Title: Effects of expected task difficulty on metacognitive confidence and multitasking

Authors: Maxine T. Sherman^{1,2,3} & Anil K. Seth^{1,3,4}

Affiliations:

[1] Sackler Centre for Consciousness Science

[2] Brighton and Sussex Medical School, UK

[3] Department of Informatics, University of Sussex, UK

[4] Canadian Institute for Advanced Research, Program on Brain, Mind and Consciousness,
Toronto, Canada

Corresponding author: Maxine T. Sherman, maxinesherman@gmail.com

Keywords: confidence; metacognition; dual-task; multitasking; self-efficacy; perception; consciousness; task difficulty

Acknowledgements. Both authors are grateful to the Dr. Mortimer and Theresa Sackler Foundation, which supports the Sackler Centre for Consciousness Science.

Abstract

In daily life, repeated experiences with a task (e.g. driving) will generally result in the development of a belief about one's ability ("*I am a good driver*"). Here we ask how such beliefs, termed self-efficacy, interact with metacognitive confidence judgements. Across three pre-registered experiments, participants performed a perceptual discrimination task and reported their decision confidence. We induced contextual beliefs about performance (our operationalisation of self-efficacy) by manipulating the prior probability of an easy or hard trial occurring in each block. In Experiment 1 easy and hard trials generated the same levels of performance (a "subjective difficulty" manipulation), whereas in Experiments 2 and 3 performance differed across difficulty conditions (an "objective difficulty" manipulation). Results showed that context (self-efficacy) and difficulty interacted multiplicatively, consistent with the notion that confidence judgements combine decision evidence with a prior (contextual) belief on being correct. This occurred despite context having no corresponding effect on performance. We reasoned that performing tasks in easy contexts may reduce cognitive "load", and tested this, in Experiment 3, by instructing participants to perform two tasks concurrently. Consistent with a reduction in load, the effects of context transferred from influencing *confidence* on our primary task to improving *performance* on the secondary task. Taken together, these studies reveal that contextual beliefs about performance facilitate multitasking, potentially by reducing the load of tasks believed to be easy, and they extend psychophysical investigations of perceptual decision-making by incorporating 'higher-order' beliefs about difficulty context, corresponding to intuitive notions of self-efficacy.

Introduction

Decisions made under uncertainty are usually accompanied by a subjective feeling of confidence, and these confidence judgements are important for guiding many aspects of decision-making: high confidence accompanies perceptual learning (Guggenmos, Wilbertz, Hebart, & Sterzer, 2016; Zizlsperger, Kümmel, & Haarmeier, 2016), signals a stopping point for evidence accumulation, and allows us to identify and learn from mistakes (Charles, Van Opstal, Marti, & Dehaene, 2013; Murphy, Robertson, Harty, & O'Connell, 2015). These benefits can be seen in everyday life, for example, when choosing when to cross a busy road it would be wiser to wait until you are highly certain that it is safe than to cross as soon as evidence favours a safe over an unsafe crossing. We also see an essential function of metacognitive reports in social interactions: explaining why we have made a particular action, or reporting and explaining abstract ideas both require the ability to reflect upon our mental states (Frith, 2012; Shea et al., 2014).

Understanding how confidence judgements are formed is thus key to understanding decision making, yet the factors that influence confidence remain incompletely characterized. One key issue is how confidence can be constructed in the absence of explicit feedback. While in many cases we receive feedback from the world that signals whether a decision was “correct” or “incorrect” (for example, crossing the road and finding that you misjudged the cars’ speeds), such feedback is by no means universal. Yet, even in the absence of explicit feedback, it is clear that we can make effective confidence judgements: lab-based studies repeatedly and robustly show that confidence is highly correlated with decision accuracy, even when no feedback is given (Grimaldi, Lau, & Basso, 2015).

How do we form our confidence judgements in the absence of explicit feedback? In the context of perceptual decisions, there is considerable evidence that confidence is influenced by many factors that also influence the perceptual decision itself. Notably, confidence judgements are shaped by both sensory (Fetsch, Kiani, Newsome, & Shadlen, 2014; Spence et al., 2015) and environmental (Zylberberg, Fetsch, & Shadlen, 2016) uncertainty, by expectations about target identity (Alilović, Slagter, & van Gaal, 2020; Sherman, Seth, Barrett, & Kanai, 2015), and by endogenous attention (Denison, Adler, Carrasco, & Ma, 2018; Kurtz, Shapcott, Kaiser, Schmiedt, & Schmid, 2017; Schröger, Marzecová, & SanMiguel, 2015; Zizlsperger, Sauvigny, & Haarmeier, 2012). These studies indicate that stronger decision evidence generally corresponds to higher confidence because – at least to some degree – both objective and metacognitive reports, unsurprisingly, use overlapping sources of information.

Other studies reveal additional factors, beyond decision evidence, that influence confidence

judgements. For example, properties of motor responses (Fleming et al., 2015) and sensitivity to stimulus variance (Zylberberg, Roelfsema, & Sigman, 2014) both differentially affect objective decisions and confidence judgements.

One potentially important but understudied factor in determining subjective decision confidence is *self-efficacy*. Self-efficacy has a long history in psychology, though it has been rather neglected in the recent perceptual decision-making and (cognitive) metacognition literature. It refers to a person's belief that they can perform some task to a particular standard (Bandura 1977), usually assessed via questionnaires. Bandura proposed that the extent to which one will maintain effort and motivation to complete a task is mediated by their self-efficacy, such that individuals will not persevere at tasks which they believe they cannot perform adequately in some reasonable timeframe (Bandura, 1982). Continuing with the example of crossing the road, if a person has repeated negative experiences attempting to cross safely then they will likely build a belief that they are “bad” at this task, i.e., have low self-efficacy in road-crossing. A reasonable prediction then is – all else being equal – that this individual will have lower confidence for any specific road-crossing decision. Furthermore, when the road is particularly busy (i.e. road-crossing is objectively hard) they may be more likely than someone with high self-efficacy to cross elsewhere, at a designated crossing, rather than waiting for a seemingly clear path.

The prediction of a positive association between self-efficacy and subjective decision confidence, though seemingly trivial, has until recently received mostly circumstantial support. First, confidence on a particular trial is higher if confidence on the previous trial is high, even if the two trials came from different perceptual tasks (Rahnev, Koizumi, Mccurdy, Esposito, & Lau, 2015), and similarly, when the variance of a stimulus array is primed, confidence for a decision on that stimulus is higher (Michael, de Gardelle, & Summerfield, 2014). Second, at least some information that is relevant for confidence judgements must be available prior to stimulus presentation, since prospective confidence reports correlate with both decision accuracy and with retrospective confidence (Siedlecka, Paulewicz, & Wierzchoń, 2016). Both these lines of evidence are suggestive that some form of self-efficacy is associated with confidence, though they do not test this association directly.

One recent study gets at the issue a little more directly, Rouault and colleagues found that, given a choice between two tasks, participants will disproportionately choose the one they believe they are performing better in (Rouault, Dayan, & Fleming, 2019; Rouault & Fleming, 2020), indicating that people both develop and utilise beliefs about their performance. They then showed that if participants learn the distribution of either their external feedback, or their confidence ratings over the course of a block, then ratings of self-efficacy at the end of a block can be constructed

from those distributions. In other words, self-efficacy ratings can be modelled as expected task performance, learned either from external feedback or indirectly from confidence judgements. These findings indicate that there is some “global” (Rouault et al., 2019) subjective estimate of performance that is separable from trial-specific factors, and which is instead associated with the individual’s prior history of inferring their overall performance (self-efficacy) on a task. However, this study does not explicitly manipulate self-efficacy in order to test its influence on confidence.

We conducted three pre-registered experiments to explicitly test how self-efficacy and task difficulty interact to influence subjective decision confidence, in the realm of perceptual decisions. We operationalized self-efficacy by inducing (statistical) expectations that task performance will be high versus low, learned over the course of a block. In other words, we manipulated *difficulty context* so that participants performed blocks where either most of trials were relatively easy (easy context, i.e. high self-efficacy), or most were relatively difficult (hard context, i.e. low self-efficacy). This allowed us to compare both objective performance and subjective confidence as a function of both task difficulty and statistical beliefs about task performance – the latter being self-efficacy.

In Experiment 1, we aimed to manipulate task difficulty without changing task accuracy. We did this by presenting grating stimuli with particular contrast ratios that in turn induce either high or low confidence but which maintain (broadly) equal performance (Koizumi, Maniscalco, & Lau, 2015; Peters et al., 2017). Because performance did not change, these two classes of grating constituted a “subjective difficulty” manipulation. In Experiment 2 we used random dot motion stimuli instead. Here, easy and hard trials were simply trials with respectively low or high dot motion variability – an “objective difficulty” manipulation. Finally, in Experiment 3 participants performed a dual-task version of Experiment 2 in order to test whether effects of difficulty context on confidence differ during multitasking relative to single-task designs.

Methods

Experiment One

Experiment 1 tested whether high self-efficacy increases confidence in perceptual decision making. Participants performed an orientation (left/right) discrimination task. On each trial, they were presented with a grating embedded in visual noise, and gave a binary orientation report followed by a decision confidence rating on a continuous scale (from 50% chance to 100% fully confident). On each trial the grating was either subjectively easy or subjectively hard. Subjectively easy/hard trials were those on which the grating stimulus elicited higher/lower confidence, but not higher/lower accuracy (details below). To manipulate difficulty context we created blocks with either mostly subjectively easy trials or mostly subjectively hard trials, the

idea being that exposure to many subjectively easy/hard trials would drive beliefs that one is performing well/poorly (i.e., would modulate self-efficacy).

Experiment One, Participants

Data were collected from 46 naïve participants (18-28 years old, mean 19.9 years, 39 female, 1 non-binary) with normal or corrected-to-normal vision. This sample size was determined by a Bayesian stopping rule, in which we collected data until a minimum of 20 participants were retained and until the Bayes factor for the effect of self-efficacy context on confidence was sensitive. We calculated Bayes Factors at the end of each testing day.

