relative to sensitivity, when stimuli are calibrated to equate levels of confidence direction sensitivity should *improve* for stimuli containing an increasingly broad range of direction signals. Put differently, direction judgments will have to be made progressively easier in order to compensate for increasing levels of uncertainty triggered by stimuli containing increasingly broad ranges of direction signals.

Experiment 3

Details for Experiment 3 were as for Experiment 2 ($N = 8$), with the following exceptions: Trials during the preliminary calibration task were like those in the test phase of Experiment 2, with successive presentations of two test stimuli. One of these, the standard, had test directions uniformly distributed ±5° from the mean test direction, which was set to the participant's 70% threshold for this stimulus from Experiment 2 (mean offset 5.1°, $SD = 2.5°$). The global direction offset from vertical for comparison stimuli (± 10°, 20°, or 30° range conditions) was adjusted on a trial-by-trial basis according to one-up one-down staircase procedures, up by 1° if the participant chose the standard as the stimulus in which they had felt the most confident when judging direction, and down by 1° otherwise. Two staircase procedures were conducted for each comparison stimulus, one was instigated at the maximal offset (22°), the other at the minimal offset (1°).

During a block of trials each of the two staircase procedures for each of the three standard comparison combinations was sampled on 30 individual trials, with blocks of trials containing 180 individual trials all sampled in random order. Data for each comparison were collated across two staircase procedures for that test range, and logistic functions were fitted to proportion confident scores as a function of comparison global direction offset from vertical. We took 50% points on fitted functions as estimates of the global direction offset from vertical resulting in equal levels of confidence relative to the standard. Global direction offset magnitudes determined via this procedure were 6.48 (±3.79), 7.54 (±2.37), and 10.51 (±3.56) for ± 10°, 20°, and 30°, respectively. Details for subsequent blocks of test trials were as for Experiment 2, but in this case stimuli had been calibrated to elicit equal levels of confidence, as opposed to equal levels of performance.

Results

As intended, confidence levels were invariant for stimuli containing different ranges of element direction signals, $F(1, 7) = 2.77$, $p = .088$, $\eta_p^2 = 0.28$, but task performance was variable, $F(1, 7) = 3.85$, $p = .048$, $\eta_p^2 = 0.36$ (see Figure 3C), improving for tests containing increasingly broad ranges of element direction signals: linear trend analysis, $F(1, 7) = 13.05$, $p = .009$, $\eta_p^2 = 0.65$ (see online supplementary material Tables S5–6 for individual participant data). Means and *SEM* for proportion correct were 0.72 (0.02), 0.74 (0.05), 0.81 (0.02), and 0.84 (0.03) for ±5°, 10°, 20°, and 30°, respectively, which equated to d' values of 1.14 (0.08), 1.14 (0.29), 1.76 (0.25), and 1.85 (0.21).

In addition to one-way repeated-measures ANOVAs, proportion correct and proportion confident scores (see online supplementary material Tables S5–6 for means and standard deviations) were subjected to BANOVAs. For proportion confident data this revealed a Bayes factor BF10 of 2.95 (±0.40%), providing support for neither the null nor the alternate hypothesis, that confidence would vary as a function the range of stimulus direction signals. Note that this lack of evidential support contrasts with Experiments 1 and 2. For proportion correct data, this yielded a Bayes factor BF10 of 4.82 (±0.39%), providing support for the alternate hypothesis that performance would vary across range conditions.

Discussion

The results of Experiment 3 complement the findings of Experiments 1 and 2, in that they show that the range of directional signals has a disproportionate impact on decisional confidence relative to objective task performance. Here this is evident as stimuli had to be calibrated such that direction judgments became increasingly easy as the range of directional signals increased, in order to compensate for the adverse impact this had on confidence.

One potential problem with Experiments 1–3 is that they all rely on people making binary decisions, either about whether they had felt a low or high level of confidence (Experiment 1), or about which of two decisions they had felt the most confidence (Experiments 2–3). This could be problematic if there had, in fact, been no systematic change in decisional confidence, and participants instead responded in a “Clever Hans” fashion, adopting the differential appearance of stimuli as a selection criterion when forced to make binary classifications regarding a quality (confidence) that was in fact invariant. This would have constituted a task-demand-related confound unrelated to the desired dependent measure—confidence. To address this concern, in Experiment 4 we adopted a continuous measure of decisional confidence, rather than having participants make binary categorical judgments.

