Stage 2 manuscript accepted at Cortex, 12.12.2022

Feasibility of Unconscious Instrumental Conditioning: A Registered Replication

Lina I. Skora ^{a,c,d}, James J.A. Livermore ^{a,e}, Zoltan Dienes ^{a,c}, Anil K. Seth ^{b,c,f}, & Ryan B. Scott ^{a,c}

^a School of Psychology, University of Sussex, Brighton, UK, BN1 9RH

^b School of Informatics, University of Sussex, Brighton, UK, BN1 9RH

^c Sackler Centre for Consciousness Science, University of Sussex, Brighton, UK, BN1 9RH

^d Biological Psychology of Decision Making, Department of Experimental Psychology, Heinrich-Heine-Universität, 40225 Düsseldorf, Germany

^e Donders Institute for Brain, Cognition and Behaviour, Radboud University, 6525 HR Nijmegen, Netherlands

^f Brain, Mind, and Consciousness Program, Canadian Institute for Advanced Research (CIFAR), Toronto, Ontario, Canada

Declarations of interest: none reported.

Corresponding author:

Lina Skora, School of Psychology, Pevensey 1 Building, University of Sussex, Falmer, BN1 9QH.

Email: L.Skora@sussex.ac.uk

Feasibility of Unconscious Instrumental Conditioning: A Registered Replication**Abstract**

The extent to which high-level, complex functions can proceed unconsciously has been a topic of considerable debate. While unconscious processing has been demonstrated for a range of low-level processes, from feature integration to simple forms of conditioning and learning, theoretical contributions suggest that increasing complexity requires conscious access. Here, we focus our attention on instrumental conditioning, which has been previously shown to proceed without stimulus awareness. Yet, instrumental conditioning also involves integrating information over a large temporal scale and distinct modalities in order to deploy selective action, constituting a process of substantial complexity. With this in mind, we revisit the question of feasibility of instrumental conditioning in the unconscious domain. Firstly, we address the theoretical and practical considerations relevant to unconscious learning in general. Secondly, we aim to replicate the first study to show instrumental conditioning in the absence of stimulus awareness (Pessiglione et al., 2008), following the original design and supplementing the original crucial analyses with a Bayesian approach (Experiment 1). We found that apparent unconscious learning took place when replicating the original methods directly and according to the tests of awareness used. However, we could not establish that the full sample was unaware in a separate awareness check. We therefore attempted to replicate the effect yet again with improved methods to address the issues related to sensitivity and immediacy (Experiment 2), including an individual threshold-setting task and a trial-by-trial awareness check permitting exclusion of individual aware trials. Here, we found evidence for absence of unconscious learning. This result provides evidence that instrumental conditioning did not occur without stimulus awareness in this paradigm, supporting the view that complex forms of learning may rely on conscious access. Our results provides support for the proposal that perceptual consciousness may be necessary for complex, flexible processes, especially where selective action and behavioural adaptation are required.

Keywords: unconscious processing, unconscious learning, instrumental conditioning, consciousness

1. Introduction

Ever since the earliest demonstration of subliminal perception (Peirce & Jastrow, 1886), the extent to which information can be processed in the brain without conscious awareness has been a widely studied question. Unconscious processing has been demonstrated for many low-level processes such as feature detection and integration (e.g. integrating colour, shape and texture of an object into one coherent percept; Blake & Fox, 1974), as well as simple forms of learning, for instance emotional (Olsson & Phelps, 2004), visuospatial (Rosenthal et al., 2010) or associative (Scott et al., 2018). However, the extent to which unconscious processing is possible for higher-level, more complex functions remains a topic of debate (Axelrod & Rees, 2014; Mudrik et al., 2014; Sterzer et al., 2014). One such example is learning the contingencies between stimuli and outcomes, especially in instrumental scenarios, where the agent must learn from multiple temporally separated events: the stimulus itself, their action, and its consequence. This kind of learning has apparently been shown to be feasible in the absence of stimulus awareness (Mastropasqua & Turatto, 2015; Pessiglione et al., 2008). However, following recent evidence to the contrary (Reber et al., 2018) and discussions about the minimal conditions for unconscious processing (e.g. Mudrik et al., 2014), as well as developments in methods used to assess conscious awareness (Dienes, 2015; Rothkirch & Hesselmann, 2017; Shanks, 2017), we revisit the finding that instrumental learning can occur unconsciously. Here, we attempt to replicate the original result of Pessiglione and colleagues (2008), leveraging the developments in the field of unconscious processing to apply a more robust statistical approach (Experiment 1), and a more rigorous methodology (Experiment 2).