In line with our pre-registered exclusion criteria, 20 participants were excluded for having extreme decision or confidence criteria (defined as a type 1 or 2 hit or false alarm rate either less than 5% or greater than 95%), and a further 2 participants for having a meta- $d' \leq 0$ (see *Statistics*, below, for an explanation of meta- d'). This left us with 24 participants for analysis. Note that all results were qualitatively the same when relaxing the first criterion to requiring the presence of variability in confidence ratings in each condition (which leaves 39 participants: 4 excluded for mean $d' \leq 0$, and a further 3 for meta- $d' \leq 0$).

Experiment One, Stimuli

Stimuli were presented on a 19-in CRT monitor (60 Hz, 1240 × 1048 resolution) using Psychtoolbox for Matlab on Windows 7. Participants sat 60 cm from the monitor. We used the grating stimuli described in Koizumi et al (Koizumi et al., 2015) and depicted in Fig. 1A.

These grating stimuli are built from the superimposition of three components: one Gabor patch oriented at +45°; another Gabor oriented at -45° (both Gabors had diameter 6.6 dva, SD 1.6 dva, phase 0°, sf 1cpd); and a white noise patch. The absolute contrast of one Gabor was greater than the other. The task was to report the orientation of the compound stimulus. On this task, the correct response corresponds to the orientation (left/right) of the higher contrast Gabor. “Positive evidence” refers to the contrast level of this higher-contrast Gabor. “Negative evidence” refers to the contrast level of the lower-contrast Gabor.

As Koizumi et al demonstrated, using these compound stimuli one can increase confidence without changing accuracy by increasing the amount of positive evidence, and compensating for this increase in signal (to maintain accuracy at the same level) by also increasing the amount of negative evidence. In this way, one can create two stimulus sets from these gratings: one that generates low confidence (because the amount of positive evidence is low), and one that generates high confidence (because positive evidence is high). With careful adjustment of the other two components, negative evidence and noise contrast, the high and low confidence

Figure 1. *Experimental designs.* **A-C.** Paradigm for Experiment 1, which used a left/right orientation discrimination task. Panel A depicts examples of our subjectively hard and subjectively easy stimuli (low and high positive evidence, PE, respectively), matched for objective performance d' . Panels B and C depict the procedure and trial sequence respectively. **D-E.** Procedure and trial sequence respectively for Experiment 2, which used a left/right motion discrimination task. The main changes here were: first, objective rather than subjective difficulty was manipulated; second, the task was motion rather than orientation discrimination; and third, a control condition was added where easy versus hard trials were equally likely to appear. **F-G.** Procedure and trial sequence for Experiment 3. This was almost identical to Experiment 2, except for the addition of a dual-task component. In the secondary task, participants had to report the identity of a peripheral bisected disc while performing the motion discrimination task.

stimuli can generate equal accuracy.

Below we describe Koizumi et al.'s (2015) equations for jointly adjusting the three components. Following Koizumi et al, the following variables were defined, C denotes contrast and R denotes a ratio:

$$C_{Total} = C_{NE} + C_{PE} + C_{Noise} \quad (1)$$

$$R_{PE:Noise} = C_{PE} / C_{Noise} \quad (2)$$

$$R_{NE:PE} = C_{NE} / C_{PE} \quad (3)$$

$R_{PE:Noise}$ is the ratio between the absolute contrast of the “correct” Gabor and the visual noise. All else being equal, confidence will be higher when this ratio takes higher values. $R_{NE:PE}$ is the ratio between the incorrect and correct Gabor contrasts. C_{Total} was set to 1, meaning that the three component contrasts could be expressed as follows, in terms of the two ratios defined above:

$$C_{Noise} = 1 / (1 + R_{PE:Noise} \times (1 + R_{NE:PE})) \quad (4)$$

$$C_{PE} = R_{PE:Noise} / (1 + R_{PE:Noise} \times (1 + R_{NE:PE})) \quad (5)$$

$$C_{NE} = R_{PE:Noise} \times R_{NE:PE} / (1 + R_{PE:Noise} \times (1 + R_{NE:PE})) \quad (6)$$

The goal is to have one set of contrasts with high C_{PE} and one with low C_{PE} , but to keep accuracy broadly constant across the two sets. Koizumi and colleagues found that one could fix $R_{NE:PE}$ at either a high (0.7) or low (0.35) value and titrate $R_{PE:Noise}$ for each participant until accuracy for the two sets (the two values of $R_{NE:PE}$) are equated. We used the low positive evidence stimuli (that generate relatively low confidence) as “subjectively hard” gratings, and the high positive evidence stimuli as “subjectively easy” gratings. The mean values (over all 46 participants) for each variable described in Eqs. 1-6 are given in Table 1. Example subjectively easy and hard stimuli are shown in Fig. 1A.

Table 1. Mean contrasts and contrast ratios for the subjectively easy and subjectively hard stimuli

	$R_{NE:PE}$	$R_{PE:Noise}$	C_{PE}	C_{NE}	C_{Noise}	Mean confidence
Subjectively Easy	0.7	0.29	0.19	0.14	0.67	78%
Subjectively Hard	0.35	0.14	0.11	0.04	0.85	73%

Experiment One, Procedure

The procedure is illustrated in Fig. 1B. After obtaining informed consent, participants were familiarized with the subjectively easy ($R_{NE:PE} = 0.35$) and subjectively hard ($R_{NE:PE} = 0.7$) stimuli before completing approximately 10 minutes of training on the orientation discrimination task. Here, participants completed one block of 50 training trials for each subjective difficulty condition. The correct orientation (left or right) was randomly determined on each trial ($P(\text{left}) = .5$). Block order (easy versus hard) was randomised across subjects. So that participants improved over the course of training, trials became more difficult over the course of each block: $R_{PE:Noise}$ (which governs task difficulty) was titrated by an adaptive QUEST staircase (beta = 11/13 for

subjectively easy/difficult trials respectively) (Watson & Pelli, 1983) that converged to 72% correct. On each trial, participants gave their orientation judgement followed by a confidence rating. Reports were followed by the presentation of valid feedback on the accuracy of their left/right response (feedback presented for one second).

Following training participants performed two staircase procedures to equate performance across subjectively easy and difficult trials. As with the training blocks, $R_{PE:Noise}$ was titrated separately for each level of subjective difficulty to achieve 72% correct. Participants performed staircase 80 trials for each staircase. These trials were identical to the training trials, except that confidence was not reported and feedback was not given.

Finally, participants performed 6 blocks of 80 experimental trials (960 total, see figure 1B). On each trial either a subjectively easy or subjectively hard stimulus was presented. To investigate effects of self-efficacy on confidence, we changed the context in which the stimulus was presented: trials were either performed in blocks where subjectively easy trials were more probable (with 75% probability, easy context blocks) or where subjectively hard trials were more probable (with probability 75%, hard context blocks). Participants were not told which context they were in. The idea here was that over the course of each block participants would learn to expect easy or hard trials, and accordingly, would develop high or low self-efficacy for their performance on the orientation task. Orientation and difficulty were fully counterbalanced within each block. Block order (easy or hard context blocks) were randomised.

Experiment One, Trial structure

As depicted in Figure 1C, each trial began with a blank screen for a random interval of 300-600ms. Next, the target stimulus (diameter 6dva) was briefly presented foveally for 2 frames (33.33ms), followed by a blank screen for 300ms and then a response prompt. Here, participants pressed the d or f keys to report whether the left-oriented or right-oriented Gabor had higher contrast. Finally, a vertical visual analogue scale appeared, extending from 50% to 100%. Using the up and down arrow keys, participants moved a slider (with random starting location) along the scale to indicate their confidence in their orientation judgement.

Experiment Two

Experiment 2 also manipulated stimulus difficulty and difficulty context. However, in contrast to Experiment 1 here we used a motion discrimination task and manipulated objective, not subjective difficulty. Objective difficulty means that easy/hard trials were associated with better/poorer performance. We manipulated objective difficulty so that easy and hard stimuli were matched on all features except stimulus noise. On each trial, participants were presented

with a random dot motion stimulus, where the mean dot motion trajectory was either -45° or $+45^\circ$ (“up-left” or “up-right”) and the variance was either low or high (easy or hard trials). Then, as in Experiment 1, we manipulated difficulty context (self-efficacy) by creating blocks with either mostly easy or mostly hard trials. We also added a control block, in which 50% of trials were easy and 50% of trials were hard.

Experiment Two, Participants

Data were collected from 25 naïve healthy University of Sussex undergraduate and postgraduate students (18-33 years old, mean 20.8 years, 22 female, 1 non-binary) with normal or corrected-to-normal vision. None had participated in Experiment 1. This sample size was determined by a Bayesian stopping rule, in which we collected data until a minimum of 20 participants were retained and until the Bayes factors for both the effect of context (self-efficacy) on confidence, and the difficulty by context interaction under confidence were sensitive. In line with our pre-registered exclusion criteria, 3 participants were removed due to staircase failure (smaller motion SD on difficult than easy trials) and a further 1 was removed for having a difference in d' between easy and hard trials that was more than 2 SD than the mean (2.2 SD). This left 20 participants for analysis.

Experiment Two, Stimuli

The experiment was presented on a 19-in CRT monitor (100 Hz, 1240 x 1048 resolution) using Psychtoolbox for Matlab on Ubuntu. We used a random dot kinematogram (RDK) stimulus. The stimulus comprised 250 moving white dots (radius 0.1dva) presented in a transparent, square bounding box (8.3dva x 8.3dva) on a black background. Dot motion trajectories were drawn from a circular generative Gaussian with central tendency -45° or $+45^\circ$, corresponding to a mean motion direction of upper left or right diagonal. The dispersion of the generative distribution was set for each participant by a staircase procedure (see below).

Experiment Two, Trial sequence

Trials began with the presentation of a white fixation cross (0.5dva x 0.5dva) for a random interval between 0 and 500ms. Next, the RDK and fixation were presented for 300ms. Immediately following stimulus offset a response prompt appeared, and participants reported left or right average motion using the d and f keys. Immediately following the motion direction report, a vertical visual analogue scale appeared identical to that used in Experiment 1. Participants rated their decision confidence on a 50-point 50-100% scale by moving a slider with the up and down arrow keys (Fig. 1D). Note that, in contrast to Experiment 1 in which each trial involved two superimposed Gabors, here each trial involved presentation of a single RDK.