Experiment 4

Details for Experiment 4 were as for Experiment 1 ($N = 8$), with the following exceptions. Following motion direction judgments, on each trial participants indicated how confident they felt in the preceding judgment by setting the position of a marker to any point along a continuum marked by two extremities (see Figure 6). The left extremity was marked by a red bar (CIE x = 0.63 y = 0.35 Y = 17, height: 1.1 dva, width: 0.2 dva, positioned 2 dva to the left of the display center) signifying a complete lack of confidence (guessing). The right extremity was marked by a green bar (CIE x = 0.28 y = 0.62 Y = 66, height: 1.1 dva, width: 0.2 dva, positioned 2 dva to the right of the display center) signifying complete confidence. The marker, a white bar (CIE x = 0.28 y = 0.33 Y = 91, height: 0.5 dva, width: 0.2 dva), was initially positioned in the display center, and could be shifted left or right by holding down respective mouse buttons. Once the participant was satisfied that the position of the slider indicated their felt level of confidence, they pressed the middle mouse button, at which point the marker position was recorded on a scale from 0 (guessing) to 1 (complete confidence) with a resolution of 0.01. These values were not made known to the participant. At the end of a run of trials we calculated the average confidence score associated with each experimental condition.

Results

Performance levels were invariant for stimuli calibrated to promote equal levels of performance, even though stimuli contained different ranges of direction signals, $F(1, 7) = 0.84$, $p = .456$,

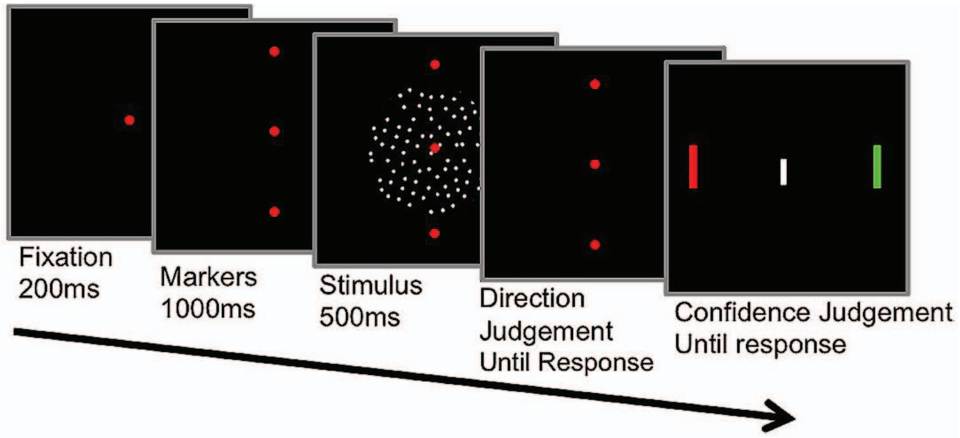

Figure 6. Graphic depicting the experimental protocol in Experiment 4. See the online article for the color version of this figure.

$\eta_p^2 = 0.11$. Confidence scores were, however, variable, $F(1, 7) = 15.88$, $p = .001$, $\eta_p^2 = 0.69$ (see Figure 3D), with average levels of confidence decreasing for tests containing increasingly broad ranges element direction signals (±5°, $M = 0.65$, $SEM = 0.02$; 10°, $M = 0.63$, $SEM = 0.03$; 20°, $M = 0.61$, $SEM = 0.02$; and 30°, $M = 0.54$, $SEM = 0.03$), linear trend analysis, $F(1, 7) = 19.35$, $p = .003$, $\eta_p^2 = .73$ (see online supplementary material Tables S7–8 for individual participant data).

In addition to the one-way repeated-measures ANOVAs, proportion correct and average confident scores were subjected to BANOVAs. For proportion correct data this yielded a Bayes factor BF10 of 0.37 (±0.44%), providing support for the null hypothesis that there would be no differences in performance across conditions. For average confidence scores this yielded a Bayes factor BF10 of 1,094.80 (±0.69%), providing strong support for the alternate hypothesis that confidence would vary across range conditions.

As in Experiment 1, we examined whether declining confidence with increasing signal range was associated with a decline in the degree to which confidence correctly distinguished between correct and incorrect perceptual decisions—metacognitive sensitivity. Trial-by-trial confidence ratings were first categorized as high (confidence ratings ≥ 0.5) or low (confidence ratings <0.5). Proportion correct performance on high- and low-confidence trials were then plotted (see Figure 7) and analyzed.