While there is no clear agreed theory demarcating what conscious versus unconscious mental states may represent (Breitmeyer, 2015; Dupoux et al., 2008; Kouider & Dehaene, 2007; Seth & Bayne, 2022), there have been attempts at outlining conditions under which unconscious processing can take place (Mudrik et al., 2014). A number of theoretical contributions consider consciousness a necessary component for higher-level processing, including (but not limited to) semantic knowledge, complex visual processing, as well as problem solving and decision-making (Baars, 2002; Treisman, 2003). In those views, consciousness plays a role in enabling information to be integrated across distinct brain regions through long-range feed-back and feed-forward connections (Baars, 2002; Dehaene & Changeux, 2011; Dehaene, Sergent, & Changeux, 2003). In contrast, unconscious processing appears to be confined to separate areas, and does not result in a global spread of activity (Baars, Ramsøy, & Laureys, 2003; Dehaene et al., 2001; Melloni et al., 2007).

In support of this view, previous research in the fields of associative learning and priming suggests that low-level or short-range spatiotemporal (Lin & He, 2009; Van Den Bussche et al., 2009) and multisensory (Faivre et al., 2014; Scott et al., 2018) information integration can proceed without conscious

awareness of the stimuli (typically achieved with subliminal presentation methods such as masking or continuous flash suppression). Conversely, higher-level or longer-range spatial and temporal processing (e.g. for tasks requiring longer-term information maintenance or selective, flexible decision-making) should require conscious access (Dehaene & Naccache, 2001; Dehaene et al., 2003; Kouider & Dupoux, 2001). Previously reported instances of unconscious learning are in line with those assumptions. For example, classical conditioning can be achieved without awareness in delay scenarios (where stimuli to be integrated overlap temporally), but not in trace scenarios (where stimuli are temporally separated (R. E. Clark & Squire, 1998)), with similar results in other associative learning tasks (Knight et al., 2003; Raio et al., 2012; Seitz et al., 2009).

As such, the idea that instrumental learning can proceed with unconsciously perceived stimuli is an interesting case. In the first experiment to demonstrate it (Pessiglione et al., 2008), participants learned to adjust their behaviour, through Go/NoGo decisions, in line with subliminally presented rewarding and punishing cues, learning to approach the rewarding and avoid the punishing stimulus without ever consciously perceiving them (constituting an example of trace instrumental learning). In order to successfully learn when to act and when to refrain from acting, participants had to learn from two temporally separated events across the length of the trial (up to 4 seconds): the subliminal stimulus itself, and its consequences, presented supraliminally as monetary reinforcement. Such a form of learning thus involves a fairly complex process of integrating information over a large temporal scale and distinct modalities, necessary to process the visual input, deploy selective action in response to the predictive cue, and process the reinforcement. The task is then considerably more complex than the aforementioned classical conditioning or associative learning scenarios, where there are fewer events, often in closer temporal proximity. Assuming that subliminally presented cue is not capable of evoking large-scale activity to be integrated with subsequent processes, the case of unconscious instrumental conditioning might appear at odds with the theory and previous experimental evidence covered above. Yet, instrumental conditioning is also one of the earliest and most fundamental forms of adaptive behaviour, both phylogenetically and ontogenetically. As such, the extent to which it requires conscious access is a question of considerable theoretical value.

A key challenge in any research into unconscious influences on behaviour lies in reliably asserting that processing is genuinely unconscious (Newell & Shanks, 2013; Rebuschat, 2013; Timmermans & Cleeremans, 2015). Although it is frequent practice in this line of research to infer unconscious processing when a behavioural measure (e.g. conditioning, priming, etc.) is above chance, while a separate measure of awareness is non-significantly different from chance performance (e.g. a non-significant result in a discrimination task), this approach has been heavily criticised (Dienes, 2015; Vadillo, Konstantinidis, & Shanks, 2016). A non-significant result alone cannot disambiguate between no evidence for an effect (i.e.

insensitive data, e.g. due to the small sample size) and absence of an effect (i.e. support for the null hypothesis). As such, finding that performance on an awareness check does not significantly differ from chance is not enough to assert true absence of awareness – an assertion which must be fulfilled to enable any inferences about the effect of interest, such as presence of unconscious conditioning in the original Pessiglione and colleagues study (Dienes, 2015; Shanks, 2017). This fallacy can be rectified in two ways: 1) ensuring that the methods are relevant, and sufficiently sensitive (Berry & Dienes, 1993, p.38; Shanks & St. John, 1994), and 2) with use of statistical methods, most prominently the Bayes factor, which enables stronger inferences about whether a null result indicates support for the null (e.g. awareness absent) over the alternative hypothesis (e.g. awareness present), or whether the data are insensitive (Dienes, 2014, 2016; Sand & Nilsson, 2016).