Experiment Two, Procedure

The procedure is illustrated in Fig. 1D. After obtaining informed consent, participants performed approximately 7 minutes of practice trials. At the end of each trial, participants were presented with one second of on-screen feedback about their motion discrimination accuracy. Feedback was only given in these training trials. The difficulty of the practice trials, parametrised as the variance of the generative distribution for dot motion trajectories, was governed by a 3 down 1 up psychophysical staircase procedure (initial variance 45°, step size 10°, range 1° - 180°) that terminated after 5 reversals.

Following training, participants completed two psychophysical staircase procedures in random order. In contrast with Experiment 1 where we titrated both stimuli to generate approximately 72% accuracy, here we wanted easy and hard stimuli to be associated with high and low performance respectively. The "easy" staircase determined the variance required for approximately 80% accuracy (4 down 1 up, initial variance 60°, step size 5° range 1° - 180°) and the "hard" staircase set performance to approximately 60% accuracy (2 down 1 up, initial variance 75°, step size 5° range 1° - 180°). Confidence was not reported in staircase trials. Staircasing lasted approximately 7 minutes.

The final ~45 minutes of the session comprised experimental trials. Each block consisted of 56 trials, in which the probability of an easy (versus hard) trial was either 0.25 (low self-efficacy condition), 0.50 (control condition) or 0.75 (high self-efficacy condition). Motion direction (left/right) and difficulty (easy/hard) was fully counterbalanced within each block, and blocks were presented in fully randomised order. Participants completed between 3 and 5 repetitions of each block type.

Experiment Three

The purpose of Experiment 3 was to test effects of self-efficacy on confidence in a dual-task setting. The paradigm was identical to that of Experiment 2, except that on each trial participants performed two tasks concurrently: the first was the motion discrimination task previously described; the second was a disc discrimination task. For this secondary task a bisected red and green disc was briefly presented during the trial in the lower-left quadrant of the screen, and participants had to judge its identity (red on the left and green on the right, or vice versa). The difficulty of this disc discrimination task was held constant across the experiment.

Experiment Three, Participants

Data were collected from 34 naïve healthy University of Sussex undergraduate and postgraduate students (18-43 years old, mean 25.5 years, 25 female, missing demographic data for 2 participants) with normal or corrected-to-normal vision and who were not red-green colour-blind.

None had participated in Experiments 1 or 2. In line with our pre-registered exclusion criteria we excluded 12 participants: 4 for completing fewer than 500 trials in the one hour session due to slow responding, 1 for negative d' on the disc task, and 7 for their RDK performance (negative d' in any condition, higher d' or lower motion SD on hard than easy trials, or d' difference between easy and hard trials > 2 SD from mean). This left 22 participants for analysis.

Experiment Three, Stimuli

The experiment was presented on a 19-in CRT monitor (100 Hz, 1240 x 1048 resolution) using Psychtoolbox for Matlab on Ubuntu. We used an identical RDK to that in Experiment 2, with one exception: at 100ms to 250ms (in 50ms intervals) post-RDK onset, a small bisected red-green disc appeared at the lower-left corner of the RDK bounding box for 100ms. The disc was either red on the left and green on the right, or vice versa.

Experiment Three, Trial Structure

Trials began with the presentation of a white fixation cross (0.5dva x 0.5dva) for a random interval between 0 and 500ms. Next, the RDK and fixation were presented for 350ms. At 100, 150, 200 or 250ms after RDK onset the bisected red-green disc appeared. Its onset time was randomly determined on each trial (uniform distribution) and was removed after 100ms. Immediately following RDK offset a response prompt appeared, and participants reported left or right average motion using the d and f keys. Immediately following the motion direction report, a vertical visual analogue scale appeared identical to that used in Experiment 1. Participants rated their decision confidence on a 50-point 50-100% scale by moving a slider with the up and down arrow keys. Finally, using the d and f keys participants reported which of the two discs had been presented in the periphery (Fig. 1F).

Experiment Three, Procedure

The procedure is illustrated in Fig. 1G. After obtaining informed consent, participants performed approximately 7 minutes of practice trials. During this training period, no bisected disc was presented and no disc reports were given (i.e., there was no secondary task). At the end of each trial, participants were presented with one second of on-screen feedback about their motion discrimination accuracy. Feedback was only given in these training trials. The difficulty of the practice trials, parametrised as the variance of the generative distribution for dot motion trajectories, was governed by a 3 down 1 up psychophysical staircase procedure (initial variance 45°, step size 10°, range 1°-180°) that terminated after 5 reversals.

Following training, participants completed two psychophysical staircase procedures in random order. The "easy" staircase determined the variance required for approximately 80% accuracy (4 down 1 up, initial variance 60°, step size step size 5°, range 1°-180°) and the "hard" staircase set performance to approximately 60% accuracy (2 down 1 up, initial variance 75°, step size step

size 5°, range 1°-180°). Confidence was not reported in staircase trials. No bisected disc was presented and no disc reports were given. Staircasing lasted approximately 7 minutes. After stimulus titration, participants completed 20 training trials on the full dual-task paradigm (i.e. with presentation of the bisected disc). Task difficulty was set according to the "easy" (i.e. low variance) threshold acquired during staircasing. After reporting motion discrimination confidence participants received valid feedback on their decision accuracy for 500ms. Next, participants reported their disc discrimination judgement, after which they received 500ms of valid feedback on this report.

The final ~45 minutes of the session comprised experimental trials. Each block consisted of 56 trials, in which the probability of an easy (versus difficult) trial was either 0.25 (low self-efficacy condition), 0.50 (control condition) or 0.75 (high self-efficacy condition). Motion direction (left/right) and difficulty (easy/hard) was fully counterbalanced within each block and blocks were presented in fully randomised order. Disc identity (red-green or green-red) was randomly determined on each trial. Participants completed between 3 and 5 repetitions of the three block types.

Statistics (all Experiments)

A comprehensive analysis plan was pre-registered on the Open Science Framework. This, alongside raw data and analysis code, can be found at <https://osf.io/vxjp8> (pre-registration for Experiment 1: <https://osf.io/yhv2k>, pre-registration for Experiments 2 and 3: <https://osf.io/b9fre>). Pre-registered versus exploratory analyses are clearly delineated in the Results section.

To quantify objective performance we used the Signal Detection Theory (SDT) measure d' (Green & Swets, 1966). This represents the degree to which people can discriminate between two stimulus classes, S1 and S2 (e.g. left/right) when they make a binary report, R1 or R2. To calculate d' , first we count the number of hits (R1 | S1), misses (R2 | S1), false alarms (R1 | S2) and correct rejections (R2 | S2). We padded each cell with 0.5 to avoid having zero counts (the log-linear rule, Goodman, 1970). Then, d' is calculated as follows, where Z is the standard z -score:

$$\text{hit rate} = \frac{\text{hits} + 0.5}{\text{hits} + \text{misses} + 1} \quad (7)$$

$$\text{false alarm rate} = \frac{\text{false alarms} + 0.5}{\text{false alarms} + \text{correct rejections} + 1} \quad (8)$$

$$d' = Z(\text{hit rate}) - Z(\text{false alarm rate}) \quad (9)$$

To quantify metacognitive performance we used the measure m -ratio, which reflects the SDT-optimality of participants' type 2 (metacognitive) performance. Higher values of m -ratio correspond to better metacognition. Specifically, m -ratio is given by $\text{meta-}d'/d'$, where $\text{meta-}d'$ is

the type 1 d' of the SDT-optimal observer who uses all of the type 1 information to form their confidence (type 2) judgement. In this way, m-ratio reflects how suboptimal the participant is, because lower values correspond to the SDT-optimal observer being able to accomplish the participant's type 2 (metacognitive) performance with less type 1 (objective) sensitivity.

To estimate m-ratio we used Matlab code provided by Maniscalco & Lau (2015). This deviates from our Experiment 1 pre-registration, where we stated we would run analyses in using the meta-SDT R package (an R implementation of the Matlab code we used). We made this change due to incompatibility issues with the R package.

For all analyses, inferences were based on Bayes Factors (BFs) computed on main effect and interaction terms. BFs were interpreted in line with Jeffreys (Jeffreys, 1961), such that $BF \leq 1/3$ or $BF \geq 3$ corresponds to evidence for the null hypothesis (H_0) and alternative hypothesis (H_1) respectively. Bayes Factors between $1/3$ and 3 are interpreted as 'insensitive'. Priors for interaction terms were set as uniform distributions from 0 to the largest simple effect (Dienes, 2014). For readers unfamiliar with Bayes, frequentist statistics are also presented (but not interpreted). The following notation is used to describe priors:

- $BF_{U(a,b)}$ is the Bayes Factor when the prior is a uniform distribution with lower-bound a and upper bound b .
- $BF_{N(a,b)}$ is the Bayes Factor when the prior is a normal distribution with mean a and standard deviation (SD) b .
- $BF_{H(a,b)}$ is the Bayes Factor when the prior is a half-normal distribution with mean a and SD b .

Each Bayes Factor is presented with its corresponding robustness region (RR), which specifies the values of the prior SD (for half-normal priors) or the prior's upper bound (for uniform priors) for which the interpretation of the BF (i.e. evidence for the null/insensitive/evidence for H_1) remains the same. Justifications for all priors are presented in S1.

Unless otherwise stated, reported p -values are taken from repeated-measures ANOVAs, which were run on Matlab 2018a using the Statistics Toolbox and custom code, and p -values for linear mixed model were acquired using the *lmer* and *car* packages in R.