Results of a two-way repeated measures ANOVA revealed a main effect of confidence level, such that participants had greater levels of performance on high- relative to low-confidence trials, $F(1, 7) = 36.79$, $p = .001$, $\eta_p^2 = 0.84$. There was, however, no main effect of direction signal range, $F(1, 7) = .66$, $p = .523$, $\eta_p^2 = 0.09$, and no Confidence Level × Signal Range Condition interaction, $F(1, 7) = .82$, $p = .471$, $\eta_p^2 = 0.11$. These results are therefore similar to Experiment 1, suggesting the effect of direction signal range on confidence is to bring about a shift in the criteria used to adjudge confidence.

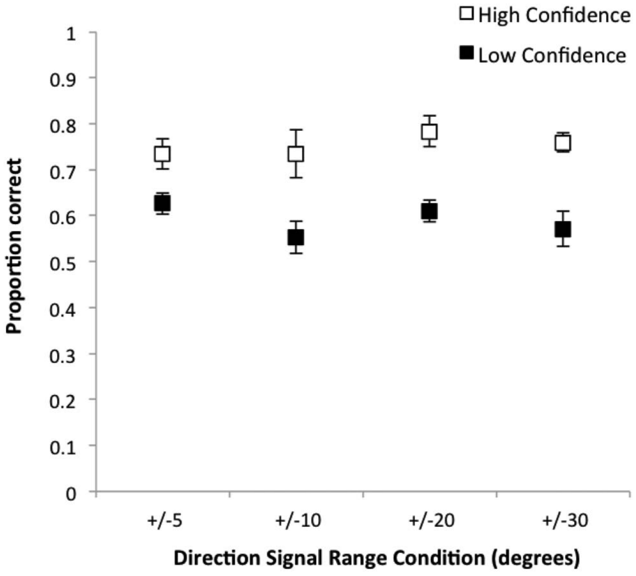

Figure 7. Graph depicting the conditional accuracy plot for Experiment 4. In both cases, error bars depict ±1 *SEM*.

Discussion

The results of Experiment 4 are consistent with the findings of Experiments 1–3, in that they show that the range of direction signals had a disproportionate impact on decisional confidence relative to sensitivity. These results cannot readily be ascribed to a simple "Clever Hans" scenario unrelated to our desired dependent measure—the level of confidence felt in global direction judgments. This is the case as there was no requirement to classify confidence as high or low, or to choose one of two stimuli as having elicited more confidence. If participants had felt a constant level of confidence across conditions, they were free to set the confidence marker to a constant position across experimental conditions. They did not do so, demonstrating that participants had felt different levels of decision confidence for our different stimuli, despite these having elicited a constant level of objective task performance. Hence we believe both these data, and the results of our previous experiments, speak to the influence that direction signal range has on confidence in global direction judgments.

If our data speak to the influence that direction signal range has on global direction judgments, we should find this has no impact on an unrelated perceptual judgment, such as stimulus brightness. If, however, our data result from participants coopting stimulus appearance as a proxy for confidence (which is actually invariable due to sensitivity having been equated across experimental conditions), participants should display a similar response bias when judging properties unrelated to global direction. We assessed this possibility in Experiment 5.

Experiment 5

Details for Experiment 5 were as for Experiment 3 ($N = 8$) with the following exceptions. Participants were presented two successive dot kinematograms presentations (the standard and a comparator). These moved in an average vertical motion direction, up or down on successive trials. In the first phase of the experiment ~70% correct thresholds for detecting which of these two stimuli had been brighter were determined for each range condition, using one-up two-down staircase procedures. The standard (white) was set at CIE x = 0.29 y = 0.34 Y = 90.9. Comparators had the same chromaticity coordinates but a lesser level of luminance intensity, which was instigated at a level of 45.4 cd/m^2 and adjusted in steps of 3.8cd/m^2 according to staircase procedures with participants selecting which of the two sequential tests had seemed brightest, before making a binary confidence judgment (see Figure 2F for a depiction of the paradigm). There were 50 calibration trials per range condition.

For the second critical phase of the experiment, comparators were set to the individuals' brightness threshold for that stimulus type (with different ranges of direction signal, as in Experiment 3). Each combination of stimulus condition was sampled on 60 individual trials, with a block of trials containing 240 individual trials, all sampled in random order. Participants completed two blocks of trials.