With these considerations in mind, we revisit the suggestion that instrumental learning can proceed without stimulus awareness. Experiment 1 will attempt to replicate the effect found by Pessiglione and colleagues (2008), following the original design and supplementing the original analyses with a Bayesian approach geared to determine a genuine absence of awareness, at least as measured by their test of awareness (whether this measure is a justified measure of awareness is an issue we will return to). Should the replication be successful, Experiment 2 will attempt to replicate the effect once again, this time with improved methods, to address the methodological issues related to the criteria of sensitivity and relevance in the original study.

In order to test whether stimuli that produce a certain level of learning are subliminal, one needs to know how much conscious perception would be needed to produce that level of learning (Dienes, 2015). Thus, a pilot study was conducted in which stimuli were presented moderately above the objective threshold in order to determine a relationship between the level of awareness (given the test of awareness used by Pessiglione et al, 2008) and learning. Thus, first we ran a pilot study to norm the relationship between learning and required awareness levels when the learning is based on conscious perception.

2. Pilot: Relationship between level of awareness and learning above the objective threshold

The pilot study aimed to assess both perceptual discrimination accuracy when awareness is present in a same/different discrimination task (as ensured with supraliminal stimulus presentation), and the corresponding level of learning subsequently achieved in a Go/NoGo task with the same stimulus exposure duration. This will be assessed employing a methodology identical to that of the replication study, Experiment 1. The observed relationship between awareness and learning will be used to identify the rough appropriate effect size for Bayes Factor calculation in the corresponding task conducted without awareness. The pilot was pre-registered at <https://osf.io/rwnt7>.

2.1. Method

2.1.1. Participants

26 participants (3 males, $M_{AGE} = 22$, $SD_{AGE} = 4.3$) were recruited. Sample size was determined with G*power (Faul et al., 2007), using a Cohen's d of 0.7 (a large effect size is justified given the supraliminal nature of the stimuli and the simplicity of the task), with 95% power. One participant was excluded after reporting to have misunderstood the learning task during debrief, yielding a final sample of 25 participants.

2.1.2. Stimuli and Materials

The stimuli included 9 randomly selected characters from the Agathodaimon font presented in a white typeface on a black background, with a size of 70 x 70 pixels. For each participant, 3 were randomly assigned to the first perceptual discrimination task (PDT1; threshold-setting), 3 to PDT2 (awareness check), and 3 to the main leaning task (1 to be associated with rewarding, 1 with punishing, and 1 with neutral outcome). Two black-and-white visual noise masks of the same size as the stimuli were generated by scrambling one character image into 8.75 by 8.75 pixels squares. The same two masks were used for all participants in the same fashion (one preceding and one following the target stimulus). The outcome images were a circled £1 coin image for reward, a crossed-out £1 image for punishment, and a greyed-out coin for neutral.

The task was programmed using Matlab 2018b (MathWorks, 2018), running Psychophysics Toolbox (Brainard, 1997), and presented on a Samsung 2233RZ LCD monitor with a 120Hz refresh rate (following recommendations for precise visual presentation; Wang & Nikolić, 2011). Responses were collected with a standard keyboard.

2.1.3. Procedure

2.1.3.1. Perceptual discrimination task 1: Threshold finding

This task aimed to establish a cue display duration that permitted conscious discrimination at above-chance levels without reaching ceiling. Participants were seated at a 50cm distance from the screen (ensured with a chinrest). Each trial began with a fixation cross (500ms), followed by presentation of two cues (display duration starting at 600ms), both forward-backward masked (67ms), separated by a 3s interval indicated by a fixation cross, following the method of Pessiglione et al (2008). Following the displays, participants were asked to indicate whether the cues presented were the same or different, and judge their confidence in that decision (on a binary scale between "some confidence" and "total guess"). Both responses were made using the arrow keys. Cue display duration started at 600ms, and dropped by 50ms with every correct and confident discrimination. Once participants reached 100ms or indicated guess for the first time, the display duration was increased by one increment (+50ms), and proceeded to decrease

by smaller increments (8ms, corresponding to a single screen refresh duration on a 120Hz monitor). Once participants responded guess 6 consecutive times (irrespective of accuracy), the corresponding display duration was taken to be their threshold of conscious perception.