Results

Experiment One: Confidence is more sensitive to difficulty in subjectively easy contexts (exploratory exclusion criteria)

In Experiment 1, we tested whether performing an orientation discrimination task in subjectively easy contexts, where participants should believe they are doing well overall (have high self-efficacy), leads to higher confidence (Fig. 1A-C). To do this we created blocks of trials where 75% of grating stimuli generated low confidence (i.e. subjectively hard trials), and other blocks where 75% of stimuli generated high confidence (i.e. subjectively easy trials). We reasoned that if participants repeatedly encountered instances of (perceived) high performance, i.e. frequently believe they are performing well, they will develop high self-efficacy for the task. In other words, we attempted to manipulate self-efficacy by changing the task statistics rather than through a manipulation such as false feedback.

The grating stimuli used, introduced in Koizumi et al. as high and low “positive evidence” stimuli (Koizumi et al., 2015), were chosen because of their capacity to induce different levels of confidence for matched decision accuracy. This allowed us to examine confidence across levels of (subjective) difficulty, and difficulty contexts, without confounding effects of differences in task performance. Because these stimuli should not change objective accuracy we call them *subjectively* (rather than objectively) easy and hard.

We made two changes to our pre-registered analysis plan. First, for some participants d' was not successfully equated between subjectively easy and hard trials: for five, the absolute difference was greater than 1 ($d'_{diff} = 1.08, 1.24, 1.41, 1.74$). These participants were excluded (not pre-registered), leaving 19 for analysis. Second, in our pre-registration we planned to analyse trial-by-trial confidence ratings in a linear mixed-model (LMM), however this model did not converge under any choice of optimizer. Therefore, condition-averaged confidence was analysed with a repeated-measures ANOVA. Full statistical results from our pre-registered analyses can be found in Table S1, and analysis of metacognitive sensitivity measure m -ratio can be found in the Supplementary Materials (SM2). To accompany the Bayesian analysis presented below, frequentist statistical results (i.e. p -values) can be found in Table 2.

To determine whether d' (objective orientation discrimination performance) was equated across conditions, we set a strict minimal interesting difference of -0.06 to 0.06 , below which values are deemed effectively equal (justification in SI). As expected for the stimuli used, d' was equated both between subjectively easy and hard gratings, (no main effect of difficulty: $BF_{MID(\pm 0.06)} = 0.26$, Table 2, row a), and between the easy and hard context blocks (no main effect of context: $BF_{MID(\pm 0.06)} = 0.01$, Table 2, row b, Fig. 2A). However, the subjective difficulty x context interaction

Table 2. Exploratory results from Experiment 1.

	M_{diff}	SE_{diff}	Prior	<u>Bayesian</u>		<u>Frequentist</u>	
				BF	RR	F	p
d'							
a. Subjective difficulty	0.057	0.106	MID(± 0.06)	0.26	[0.11, ∞)	0.290	.597
b. Difficulty context	0.023	0.070	MID(± 0.06)	0.01	[0.06, ∞)	0.107	.748
c. Self-efficacy x subjective difficulty	-0.229	0.076	U(0,0.06)	3.29	[0.05, 5.89]	9.013	.008
Confidence*							
d. Accuracy	8.922	1.202	-	-	-	65.000	< .001
e. Subjective difficulty	3.927	0.907	H(0,0.25)	4038	[0.002, ∞)	21.354	< .001
f. Difficulty context	0.865	0.487	H(0,0.25)	0.19	[0.14, ∞)	2.584	.125
g. Accuracy x subjective difficulty	-0.734	0.927	U(0,0.13)	0.52	($-\infty$, 0.21]	1.590	.223
h. Accuracy x context	-0.111	0.476	U(0,0.13)	0.10	[0.04, ∞)	0.191	.667
i. Subjective difficulty x context	1.422	0.421	H(0,0.06)	27.76	[0.004, 0.54]	9.106	.007
j. Accuracy x context x subjective difficulty	0.855	1.300	U(0,0.02)	1.21	($-\infty$, 0.14]	0.510	.484
m-ratio							
k. Subjective difficulty	-0.019	0.068	U(0,0.50)	0.217	[0.32, ∞)	0.082	.778
l. Difficulty context	-0.086	0.073	U(0,0.50)	0.65	($-\infty$, 0.99]	1.384	.255
m. Subjective difficulty x context	0.169	0.112	U(0,0.09)	1.67	($-\infty$, 2.48]	2.274	.149

M_{diff} = mean difference; SE_{diff} = standard error of difference; BF = Bayes Factor; RR = robustness region; MID = minimal interesting difference; U = uniform prior; H = half-normal prior

*BFs and ANOVAs are calculated on log-confidence, and priors and RRs are given in log-confidence units. For interpretability, descriptives (M_{diff} and SE_{diff}) are given in raw (untransformed) units

exceeded our minimal interesting difference (interaction: $BF_{U(0,0.06)} = 3.29$, Table 2, row c): while d' was equated across subjective difficulty levels for easy contexts ($BF = 0.31$), it was not equated in hard contexts ($BF = 2.64$).

Our key *a priori* hypothesis was that confidence would increase under high self-efficacy. We tested this hypothesis with an accuracy x difficulty context x subjective difficulty repeated-measures ANOVA on log-confidence. First, we verified that our grating stimuli successfully manipulated confidence. This was indeed the case: confidence was higher for subjectively easy than subjectively hard gratings (main effect of difficulty: $BF_{H(0,0.25)} = 4038$, Table 2, row e).

Contrary to our hypothesis, overall confidence was unaffected by difficulty context (no main effect of context, $BF_{H(0,0.25)} = 0.19$, Table 2, row f). Instead, we found that context and subjective difficulty interacted, (interaction effect: $BF_{H(0,0.06)} = 27.76$, Table 2, row i). As can be seen in Figure 2B, the effect of subjective difficulty on confidence was larger in the easy context than the hard context. In other words, subjective difficulty exerted a larger effect on confidence in blocks in which participants expected to perform well. This held despite no corresponding pattern in objective performance, as measured by d' .

Figure 2. Results of Experiments 1-3. **A-C.** Experiment 1 results. Type 1 d' , confidence and m-ratio respectively as a function of difficulty context and subjective difficulty (hard/easy are coloured red/green respectively). **D-F.** Experiment 2 results. Type 1 d' , confidence and m-ratio respectively as a function of difficulty context and objective difficulty (hard/easy are coloured red/green respectively). **G-J.** Experiment 3 results. Type 1 d' , confidence and m-ratio for the RDK task, and type 1 d' for the disc task respectively. These are plotted as a function of difficulty context and objective difficulty (hard/easy are coloured red/green respectively). Error bars represent within-subject SEM.

In summary, results from Experiment 1 suggest that high self-efficacy amplifies the effects of subjective difficulty (in a multiplicative fashion) on confidence, without improving objective performance on the task.

Experiment Two: Replication of effect under pre-registered exclusion criteria and visually-matched stimuli

In Experiment 1 we wanted to manipulate trial difficulty and difficulty context in such a way that objective performance was kept constant. For this reason, we manipulated what we termed “subjective difficulty”. However, this manipulation was primarily effective in the easy context condition, and also had the disadvantage of necessitating stimuli that presented rather large differences in visual appearance. We therefore created a second pre-registration for Experiments 2 and 3 (described later), which set out to manipulate objective rather than subjective trial difficulty.

In Experiment 2 we manipulated objective difficulty by having easy or hard trials that were associated with relatively high or low performance respectively (Fig. 1, D-E). Participants performed a motion discrimination task where the generative distribution for the dot motion trajectories had a mean of either -22.5 or +22.5 degrees (corresponding to “up-left” or “up-right” motion), and either low or high variance (corresponding to ‘objectively easy’ and ‘objectively hard’ trials respectively). In this way easy and hard trials were matched on all features except the variance in dot motion direction, but this also meant that participants’ discrimination performance was – by design - better on easy trials than on hard trials. This was therefore an *objective* (not *subjective*) difficulty manipulation. Experiment 2 also included a control condition, where 50% of trials were easy and 50% were hard. This condition could be compared against the easy and hard context conditions to determine whether the predictability of trial difficulty influences confidence or metacognition, though we did not pre-register such tests. Full statistical results from this experiment can be found in Table 3, and results for metacognition m -ratio in the supplementary materials (SM2).

Results confirmed that this task successfully manipulated objective difficulty: by design, d' was on average 0.95 ± 0.12 higher on easy than hard trials (Fig. 2D), and was almost equal across easy and hard contexts (no main effect of context: $BF_{MID(\pm 0.06)} = 0.36$, Table 3, row b). We could not determine whether the interaction effect was within our minimal interesting difference (interaction: $BF_{H(0,0.23)} = 0.67$, Table 3, row c).

In Experiment 1, our key result was that subjective difficulty had a larger effect on (log) confidence in an easy context than in a hard context. We replicated this finding with the objective difficulty manipulation in Experiment 2 (difficulty x context: $BF_{H(0,0.02)} = 24.29$, Table 3, row i, Fig. 2E).

Participants also exhibited higher confidence on objectively easy than objectively hard trials (main effect of difficulty: $BF_{U(0.08)} = 355.05$, Table 3, row e) and in easy than hard contexts (main effect of context: $BF_{H(0.26)} = 4.14$, Table 3, row f). In Experiment 1, this effect of context on confidence was found only on subjectively easy trials (Fig. 2E vs. Fig. 2B).

Together with results from Experiment 1, these data tell us that, irrespective of whether participants are truly performing better on easy trials or not, easier contexts (higher self-efficacy) are associated with larger effects of task difficulty on confidence, without there necessarily being corresponding improvements in objective performance.

Table 3. Pre-registered results from Experiment 2.