Results

By design, brightness discrimination performance was invariant for stimuli containing different ranges of direction signals, $F(1,$

7) = 0.42, p = .737, $\eta_p^2 = 0.06$. This mimicked the results of Experiments 1, 2, and 4. Confidence ratings, however, were also invariant across directional signal range conditions, $F(1, 7) = 0.229$, p = .875, $\eta_p^2 = 0.03$ (see Figure 8). Individual participant data is reported in online supplementary materials Tables S9–10.

In addition to the one-way repeated-measures ANOVAs, proportion correct and average confident scores were subjected to BANOVAs. For proportion correct data this yielded a BF10 of 0.24 (±1.38%), providing support for the null hypothesis that there was no differences in performance across conditions. For average confidence scores, this analysis revealed a BF10 of 0.19 (±0.77%), providing support for the null hypothesis that confidence would not vary across directional range conditions.

The results of Experiment 5 suggest the results of all previous experiments speak to a systematic influence of direction signal range on confidence for global direction judgments. In Experiment 5, as in Experiments 1, 2, and 4, objective task performance was held constant for stimuli containing different ranges of directional signals. If people had adopted stimulus appearance as a proxy for confidence, which was invariant due to a constant level of task performance, they should have shown the same response bias in this as in previous experiments. They did not. This signifies that direction signal range can disproportionately undermine confidence for global direction judgments, relative to objective task performance, but this has no impact on confidence when people judge an unrelated property—brightness.

General Discussion

We have shown that decisional confidence in global direction judgments is inversely related to the range of direction signals within the stimulus. Importantly, this relationship was disproportionate relative to the impact signal range had on objective task performance. Despite constant levels of objective task performance in Experiments 1, 2, and 4, increasingly broad ranges of direction signals negatively impacted confidence in global direction judgments. Note, however, that these data speak to a disproportionate impact, not to a complete dissociation. Stimulus calibrations in these experiments successfully manipulated the magnitude by which mean global test directions were offset from vertical (the decisional boundary) in order to equate objective task performance across conditions. This required increasingly large offsets for stimuli containing increasingly broad ranges of direction signals (Experiments 1, 2, and 4). So the range of directional signals impacted on both objective sensitivity and confidence, but the impact on confidence was disproportionate.

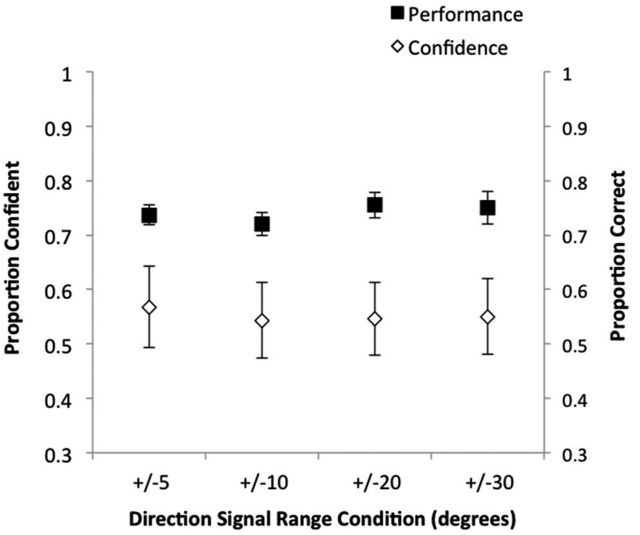

Figure 8. Graph depicting the results of Experiment 5. Error bars depict ±1 *SEM*.

The relationship between directional signal range and confidence in global direction judgments was also evident when offset magnitudes were calibrated (further increased as directional signal range increased) to equate felt levels of confidence across experimental conditions (Experiment 3). This resulted in increased levels of performance, despite constant levels of confidence. These complementary results reveal a differential weighting of signal range in computations that limit visual sensitivity and decision confidence for global direction judgments. More broadly, in combination, Experiments 1–4 show that the precision of perceptual decisions, and confidence in those decisions, can rely on relatively independent transformations of sensory input.

The disproportionate impact of signal range on decisional confidence, relative to the sensitivity of direction judgments, cannot reasonably be ascribed to a people adopting stimulus appearance as a proxy for confidence when asked to either categorize confidence as low or high (Experiment 1) or when deciding in which of two decisions they had felt a greater level of confidence (Experiments 2, 3, and 5). In each case, if confidence had been constant, there was a risk that participants might not have responded randomly, but instead made confidence decisions based on stimulus appearance, resulting in a confound that was unrelated to felt confidence. Our data, however, suggest a reliable and systematic impact of signal range on decisional confidence across participants. Observers in all experiments reported less confidence for broader ranges of signals, whereas a confound like that described above could reasonably be expected to be variable across participants, with some arbitrarily rating more variable signals as eliciting greater confidence, and others vice versa.

Further evidence against an alternate “Clever Hans” explanation of our data, independent of the computations governing confidence, was provided by the results of Experiments 4 and 5. In Experiment 4 we implemented a continuous measure of confidence, which should be immune to biases purely resulting from demand characteristics associated with binary forced categorization tasks that were not associated with actual confidence. If there was indeed no real variance in felt confidence across experimental conditions, participants were free to set a marker to a constant position to reflect this. The data show that they in fact did not; the results of Experiment 4 were consistent with all previous experiments. In Experiment 5 we investigated the effect of directional signal range on reported confidence when judging brightness. As expected, there was no effect of signal range on confidence when judging a quality unrelated to encoded global direction (brightness), so we can ascribe the results of all previous experiments to the systematic impact the range of directional signals had on confidence when reporting encoded global direction, rather than to a task-demand-related confound.

Some recent observations are pertinent to our data. Zylberberg, Roelfsema, and Sigman (2014) examined global orientation judgments concerning stimuli containing multiple oriented line segments. In contrast to our findings, they reported that stimuli containing a broader range of differently oriented elements could result in illusions of increased confidence, despite poor objective performance. This finding could speak to a discrepancy between the encoding of confidence for a spatial (orientation) relative to a temporal (global direction) stimulus dimension. Alternatively, this result could ensue if, instead of integrating multiple oriented signals, participants had based confidence judgments on a particularly salient individual element. We plan to address these possibilities in future experiments.

de Gardelle and Mamassian (2015) have reported a lack of a single systematic relationship between stimulus variability and confidence in global direction judgments. Specifically, some participants exhibited greater confidence for more variable global direction signals, whereas others displayed lesser levels of confidence for more variable signals. Unlike our experiments, this study did not precalibrate stimuli to equate performance across experimental conditions. Instead, they compared conditions containing different ranges of global directions by selecting, for each condition, one of a number of different sampled global directions on the basis that these were approximately matched in terms of performance. While it is difficult to interpret a null result, the failure to detect a reliable relationship between signal range and global direction judgment confidence may have been due to a lack of power, with insufficient trials concentrated on critical stimulus conditions precisely matched in terms of sensitivity. This was a key feature of our experimental design, and we consistently found evidence for a systematic relationship between signal range and confidence in global direction judgments.

Our findings should also be considered in relation to recent studies using TMS to introduce neural noise. Rahnev, Maniscalco, Luber, Lau, and Lisanby (2012) found that direct injection of noise using subthreshold, low intensity single-pulse occipital TMS, slightly decreased perceptual accuracy in an orientation discrimination, but increased confidence. Others have observed decreased confidence with higher intensity TMS to area MT, an approach that is perhaps more analogous to the present study (e.g., Koivisto, Mäntylä, & Silvanto, 2010; Koivisto, Railo, & Salminen-Vaparanta, 2011). Recent evidence has also shown that confidence can be disrupted without impairing visual discrimination performance by stimulating the motor response for the alternative choice in premotor cortex (Fleming et al., 2015). Adjustments to a confidence criterion might also occur under different instructions, resulting in interactive effects of noise injection and behavior (Rahnev et al., 2012). Taken together, these results highlight that diverse effects of external noise injection on confidence can often be obtained using different experimental paradigms and stimuli (e.g., Harris, Clifford, & Miniussi, 2008 & Ruzzoli et al., 2011). How these findings relate to our own will need to be clarified systematically by further research.

Also pertinent are recent studies that have examined the role of decision time in determining decisional confidence. Both Kiani, Corthell, and Shadlen (2014) and Zylberberg, Barttfeld, and Sigman (2012) have suggested that decisional confidence can be negatively related to response times, such that more rapid decisions are associated with greater levels of confidence. Although response times were not measured in this study, interesting avenues for future research may include an examination of response time with respect to the dissociation of performance and confidence due to signal range reported here.