	M_{diff}	SE_{diff}	Prior	Bayesian BF	RR	Frequentist F	p
d'							
a. Objective difficulty	0.947	0.119	-	-	-	62.934	< .001
b. Difficulty context	0.084	0.091	MID(± 0.06)	0.36	$(-\infty, 0.13]$	0.717	.494
c. Self-efficacy x subjective difficulty	-0.057	0.120	H(0,0.23)	0.67	$(-\infty, 0.55]$	0.075	.928
Confidence*							
d. Accuracy	6.027	0.830	-	-	-	51.327	< .001
e. Objective difficulty	1.534	0.366	U(0,0.08)	355.05	$[0.002, \infty)$	15.511	< .001
f. Difficulty context	2.188	0.773	H(0,0.26)	4.14	$[0.005, 0.37]$	4.949	.012
g. Accuracy x objective difficulty	2.563	0.768	U(0,0.08)	44.54	$[0.007, 1.23]$	9.889	.005
h. Accuracy x context	0.363	0.646	U(0,0.08)	0.15	$[0.04, \infty)$	1.295	.285
i. Objective difficulty x context	1.528	0.479	H(0,0.02)	24.29	$[0.003, 0.27]$	5.555	.007
j. Accuracy x context x objective difficulty	2.437	0.912	U(0,0.03)	16.68	$[0.01, 0.42]$	3.710	.033
m-ratio							
k. Subjective difficulty	-0.034	0.059	U(0,0.50)	0.25	$[0.37, \infty)$	0.328	.573
l. Difficulty context	0.058	0.082	U(0,0.50)	0.40	$(-\infty, 0.61]$	1.395	.260
m. Subjective difficulty x context	-0.015	0.127	U(0,0.06)	0.99	$(-\infty, 0.54]$	0.067	.935

M_{diff} = mean difference; SE_{diff} = standard error of difference; BF = Bayes Factor; RR = robustness region; MID = minimal interesting difference; U = uniform prior; H = half-normal prior

*BFs and ANOVAs are calculated on log-confidence, and priors and RRs are given in log-confidence units. For interpretability, descriptives (M_{diff} and SE_{diff}) are given in raw (untransformed) units

Experiment Three: Under dual-task conditions, self-efficacy effects transfer to the secondary task

One reason why performing tasks in easy contexts (high self-efficacy) may make confidence more sensitive to difficulty (without driving corresponding improvements in performance) is that high self-efficacy involves a reduction in the cognitive ‘load’ associated with completing the task (Lavie, Hirst, Fockert, & Essi Viding, 2004). On this view, low load would make the task feel less effortful, driving higher confidence particularly for those trials which were indeed easy. This possibility can be tested using a dual-task design. If high self-efficacy indeed reduces load then there should be “spare capacity” for other tasks (“distractor” tasks), and so performance on these other tasks should improve with lower load (higher self-efficacy).

We tested this hypothesis in our pre-registered Experiment 3. Participants performed the same motion discrimination task as in Experiment 2, but additionally had to determine the identity of a bisected red-green disc presented in the lower-right quadrant of the screen (the “distractor” task), reporting whether it was green-red or red-green (Fig. 1, F-G). Participants were instructed to prioritise the motion discrimination task. We predicted that in the presence of this secondary task

we would no longer see larger effects of difficulty on confidence in easy context blocks (compared to the hard context). We predicted that instead, this context x difficulty interaction would be seen on secondary task performance. Full statistical results can be found in Table 4, and results on metacognitive sensitivity m-ratio in the supplementary materials (SM2).

During the staircase procedure participants were instructed to ignore the bisected disc in the periphery, meaning that we titrated motion discrimination difficulty under single-task conditions. Nonetheless, in the dual-task experimental trials, participants still performed better on easy than hard trials, 95%CI = [0.55, 0.99] (Table 4, row a). We could not determine whether d' was equated (within ± 0.06 prime units) across easy and hard contexts because evidence was insensitive (main effect of context: $BF_{MID(\pm 0.06)} = 0.73$, Table 4, row b, Fig. 2G). Evidence for an interaction term within ± 0.06 d' units was also insensitive (difficulty x context: $BF_{H(0.06)} = 0.72$, Table 4, row c).

Of main importance here is that the pattern of objective performance here is similar to Experiment 2, indicating that this dual-task paradigm did not qualitatively change participants' objective decision-making.

Self-efficacy effects transfer to a secondary task

Finally, we tested our key hypothesis that motion discrimination confidence would no longer be affected by context in this dual-task paradigm. We hypothesised that instead, performing the motion task in an easy context would lead to better performance on the secondary disc discrimination task. In other words, we hypothesised that effects of difficulty context would *transfer* from primary motion confidence to secondary disc performance. Note that the difficulty of the disc task was constant over all trials: in the easy context, only the motion discrimination task was more likely to be easy, and in the hard context, again only the motion discrimination task was more likely to be hard..

To test for transfer we utilised three criteria: first, there should be no effect of difficulty context on confidence; second, the effect of difficulty context on confidence should be smaller in Experiment 3 than in Experiment 2; third, there should be a main effect of difficulty context on disc d' .

As well as hypothesising that the main effect of context would transfer, we hypothesised that the difficulty x context interaction effect would transfer as well. Therefore, the above 3 criteria were also applied to the interaction term.

First, we considered the main effect of context on confidence. This effect was not present in Experiment 1, and in Experiment 2 was such that confidence was higher in the easy context than hard context condition.

Table 4. Pre-registered results from Experiment 3.

	Bayesian				Frequentist		
	M _{diff}	SE _{diff}	Prior	BF	RR	F	p
d' (RDK)							
a. Objective difficulty	0.770	0.107	-	-	-	51.863	< .001
b. Difficulty context	-0.107	0.101	MID(±0.06)	0.73	(-∞, 0.15]	1.083	.348
c. Self-efficacy x subjective difficulty	0.223	0.178	H(0,0.06)	0.72	(-∞,0.22]	0.930	.402
Confidence*							
d. Accuracy	6.981	1.025	-	-	-	45.209	< .001
e. Objective difficulty	0.283	0.338	U(0,0.09)	0.10	[0.03, ∞)	0.239	.630
f. Difficulty context	0.304	1.117	H(0,0.26)	0.09	[0.07, ∞)	0.833	.442
g. Accuracy x objective difficulty	2.848	0.677	U(0,0.09)	851.68	[0.005, ∞)	16.225	< .001
h. Accuracy x context	-1.238	1.369	U(0,0.09)	0.78	(-∞,0.23]	2.649	.083
i. Objective difficulty x context	0.017	0.680	H(0,0.02)	0.57	(-∞,0.04]	0.425	.657
j. Accuracy x context x objective difficulty	0.016	1.763	U(0,0.04)	0.85	(-∞,0.12]	0.238	.789
m-ratio							
k. Subjective difficulty	0.009	0.052	U(0,0.50)	0.15	[0.22, ∞)	0.028	.869
l. Difficulty context	0.034	0.078	U(0,0.50)	0.29	[0.43, ∞]	4.068	.024
m. Subjective difficulty x context	-0.076	0.154	U(0,0.03)	1.05	(-∞, 0.91]	0.156	.856
d' (disc)							
n. RDK accuracy	0.380	0.113	H(0,0.9)	14.96	[0.04, ∞)	11.416	.003
o. Objective RDK difficulty	-0.008	0.053	H(0,0.03)	0.85	(-∞,0.14]	0.022	.884
p. Difficulty context	0.056	0.083	H(0,0.04)	1.15	(-∞,0.46]	1.333	.275
q. RDK accuracy x objective RDK difficulty	0.295	0.119	-	-	-	6.177	.021
r. RDK accuracy x context	-0.031	0.123	-	-	-	0.875	.425
Objective RDK difficulty x context	0.501	0.163	H(0,0.75)	38.40	[0.06, ∞)	4.874	.012
RDK accuracy x context x objective RDK difficulty	-0.069	0.280	-	-	-	0.463	.633

M_{diff} = mean difference; SE_{diff} = standard error of difference; BF = Bayes Factor; RR = robustness region; MID = minimal interesting difference; U = uniform prior; H = half-normal prior

*BFs and ANOVAs are calculated on log-confidence, and priors and RRs are given in log-confidence units. For interpretability, descriptives (M_{diff} and SE_{diff}) are given in raw (untransformed) units

In dual-task Experiment 3, confidence in the motion judgement was unaffected by context (no main effect of context: $BF_{H(0.26)} = 0.09$, Table 4, row f). This indicates that the first criterion for transfer was met. However, the second two were not: we found only anecdotal evidence for a smaller effect of context in Experiment 3 than 2 (Bayesian independent t-test, $BF_{U(0.2.19)} = 2.03$, $RR = (-\infty, \infty)$), and insensitive evidence for an effect of context on disc d' (main effect of context: $BF_{H(0.04)} = 1.15$, Table 4, row p). Therefore, we cannot conclude that the main effect of difficulty context transferred from RDK confidence to the secondary disc discrimination task.

Next we whether the self-efficacy x difficulty interaction - the key effect that was present in both Experiments 1 and 2 - transferred. With regards to transfer criterion 1, we found only anecdotal evidence against an interaction on RDK (difficulty x context: $BF_{H(0,02)} = 0.57$, Table 4, row i). Transfer criterion 2 was met: the magnitude of the interaction here in Experiment 3 was indeed smaller than in Experiment 2 (Bayesian independent t-test, $BF_{U(0,1.53)} = 3.30$, $RR = [1.25, 3.51]$). Finally, criterion 3 was met as well: difficulty and context interacted on the secondary disc task (disc d'). Here, disc discrimination performance affected by RDK difficulty more in the easy than the hard contexts (difficulty x context: $BF_{H(0,0.75)} = 38.40$, Table 4 row r, Fig. 2J). Crucially, this means that difficulty context for the primary task affected the *secondary* disc task performance, despite context only predicting the difficulty of the primary RDK stimulus.

Note that we did not measure confidence in the secondary task so we cannot ascertain whether effects of primary task difficulty context would have transferred to secondary task confidence. However, taken together, these results show that in dual-task Experiment 3, the difficulty x context interaction transferred from the primary RDK task to the secondary disc discrimination task.

Control analyses (exploratory)

The transfer of context effects is not driven by the relative ease of the disc task

One might worry that the influence of difficulty context on disc discrimination performance arose because – despite instructions – participants prioritised this secondary task. On this interpretation, effects of self-efficacy would not have *transferred* to the secondary task, but rather have applied *only to the primary* (disc) task.