We have conducted a sequence of experiments using carefully calibrated stimuli, and found consistent results across all experiments. We regard our data as evidence that the precision of perceptual decisions and the determination of perceptual confidence can rely disproportionately on different aspects of neural population coding (Kiani & Shadlen, 2009). The accuracy of perceptual decisions is more influenced by the mean value to which active neurons respond leading up to a decision, whereas confidence is more governed by the range of differently tuned neurons active during the evidence accumulation. This could be adopted as a proxy for the reliability of the encoded signal, and thereby inform confidence ratings (de Gardelle & Summerfield, 2011; Jazayeri & Movshon, 2006; Pouget et al., 2000; Ma & Jazayeri, 2014; Yang & Shadlen, 2007; Alais & Burr, 2004; Beck et al., 2008; Ernst & Banks, 2002; Ma et al., 2006; Solomon, Cavanagh, & Gorea, 2012). This relationship would allow for the brain to generate relatively independent estimates of a visual attribute and of the degree to which one should have confidence in that decision (Yeung & Summerfield, 2012; Kiani & Shadlen, 2009; Lak et al., 2014). The computational process proposed here is analogous to Bayesian accounts of sensory cue combination. It has been shown that when sensory decisions are informed by signals from multiple sensory modalities, the decision is governed by a weighted summation of sensory estimates, with weightings determined by the range associated with each of the initially independent signals (Ma & Jazayeri, 2014; Alais & Burr, 2004; Beck et al., 2008; Ernst & Banks, 2002; Ma et al., 2006; Lak et al., 2014). Data presented here show that the reliability estimates necessary for this process could provide the basis for confidence when judging global direction.

An important implication of our findings is that models that quantify decision confidence as a "readout" of objective sensitivity are insufficient to account for the transformations of sensory information implicated in computations governing confidence. There has, for instance, been a recent uptake of signal-detection theoretic (SDT) approaches to confidence measurement that constrain "metacognitive (Type II) sensitivity" as the degree to which reported confidence ratings correctly discriminate between correct and incorrect responses ("meta *d* prime," Maniscalco & Lau, 2012; Barrett et al., 2013; for a review, see Fleming & Lau, 2014). This framework makes the assumption that the information limiting decisional sensitivity (d') is exhaustive of the information available for the confidence judgment for a metacognitively efficient observer (Maniscalco & Lau, 2012). Contrary to this assumption, recent findings from this model show that observers' metacognitive sensitivity (Type II sensitivity) is less than what is predicted by their objective sensitivity (Type I sensitivity; e.g., Maniscalco & Lau, 2012). To explain this discrepancy, Maniscalco and Lau (2012) suggested that different representations or transformations of the same underlying information might pertain to confidence relative to sensitivity. Consistent with this reasoning, we show that the inadequacy of this measure to account entirely for the relationship between Type I and Type II sensitivity is due to the differential impact of signal range in computations governing confidence. Therefore, a lack of metacognitive sensitivity quantified in this way does not necessarily indicate an impaired ability to introspect upon the effectiveness of their performance (e.g., Fleming & Lau, 2014).

These findings are focused on visual confidence, using global direction judgments to demonstrate a disproportionate influence of signal range on confidence, relative to task performance. However, we anticipate the same relationship will hold for judgments concerning other sensory attributes similarly linked to population coding, such as orientation (Pouget et al., 2000; Tolhurst, Movshon, & Dean, 1983; Zemel, Dayan, & Pouget, 1998), spatial frequency (Deneve, Latham, & Pouget, 1999), disparity (Cottereau, McKee, Ales, & Norcia, 2011) and sound localization (Fitzpatrick, Batra, Stanford, & Kuwada, 1997; Harper & McAlpine, 2004). Moreover, decisional confidence might generally be governed by these computational principles. Whether the decision be about which motion direction we have seen, about which word we have heard, or about which item we might wish to buy (Iyengar & Lepper, 2000), we anticipate that decisional confidence might be governed by the range of differently tuned units within subpopulations of neurons that contribute to evidence accumulation (Gold & Shadlen, 2007).