To test this possibility, we require a means to infer which task participants treated as primary. We used the difference in performance ($d'_{RDK} - d'_{disc}$) as a proxy, reasoning that participants showing substantially greater performance on the RDK than disc task were prioritising the RDK task, and vice versa. At the group level, the RDK task out-performed disc discrimination, $t(21) = 4.16$, $p < .001$, 95%CI [0.33, 0.99] and this held for 20/22 subjects, consistent with the notion that the motion discrimination task was indeed participants' priority, as instructed. Furthermore, the difference in performance between the two tasks – the proxy for RDK priority - predicted neither the *presence* of difficulty x context interaction under d'_{disc} $\rho(20) = -0.16$, $p = .479$, nor the *absence* of that interaction under confidence, $\rho(20) = 0.12$, $p = .600$.

In summary, we found no evidence that the transfer of self-efficacy effects to the secondary disc task were driven by differences in how subjects allocated priority between the tasks. This

supports our interpretation of Experiment 3: that high self-efficacy on one task boosts objective performance on another, irrelevant task.

Smaller effects of self-efficacy on difficult trials are not caused by failure to learn the contingency

We have interpreted our confidence results in Experiments 2 and 3 as showing that effects of task difficulty on confidence are *larger in easy contexts*. However, on an alternative interpretation, difficulty context has a *smaller effect when trials are hard*. This could arise if hard trials (with high sensory uncertainty) impaired participants' ability to learn the task statistics (i.e. the difficulty context).

To test whether participants were able to learn the task statistics (difficulty context) on hard trials, we tested the extent to which confidence on trial t was influenced by trial history, specifically, the difficulty of trials $t-1$, $t-2$ and $t-3$. The idea here is that if participants do learn the task statistics then they should adapt their behaviour according to what they recently saw: more exposure to difficult trials should decrease confidence, just as more exposure to easy trials should increase confidence. This “serial dependence” style of analysis was run separately for each of Experiments 1, 2 and 3, and tested whether participants' confidence was higher/lower following an increasing number of successive easy/hard trials. If high sensory uncertainty impaired participants ability to learn, then we should not see an effect of trial history on confidence for hard trials.

To control for accuracy effects on confidence, confidence ratings were normalised within-subjects: for each participant we z-scored their confidence ratings separately for correct versus incorrect trials. Then, we split trials according to whether the current trial was easy or hard. Finally, we created 4 levels of trial history: 0, 1, 2 and 3, corresponding to whether the difficulty of the current trial differed from the previous trial (0), or was the same as the previous 1, 2 or 3 trials. Results remained qualitatively the same when using a longer or shorter trial history.

As shown in figure 3 (top), for single-task Experiments 1 and 2 normalised confidence appeared to increase with the number of successive easy trials as expected. Furthermore, normalised confidence also seemed to decrease with successive difficult trials, suggesting that participants were able to learn the context when presented with hard trials. This was clearly not the case in Experiment 3, consistent with the lack of an effect of self-efficacy on confidence in the dual-task, “high load” paradigm. We tested this statistically as follows: In case participants varied somewhat in how they weighted trial history, we tested whether participants' normalised confidence ratings monotonically increased (or decreased) with trial history (0, 1, 2, or 3 previous trials of the same

Figure 3. Effects of difficulty history on confidence. **Top.** Group-level mean normalised confidence as a function of trial history for difficult (red) and easy (green) trials, separately for Experiments 1-3. **Bottom.** Fisher-z transformed Spearman's correlations between trial history and normalised confidence, separately for Experiments 1-3. Light red/green dots are individual participants' correlations for hard/easy trials, and solid dots with outline represent the group mean $\pm$ SEM Fisher-z transformed correlation.

difficulty). Fisher-z transformed spearman correlations were then compared to zero, separately for easy and hard trials (figure 3, bottom).

For Experiment 1, there was no significant monotonic relationship between lag and confidence for hard trials, $t(18) = -1.45$, $p = .164$, with only 11/19 participants showing a negative association. For easy trials 14/19 participants had a positive association and at the group level these associations were greater than zero, $t(18) = 2.95$, $p = .009$. Accordingly, participants were indeed less able to learn the association on hard trials in Experiment 1.

However, lack of learning cannot explain the results of Experiment 2. Here, 15/21 could learn the association on hard trials and 19/21 on easy trials, and there was a significant difference in correlation from zero in both cases (hard: $t(20) = -2.95$, $p = .008$, easy: $t(20) = 4.93$, $p < .001$). Taken together, these results suggest that participants were able to learn the difficulty context when on hard trials.

In summary, these control analyses support our interpretation of Experiments 1 and 2: that confidence is more sensitive to task difficulty in easy contexts.

Discussion

Across three pre-registered experiments we examined how perceptual decisions and accompanying subjective confidence judgements are influenced by both trial-by-trial task difficulty, and the difficulty context those trials are embedded in. Difficulty context served as our operationalisation of self-efficacy: repeated interactions with easy (or hard) trials should promote statistical beliefs that subsequent trials will be easy (or hard), and thus beliefs that one is performing well (high self-efficacy), or poorly (low self-efficacy). Our results showed that confidence was more sensitive to task difficulty in easy contexts (high self-efficacy), despite leaving objective performance unaffected. This occurred both when performance was matched across difficulty levels (Experiment 1, “subjective difficulty” manipulation), and when easy trials were associated with better objective performance (Experiment 2, “objective difficulty” manipulation). Results from Experiments 1 and 2 together suggest that self-efficacy weights trial difficulty in the construction of confidence, such that trial difficulty matters more for confidence judgements when participants believe they are performing relatively well. These results are consistent with recent work showing that self-efficacy ratings are affected by the history of confidence reports (Rouault et al., 2019; Rouault & Fleming, 2020).

A third experiment tested whether self-efficacy increases confidence while leaving performance unaffected by making tasks feel less effortful. This was tested indirectly by appealing to ‘load theory’ (Lavie et al., 2004), which predicts that as the load of a primary task increases, performance on a second “distractor” task should decrease. Results largely supported our hypothesis: when performing two tasks simultaneously, and more so on easy trials, performance on a secondary disc discrimination task (distractor task) was higher in easy context (low load) than in a hard context (high load) blocks, consistent with what one would expect if difficulty context is a manipulation of load. In fact, in this experiment the interaction between trial difficulty and difficulty context no longer influenced confidence on the primary task, but transferred to influenced performance on the secondary task despite the context manipulation being entirely irrelevant to that secondary task.

While Experiment 3 was motivated by load theory, similar predictions are made by Bandura’s theory of self-efficacy and motivation (Bandura, 1977, 1982). Bandura proposed that perceived ‘effort’ would be a function of difficulty and self-efficacy, such that perseverance on a difficult task would require the individual to believe themselves capable of completing it; if self-efficacy is low and the individual therefore believes they are incapable, they will give up. This theory has received most of its support from high-level tasks such as reading, sports and mathematics, often in developmental settings (Bandura, 1982; Schunk, 1995; Zimmerman, Bandura, & Martinez-Pons, 1992). Here, extending this work, we show effects consistent with Bandura’s predictions on a simple perceptual task.

We have interpreted our results as showing that self-efficacy amplifies the effect of trial difficulty on confidence. One alternative interpretation is that difficulty context increases confidence, but only on easy trials. This could occur, for example, if participants were unable to learn the task statistics under high sensory uncertainty. However, a series of control analyses weighed against this alternative interpretation by showing that people *were* able to adapt their behaviour (give different confidence ratings) following repeated hard trials, albeit to a smaller degree than on easy trials. Another alternative explanation for our results is that confidence is constructed in a qualitatively different fashion for easy and hard trials. However, this is not a parsimonious explanation, nor is it compatible with normative models of decision confidence. Accordingly, our results remain most appropriately interpreted as indicating that confidence is more sensitive to task difficulty when self-efficacy is high.

This interpretation implies that confidence is proportional to the multiplication of sensory information (a function of trial difficulty) by a prior on being correct (difficulty context/self-efficacy). This idea - that confidence integrates sensory information with a prior on being correct - is consistent with (but not evidence for) Bayesian accounts of cognition. These models, which are closely related to “predictive processing” accounts (Clark, 2013; Hohwy, 2013), are becoming increasingly popular in the decision-making literature. Broadly, they propose that decisions arise from the integration of *priors* (expectations) and *likelihoods* (sensory information) to form a *posterior* belief (over the cause of sensory data). A key element of predictive processing (and more broadly Bayesian) accounts of perception and cognition is that the influence that each component has on inference is a function of their respective *precisions* (their inverse variance, or reliability). High-precision (reliable) priors will influence posterior beliefs to a greater extent than low-precision (unreliable) priors. It is important to recognise here that the precisions themselves have to be inferred, by integrating sensory data with prior beliefs about precisions (Friston et al., 2015; Parr, Benrimoh, Vincent, & Friston, 2018). These priors on precisions are called *precision expectations* and are also referred to as hyper-priors within a Bayesian framework. On this view, contextual factors (such as high or low signal-to-noise environments – high or low difficulty contexts, in our study) - shift precision expectations, which in turn give rise to behaviour consistent with the deployment of endogenous attention – since ‘paying attention’ can also be considered within this framework as transiently increasing the precision-weighting of sensory signals (Feldman & Friston, 2010).

In the present setting, the contextual manipulation of self-efficacy - where participants were exposed to high or low signal-to-noise environments - can be re-described as a manipulation of expected precision. Trials in high self-efficacy contexts would correspond to trials where expected precision was high, and vice versa for low self-efficacy contexts. Formulating self-

efficacy in this way would also predict an interaction between trial difficulty and self-efficacy: the sensory information (manipulated by changing trial difficulty through manipulating the stimulus) should interact multiplicatively with expected precision (the self-efficacy manipulation) because of the multiplicative nature of Bayesian integration. To the extent that confidence is proportional to the posterior mean (Hangya, Sanders, & Kepecs, 2016; Meyniel, Sigman, & Mainen, 2015; Sherman, Seth, & Kanai, 2016), such an effect should carry forward into confidence judgements.

Consistent with previous work (Konishi, Compain, Berberian, Sackur, & de Gardelle, 2020; Maniscalco & Lau, 2015), participants' ability to discriminate between their correct and incorrect judgements (metacognition) was similar across single- and dual-task contexts and across easy and hard trials (Allen et al., 2016; Spence et al., 2015). In (single-task) Experiments 1 and 2, metacognition was also unaffected by self-efficacy; however, this was not the case in (dual-task) Experiment 3. Here, metacognition was better when participants could predict task difficulty (i.e., in both high and low difficulty contexts), than when difficulty was unpredictable (i.e., in the control condition with 50% hard and easy trials). This is consistent with data showing that people can predict when their performance decreases due to multitasking, while being unable to predict the magnitude of those decreases (Finley, Benjamin, & Mccarley, 2014).