References

Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology, 14,* 257–262. http://dx.doi.org/10.1016/j.cub.2004.01.029

Barrett, A. B., Dienes, Z., & Seth, A. K. (2013). Measures of metacognition on signal-detection theoretic models. *Psychological Methods, 18,* 535–552. http://dx.doi.org/10.1037/a0033268

Barthelmé, S., & Mamassian, P. (2009). Evaluation of objective uncertainty in the visual system. *PLoS Computational Biology, 5,* e1000504. http://dx.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 of the United States of America, 107,* 20834–20839. http://dx.doi.org/10.1073/pnas.1007704107

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., . . . Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron, 60,* 1142–1152. http://dx.doi.org/10.1016/j.neuron.2008.09.021

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *The Journal of Neuroscience, 12,* 4745–4765.

Cottereau, B. R., McKee, S. P., Ales, J. M., & Norcia, A. M. (2011). Disparity-tuned population responses from human visual cortex. *The Journal of Neuroscience, 31,* 954–965. http://dx.doi.org/10.1523/JNEUROSCI.3795-10.2011

de Gardelle, V., & Mamassian, P. (2014). Does confidence use a common currency across two visual tasks? *Psychological Science, 25,* 1286–1288. http://dx.doi.org/10.1177/0956797614528956

de Gardelle, V., & Mamassian, P. (2015). Weighting mean and variability during confidence judgments. *PLoS ONE, 10*(3), e0120870. http://dx.doi.org/10.1371/journal.pone.0120870

de Gardelle, V., & Summerfield, C. (2011). Robust averaging during perceptual judgment. *Proceedings of the National Academy of Sciences of the United States of America, 108,* 13341–13346. http://dx.doi.org/10.1073/pnas.1104517108

De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. (2013). Confidence in value-based choice. *Nature Neuroscience, 16,* 105–110. http://dx.doi.org/10.1038/nn.3279

Deneve, S., Latham, P. E., & Pouget, A. (1999). Reading population codes: A neural implementation of ideal observers. *Nature Neuroscience, 2,* 740–745. http://dx.doi.org/10.1038/11205

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415,* 429–433. http://dx.doi.org/10.1038/415429a

Fitzpatrick, D. C., Batra, R., Stanford, T. R., & Kuwada, S. (1997). A neuronal population code for sound localization. *Nature, 388,* 871–874. http://dx.doi.org/10.1038/42246

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,* 1338–1349. http://dx.doi.org/10.1098/rstb.2011.0417

Fleming, S. M., Dolan, R. J., & Frith, C. D. (2012). Metacognition: Computation, biology and function. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences, 367,* 1280–1286. http://dx.doi.org/10.1098/rstb.2012.0021

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience, 8,* 443. http://dx.doi.org/10.3389/fnhum.2014.00443

Fleming, S. M., Maniscalco, B., Ko, Y., Amendi, N., Ro, T., & Lau, H. (2015). Action-specific disruption of perceptual confidence. *Psychological Science, 26,* 89–98. http://dx.doi.org/10.1177/0956797614557697

Fleming, S. M., Ryu, J., Golfinos, J. G., & Blackmon, K. E. (2014). Domain-specific impairment in metacognitive accuracy following anterior prefrontal lesions. *Brain: A Journal of Neurology, 137,* 2811–2822. http://dx.doi.org/10.1093/brain/awu221

Fleming, S. M., Weil, R. S., Nagy, Z., Dolan, R. J., & Rees, G. (2010). Relating introspective accuracy to individual differences in brain structure. *Science, 329,* 1541–1543. http://dx.doi.org/10.1126/science.1191883

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30,* 535–574. http://dx.doi.org/10.1146/annurev.neuro.29.051605.113038

Harper, N. S., & McAlpine, D. (2004). Optimal neural population coding of an auditory spatial cue. *Nature, 430,* 682–686. http://dx.doi.org/10.1038/nature02768

Harris, J. A., Clifford, C. W., & Miniussi, C. (2008). The functional effect of transcranial magnetic stimulation: Signal suppression or neural noise generation? *Journal of Cognitive Neuroscience, 20,* 734–740. http://dx.doi.org/10.1162/jocn.2008.20048

Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological Review, 18,* 186–201. http://dx.doi.org/10.1037/h0074579

Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology, 79,* 995–1006. http://dx.doi.org/10.1037/0022-3514.79.6.995

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience, 9,* 690–696. http://dx.doi.org/10.1038/nn1691

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association, 90,* 773–795. http://dx.doi.org/10.1080/01621459.1995.10476572