To conclude, the three experiments presented here extend studies of perceptual decision-making beyond trial-by-trial measures of accuracy and confidence, to encompass psychologically meaningful higher-order concepts such as self-efficacy and cognitive load. They therefore bridge social and cognitive models of how self-efficacy influences decision-making, and motivate future work on question of how to connect these models at the computational level.

References

- Alilović, J., Slagter, H. A., & van Gaal, S. (2020). Subjective visibility report is facilitated by conscious predictions only. *BioRxiv*, 2020.07.08.193037. <https://doi.org/10.1101/2020.07.08.193037>
- Allen, M., Frank, D., Schwarzkopf, D. S., Fardo, F., Winston, J. S., Hauser, T. U., & Rees, G. (2016). Unexpected arousal modulates the influence of sensory noise on confidence. *ELife*, 5. <https://doi.org/10.7554/eLife.18103>
- Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*. <https://doi.org/10.1037/0033-295X.84.2.191>
- Bandura, A. (1982). Self-efficacy mechanism in human agency. *American Psychologist*. <https://doi.org/10.1037/0003-066X.37.2.122>
- Charles, L., Van Opstal, F., Marti, S., & Dehaene, S. (2013). Distinct brain mechanisms for conscious versus subliminal error detection. *NeuroImage*, 73, 80–94. <https://doi.org/10.1016/j.neuroimage.2013.01.054>
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, 36(3), 181–204. <https://doi.org/10.3389/fpsyg.2013.00270>
- Denison, R. N., Adler, W. T., Carrasco, M., & Ma, W. J. (2018). Humans incorporate attention-dependent uncertainty into perceptual decisions and confidence. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.1717720115>
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*.

- <https://doi.org/10.3389/fpsyg.2014.00781>
- Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4(December), 215. <https://doi.org/10.3389/fnhum.2010.00215>
- Fetsch, C. R., Kiani, R., Newsome, W. T., & Shadlen, M. N. (2014). Effects of Cortical Microstimulation on Confidence in a Perceptual Decision. *Neuron*, 83(4), 797–804. <https://doi.org/10.1016/j.neuron.2014.07.011>
- Finley, J. R., Benjamin, A. S., & Mccarley, J. S. (2014). Metacognition of multi-tasking: how well do we predict the costs of divided attention?, 20(2), 158–165. <https://doi.org/10.1037/xap0000010>.Metacognition
- Fleming, S. M., Maniscalco, B., Ko, Y., Amendi, N., Ro, T., & Lau, H. (2015). Action-specific disruption of perceptual confidence. *Psychological Science*, 26(1), 89–98. <https://doi.org/10.1177/0956797614557697>
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., FitzGerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, (April), 150217111908007. <https://doi.org/10.1080/17588928.2015.1020053>
- Frith, C. D. (2012). The role of metacognition in human social interactions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1599), 2213–2223. <https://doi.org/10.1098/rstb.2012.0123>
- Goodman, L. A. (1970). The multivariate analysis of qualitative data: Interactions among multiple classifications. *Journal of the American Statistical Association*. <https://doi.org/10.1080/01621459.1970.10481076>
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics* (Vol. 1). New York: Wiley.
- Grimaldi, P., Lau, H., & Basso, M. A. (2015). There are things that we know that we know, and there are things that we do not know we do not know: Confidence in decision-making. *Neuroscience and Biobehavioral Reviews*, 55, 88–97. <https://doi.org/10.1016/j.neubiorev.2015.04.006>
- Guggenmos, M., Wilbertz, G., Hebart, M. N., & Sterzer, P. (2016). Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *ELife*, 5, 1–19. <https://doi.org/10.7554/eLife.13388>
- Hangya, B., Sanders, J. I., & Kepecs, A. (2016). A mathematical framework for statistical decision confidence. *BioRxiv*. Retrieved from <http://biorxiv.org/content/early/2016/01/01/017400.abstract>
- Hohwy, J. (2013). *The Predictive Mind*. Oxford: OUP.
- Jeffreys, H. (1961). Theory of probability (Third ed.). In *University Press, Oxford, England*.
- Koizumi, A., Maniscalco, B., & Lau, H. (2015). Does perceptual confidence facilitate cognitive control? *Attention, Perception, and Psychophysics*. <https://doi.org/10.3758/s13414-015-0843-3>
- Konishi, M., Compain, C., Berberian, B., Sackur, J., & de Gardelle, V. (2020). Resilience of perceptual metacognition in a dual-task paradigm. *Psychonomic Bulletin & Review*, 1–10.
- Kurtz, P., Shapcott, K. A., Kaiser, J., Schmiedt, J. T., & Schmid, M. C. (2017). The Influence of Endogenous and Exogenous Spatial Attention on Decision Confidence. *Scientific Reports*. <https://doi.org/10.1038/s41598-017-06715-w>
- Lavie, N., Hirst, A., Fockert, J. W. De, & Essi Viding. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology. General*, 133(3), 339. <https://doi.org/10.1037/0096-3445.133.3.339>
- Maniscalco, B., & Lau, H. (2015). Manipulation of working memory contents selectively impairs metacognitive sensitivity in a concurrent visual discrimination task. *Neuroscience of Consciousness*. <https://doi.org/10.1093/nc/niv002>
- Meyniel, F., Sigman, M., & Mainen, Z. F. (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron*, 88(1), 78–92. <https://doi.org/10.1016/j.neuron.2015.09.039>
- Michael, E., de Gardelle, V., & Summerfield, C. (2014). Priming by the variability of visual information. *Proceedings of the National Academy of Sciences of the United States of America*, 111(21), 7873–7878. <https://doi.org/10.1073/pnas.1308674111>
- Murphy, P. R., Robertson, I. H., Harty, S., & O'Connell, R. G. (2015). Neural evidence accumulation persists after choice to inform metacognitive judgments. *ELife*, 4(DECEMBER2015). <https://doi.org/10.7554/eLife.11946>
- Parr, T., Benrimoh, D. A., Vincent, P., & Friston, K. J. (2018). Precision and False Perceptual Inference . *Frontiers in Integrative Neuroscience* . Retrieved from <https://www.frontiersin.org/article/10.3389/fnint.2018.00039>
- Peters, M. A. K., Thesen, T., Ko, Y. D., Maniscalco, B., Carlson, C., Davidson, M., ... Lau, H. (2017). Perceptual confidence neglects decision-incongruent evidence in the brain. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-017-0139>

- Rahnev, D., Koizumi, A., Mccurdy, L. Y., Esposito, M. D., & Lau, H. (2015). Confidence Leak in Perceptual Decision Making. *Psychological Science*, 26(11), 1664–1680. <https://doi.org/10.1177/0956797615595037>
- Rouault, M., Dayan, P., & Fleming, S. M. (2019). Forming global estimates of self-performance from local confidence. *Nature Communications*. <https://doi.org/10.1038/s41467-019-09075-3>
- Rouault, M., & Fleming, S. M. (2020). Formation of global self-beliefs in the human brain. *Proceedings of the National Academy of Sciences*, 117(44), 27268–27276.
- Schröger, E., Marzecová, A., & SanMiguel, I. (2015). Attention and prediction in human audition: a lesson from cognitive psychophysiology. *The Journal of Neuroscience*, 41(5), 641–664. <https://doi.org/10.1111/ejn.12816>
- Schunk, D. H. (1995). Self-efficacy, motivation, and performance. *Journal of Applied Sport Psychology*, 7(2), 112–137.
- Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., & Frith, C. D. (2014). Supra-personal cognitive control and metacognition. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2014.01.006>
- Sherman, M.T., Seth, A. K., & Kanai, R. (2016). Predictions shape confidence in right inferior frontal gyrus. *Journal of Neuroscience*, 36(40). <https://doi.org/10.1523/JNEUROSCI.1092-16.2016>
- Sherman, Maxine T., Seth, A. K., Barrett, A. B., & Kanai, R. (2015). Prior expectations facilitate metacognition for perceptual decision. *Consciousness and Cognition*, 35, 53–65. <https://doi.org/10.1016/j.concog.2015.04.015>
- Siedlecka, M., Paulewicz, B., & Wierzchoń, M. (2016). But I Was So Sure! Metacognitive Judgments Are Less Accurate Given Prospectively than Retrospectively. *Frontiers in Psychology*, 7(February), 1–8. <https://doi.org/10.3389/fpsyg.2016.00218>
- Spence, M. L., Dux, P. E., Arnold, D. H., Spence, M. L., Dux, P. E., & Arnold, D. H. (2015). Computations Underlying Confidence in Visual Perception. *Journal of Experimental Psychology : Human Perception and Performance*, 41(6).
- Watson, A. B., & Pelli, D. G. (1983). Quest: A Bayesian adaptive psychometric method. *Perception & Psychophysics*. <https://doi.org/10.3758/BF03202828>
- Zimmerman, B. J., Bandura, A., & Martinez-Pons, M. (1992). Self-motivation for academic attainment: The role of self-efficacy beliefs and personal goal setting. *American Educational Research Journal*, 29(3), 663–676.
- Zizlsperger, L., Kümmel, F., & Haarmeier, T. (2016). Metacognitive Confidence Increases with, but Does Not Determine, Visual Perceptual Learning. *Plos One*, 11(3), e0151218. <https://doi.org/10.1371/journal.pone.0151218>
- Zizlsperger, L., Sauvigny, T., & Haarmeier, T. (2012). Selective attention increases choice certainty in human decision making. *PLoS ONE*, 7(7). <https://doi.org/10.1371/journal.pone.0041136>
- Zylberberg, A., Fetsch, C. R., & Shadlen, M. N. (2016). The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. *ELife*. <https://doi.org/10.7554/eLife.17688>
- Zylberberg, A., Roelfsema, P. R., & Sigman, M. (2014). Variance misperception explains illusions of confidence in simple perceptual decisions. *Consciousness and Cognition*, 27C, 246–253. <https://doi.org/10.1016/j.concog.2014.05.012>

Supplementary Materials

Contents:

- SM1. Justification of priors**
- SM2. Supplementary results**
- SM3. Supplementary Table S1**

SM1. Justification of priors

Experiment One:

Our manipulation of difficulty context was based upon results reported in Koizumi et al (2015), who found that higher response-congruent evidence (RCE, which they refer to as “positive evidence”) was associated with a mean 0.25 increase in confidence on a 5-point rating scale. This corresponds to a 6% change. In our paper we refer to the level of RCE as a subjective difficulty manipulation.