Kepecs, A., & Mainen, Z. F. (2012). A computational framework for the study of confidence in humans and animals. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences, 367,* 1322–1337. http://dx.doi.org/10.1098/rstb.2012.0037

Kiani, R., Corthell, L., & Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron, 84,* 1329–1342. http://dx.doi.org/10.1016/j.neuron.2014.12.015

Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science, 324,* 759–764. http://dx.doi.org/10.1126/science.1169405

Koivisto, M., Mäntylä, T., & Silvanto, J. (2010). The role of early visual cortex (V1/V2) in conscious and unconscious visual perception. *NeuroImage, 51,* 828–834. http://dx.doi.org/10.1016/j.neuroimage.2010.02.042

Koivisto, M., Railo, H., & Salminen-Vaparanta, N. (2011). Transcranial magnetic stimulation of early visual cortex interferes with subjective visual awareness and objective forced-choice performance. *Consciousness and Cognition, 20,* 288–298. http://dx.doi.org/10.1016/j.concog.2010.09.001

Lak, A., Costa, G. M., Romberg, E., Koulakov, A. A., Mainen, Z. F., & Kepecs, A. (2014). Orbitofrontal cortex is required for optimal waiting based on decision confidence. *Neuron, 84,* 190–201. http://dx.doi.org/10.1016/j.neuron.2014.08.039

Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *The Journal of the Acoustical Society of America, 49,* 467–477. http://dx.doi.org/10.1121/1.1912375

Li, Q., Hill, Z., & He, B. J. (2014). Spatiotemporal dissociation of brain activity underlying subjective awareness, objective performance and confidence. *The Journal of Neuroscience, 34,* 4382–4395. http://dx.doi.org/10.1523/JNEUROSCI.1820-13.2014

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience, 9,* 1432–1438. http://dx.doi.org/10.1038/nn1790

Ma, W. J., & Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience, 37,* 205–220. http://dx.doi.org/10.1146/annurev-neuro-071013-014017

Maniscalco, B., & Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition, 21,* 422–430. http://dx.doi.org/10.1016/j.concog.2011.09.021

Peirce, C. S., & Jastrow, J. (1885). On small differences in sensation. *Memoirs of the National Academy of Sciences, 3,* 73–83.

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience, 1,* 125–132. http://dx.doi.org/10.1038/35039062

R Core Team. (2014). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. Retrieved from http://www.R-project.org/

Rahnev, D. A., Maniscalco, B., Luber, B., Lau, H., & Lisanby, S. H. (2012). Direct injection of noise to the visual cortex decreases accuracy but increases decision confidence. *Journal of Neurophysiology, 107,* 1556–1563. http://dx.doi.org/10.1152/jn.00985.2011

Ruzzoli, M., Abrahamyan, A., Clifford, C. W., Marzi, C. A., Miniussi, C., & Harris, J. A. (2011). The effect of TMS on visual motion sensitivity: An increase in neural noise or a decrease in signal strength? *Journal of Neurophysiology, 106,* 138–143. http://dx.doi.org/10.1152/jn.00746.2010

Solomon, J. A., Cavanagh, P., & Gorea, A. (2012). Recognition criteria vary with fluctuating uncertainty. *Journal of Vision, 12,* 2. http://dx.doi.org/10.1167/12.8.2

Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research, 23,* 775–785. http://dx.doi.org/10.1016/0042-6989(83)90200-6

Vickers, D. (1979). *Decision processes in visual perception*. New York, NY: Academic Press.

Volkmann, J. (1934). The relation of time of judgment to certainty of judgment. *Psychological Bulletin, 31,* 672–673.

Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature, 447,* 1075–1080. http://dx.doi.org/10.1038/nature05852

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,* 1310–1321. http://dx.doi.org/10.1098/rstb.2011.0416

Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computation, 10,* 403–430. http://dx.doi.org/10.1162/089976698300017818

Zylberberg, A., Barttfeld, P., & Sigman, M. (2012). The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience, 21*. http://dx.doi.org/10.3389/fnint.2012.00079

Zylberberg, A., Roelfsema, P. R., & Sigman, M. (2014). Variance misperception explains illusions of confidence in simple perceptual decisions. *Consciousness and Cognition, 27,* 246–253. http://dx.doi.org/10.1016/j.concog.2014.05.012

Received May 6, 2015
Revision received October 15, 2015
Accepted October 16, 2015 ■