First, we set our test and prior for orientation performance d' . We predicted that our staircase would successfully equate performance across the subjectively easy and subjectively hard conditions. In other words, we predicted a difference of “basically zero”. To model this, we needed to set a minimal interesting difference below which an effect is categorised as zero, and above which an effect is categorised as not zero.

We set our minimal interesting difference for an effect of subjective difficulty on d' at $\pm 0.06 d'$ units, because 0.06 is equal to a 6% difference when d' is approximately 1, and 6% was the effect size found by Koizumi et al. This point prior was compared to a point null of 0 by means of a likelihood ratio test. The same prior was used to test whether d' was equated across levels of difficulty context. The logic here is that self-efficacy is manipulated via changing the likelihood of a high versus low RCE stimulus, and so we should expect a similar effect size to that arising from the subjective difficulty manipulation.

Again, following results from Koizumi et al (2015) we predicted that confidence would be 6% higher for the low versus high subjective difficulty stimulus and also, when the high versus low subjective difficulty stimulus is expected. Because we analysed log-confidence using a linear mixed model, the betas for each model term corresponds to percentage change in log confidence by that term. Therefore, our prior for both stimulus and context was modelled as a half-normal with $SD = 0.06$.

Finally, we had to determine our priors for analyses on m-ratio. For this outcome measure, we used an uninformed prior. We reasoned that, because m-ratio is theoretically bounded by 0 (no metacognitive sensitivity) and 1 (SDT-optimal metacognition), that maximum effect of both stimulus and self-efficacy would be an increase of 50%, or 0.5 m-ratio units. The prior for both of these main effects was therefore a uniform distribution from 0 to 0.5.

Experiment Two:

For the main effect of difficulty context on d' , we used the same minimal interesting difference of 0.06 as in Experiment 1, and used a likelihood ratio test to determine whether effects were indeed smaller than this value. Because we found sensitive evidence for an interaction term (or rather, evidence against no interaction) in Experiment 1, we set our prior here as a half-normal with mean 0 and SD equal to the difference found in Experiment 1.

Experiment 1 revealed no sensitive evidence for tests under m-ratio, so in Experiment 2 we used the same uniform priors (uniform distributions from 0 to 0.5) for the main effects.

For analyses on log-confidence, we reasoned that the effect of trial difficulty would probably not be larger than the effect of accuracy, so the prior here was a uniform with a lower bound of 0 and an upper bound of $\text{confidence}_{\text{correct}} - \text{confidence}_{\text{incorrect}}$. We used the same prior of a half-normal centred on zero, with SD reflecting a 6% change by context. The SD was calculated as $0.06 \times \text{mean log-confidence in the difficult context condition}$. The prior for the trial difficulty $\times$ context interaction was based on results from Experiment 1: a half-normal centred on zero with SD equal to the interaction magnitude from that experiment.

Experiment Three:

For analyses on RDK d' , confidence and m-ratio priors for Experiment 3 were determined as follows: if evidence for an effect was found in Experiment 2, the prior for that effect in Experiment 3 was modelled as a half-normal with mean 0 and SD equal to the difference found in Experiment 2. If evidence against an effect was found in Experiment 2, or if the Bayes Factor was insensitive, we used the same prior as that used in Experiment 2.

For analyses on disc discrimination d' , we did the following:

For the main effects of RDK difficulty, RDK accuracy and difficulty context, and for the difficulty $\times$ context interaction, we calculated the corresponding percentage change in log-confidence seen in Experiment 2. We predicted that, whatever the effect was on confidence in Experiment 2, that same effect would be seen under disc d' in Experiment 3. Accordingly, priors for disc d'

in Experiment 3 were modelled as half-normals, with SD scaled according to those percentages. To illustrate: Experiment 2 found a 6% increase in log-confidence in easy versus difficult contexts. The prior for disc d' in Experiment 3 was therefore a half-normal with mean 0, and $SD = 0.06 \times \text{mean disc } d'$ for hard trials.

We predicted that the main effect of context on confidence would be smaller in Experiment 3 than Experiment 2. The mean difference in confidence between easy and hard contexts for Experiment 2 was 2.19, therefore the prior for this independent Bayesian t-test was set as a uniform distribution over 0 to 2.19. This represents a belief that the effect in Experiment 3 could take any value from 0 (no effect at all), to the effect found in Experiment 2.

The same approach was taken for the difficulty $\times$ context interaction. The mean interaction effect was 1.53 in Experiment 2, therefore the prior was modelled as a uniform distribution over 0 to 1.53.

SM2. Supplementary results: Effects of difficulty and self-efficacy on metacognitive sensitivity

In Experiment 1, we examined whether subjective difficulty or difficulty context influenced metacognitive sensitivity m -ratio on our Gabor orientation discrimination task. Results showed evidence *against* an effect of subjective difficulty, $M_{\text{diff}} = -0.02 \pm 0.07$, 95% CI = [-0.16, 0.12], $BF_{U(0,0.5)} = 0.22$, $RR = [0.32, \infty)$, meaning that metacognitive sensitivity was indistinguishable across subjectively easy and hard gratings. We did not find sensitive evidence for or against an effect of difficulty context, $M_{\text{diff}} = -0.09 \pm 0.07$, 95% CI = [-0.24, 0.07], $BF_{U(0,0.5)} = 0.64$, $RR = (-\infty, 0.99]$, or for a difficulty by context interaction, $M_{\text{diff}} = 0.17 \pm 0.11$, 95% CI = [-0.07, 0.41], $BF_{U(0,0.09)} = 1.67$, $RR = (-\infty, 2.48]$. This means we cannot determine whether context exerted a main or interactive effect on metacognition.

Experiment 2 manipulated objective difficulty and difficulty context in an RDK motion discrimination task. Results were similar to Experiment 1: as previously, we found evidence for no effect of difficulty on m -ratio, $M_{\text{diff}} = -0.03 \pm 0.06$, 95% CI = [-0.16, 0.09], $BF_{U(0,0.5)} = 0.25$, $RR = [0.37, \infty)$. Evidence for an effect of context was insensitive, trending towards no effect, $M_{\text{diff}} = 0.06 \pm 0.08$, 95% CI = [-0.11, 0.23], $BF_{U(0,0.5)} = 0.40$, $RR = (-\infty, 0.61]$. We could not determine whether these factors interacted, $M_{\text{diff}} = -0.02 \pm 0.13$, 95% CI = [-0.28, 0.25], $BF_{U(0,0.06)} = 0.99$, $RR = (-\infty, 0.54]$.

Experiment 3 was the same as Experiment 2, except for the fact that participants performed a concurrent disc discrimination task. In this dual-task paradigm, difficulty context effected

metacognitive sensitivity: an exploratory Bayesian t-test found evidence for a U-shaped relationship with self-efficacy, such that participants were better able to discriminate correct from incorrect choices when difficulty was predictable, $M_{diff} = 0.15 \pm 0.05$, $BF = 82.40$ (Figure 2I). The pre-registered linear trend was not present ($BF = 0.38$), and the interaction between difficulty and self-efficacy was insensitive ($BF = 1.12$). Again, we found evidence for no effect of difficulty on m-ratio, $BF = 0.25$, showing that metacognition is unaffected by task difficulty in both single- and in dual-task settings.

SM3. Supplementary Tables S1

Table S1. Confirmatory results from Experiment 1.

			<u>Bayesian</u>			<u>Frequentist</u>	
	M_{diff}	SE_{diff}	Prior	BF	RR	F	p
d'							
Context	-0.021	0.063	MID(0.06)	0.002	[0.056, ∞)	0.107	0.747
Subjective difficulty	0.053	0.151	MID(0.06)	0.497	($-\infty$, 0.146]	0.124	0.727
Context x subjective difficulty	-0.161	0.074	H(0,0.053)	2.178	($-\infty$, 0.081]	4.798	0.039
m-ratio							
Context	-0.036	0.070	U(0,0.5)	0.278	[0.407, ∞)	0.258	0.616
Subjective difficulty	-0.066	0.062	U(0,0.5)	0.471	($-\infty$, 0.739]	1.145	0.296
Self-efficacy x subjective difficulty	0.137	0.100	U(0,0.066)	1.494	($-\infty$, 1.840]	1.882	0.183
Log-confidence							
	β	SE	Prior	BF	RR	χ^2	p
Accuracy	0.064	0.007	-	-	-	84.788	< 0.001
Context	0.003	0.003	H(0,0.06)	0.140	[0.023, ∞)	0.975	0.324
Subjective difficulty	0.032	0.007	H(0,0.06)	6951	[0.002, ∞)	25.008	< 0.001
Accuracy x Context	-0.001	0.002	U(0,0.064)	0.061	[0.012, ∞)	0.418	0.518
Accuracy x subjective difficulty	0.004	0.003	U(0,0.064)	0.259	[0.049, ∞)	1.303	0.254
Context x subjective difficulty	-0.003	0.002	U(0,0.032)	0.450	($-\infty$, 0.044)	2.401	0.121
Accuracy x Context x subjective difficulty	-0.001	0.002	U(0,0.004)	0.887	($-\infty$, 0.012]	0.336	0.562

Sensitive Bayes factors are presented in bold type. M_{diff} = mean difference; SE_{diff} = standard error of difference; BF = Bayes Factor; RR = robustness region. MID = minimal interesting difference; U = uniform prior; H = half-normal prior