The display duration was then set to be 16ms greater than the identified threshold and participants required to continue to make the same 'same' or 'different' and confidence judgments for a minimum of one block of 10 further trials. If objective discrimination accuracy for those 10 trials was between 70% and 90% (above chance, indicating that participants can reliably discriminate the cues, but are not at ceiling), the task terminated and the duration was recorded as the display duration to be used in the main task. Note that confidence was discounted in this measure, and only objective accuracy was taken into account. If discrimination accuracy for these 10 trials was greater than 90% then the display duration was reduced by 8ms and the process repeated for another 10 trials until discrimination accuracy fell into the desired range (70-90%). Similarly, if discrimination accuracy for the 10 trials was below 70% the display duration was increased by 8ms and a further block of 10 trials completed until such time as the desired discrimination accuracy was achieved.

2.1.3.2. Main conditioning task.

In keeping with the original protocol (Pessiglione et al., 2008), participants were asked to choose between making a response by pressing a spacebar (Go), or refraining from a response (NoGo), to masked cues. In each block, one cue was paired with reward, one with punishment, and one with the neutral outcome. Hence, participants could choose either to take a "risky" action (where they might win £1, lose £1, or have a neutral outcome depending on the proceeding cue) or to refrain from acting and thus ensure a neutral outcome.

Each trial began with a fixation cross (500ms), followed by a forward mask (67ms), one of the target cues (determined supra-threshold display duration), and backward mask (67ms). Subsequently, a question mark appeared on the screen, indicating that the response could be made. Regardless of the response (Go or NoGo), the response window remained open for 3000ms, after which the choice made (Go! or No!) was displayed (500ms), followed by the outcome (reward, punishment, or neutral; 2000ms). There was one block of 90 trials, with 30 rewarding, punishing and neutral trials each, in a randomised order.

2.1.3.3. Perceptual discrimination task 2: Awareness check

The second and final discrimination task was used to assess the objective level of cue awareness, as indexed by same/different discrimination accuracy. No further adjustments to display duration were made, which remained at the level determined in the perceptual discrimination task. There was one block of 100 trials, with 50 same and 50 different trials in a randomised order.

2.2. Analysis and Results

Bayes factors (B) were used to assess the strength of evidence for the alternative hypothesis, H₁, over the null, H₀ (Wagenmakers et al., 2017). All Bayes factors, B, reported here represent the evidence for H₁ relative to H₀; to find the evidence for H₀ relative to H₁, take 1/B. Here, $B_{H(0, x)}$ refers to a Bayes factor in which the predictions of H₁ were modelled as a half-normal distribution with an SD of x (see Dienes & Mclatchie, 2017); the half-normal can be used when a theory makes a directional prediction where x scales the size of effect that could be expected. With the assumptions we used for modelling H₁, as it happened, where an effect yielded a p value less than .02, the Bayes factor was above 6, though there is no guarantee of such a correspondence between B and p values (Lindley, 1957). To indicate the robustness of Bayesian conclusions, for each B, a robustness region will be reported, giving the range of scales that qualitatively support the same conclusion (i.e. evidence as insensitive, or as supporting H₀, or as supporting H₁), notated as: $RR_{B>6} [x1, x2]$ or $RR_{B<1/6} [x1, x2]$ or $RR_{1/6>B>6} [x1, x2]$ where x1 is the smallest SD that gives the same conclusion and x2 is the largest (see Dienes, 2019).

2.2.1. Data pre-processing

In order to account for potential response bias, type I d' (a Signal Detection Theoretic measure of sensitivity to signal versus noise; Stanislaw & Todorov, 1999) was computed for both PDT2 (awareness check) and the main conditioning task. Type I d' can be used to index awareness level corresponding to the objective threshold, where chance performance corresponds to lack of awareness, regardless of confidence or subjective awareness reports. Note that this measure is used here following the procedure of Pessiglione et al (2008). For the PDTs, correct same/different responses were treated as hits, and incorrect responses as false alarms. In the conditioning task, Go responses to rewarding cues were treated as hits, and Go responses to punishing cues as false alarms. Go responses to neutral cues were discounted, as participants are expected to respond arbitrarily to them due to their null outcome.

2.2.2. Awareness Check

At the group level, d' scores for the PDT2 were entered into a one-way t-test against 0, which indicates no ability to discriminate the stimuli (no sensitivity between signal versus noise, akin to chance performance). A Bayes Factor (B) was computed for the difference, with the predictions of H₁ (awareness is present) modelled as a half-normal distribution centred on 0, with an SD equal to a d' of 1 (the average expected effect size corresponding to 70% hit rate (accuracy) and 30% false alarms, an estimate of above-chance and below-ceiling performance).

The results indicate that participants were able to successfully discriminate the stimuli, with the average d' significantly greater than zero ($M = 0.946$, $SE = 0.197$, $p < 0.001$, $B_{H(0,1)} = 25554$, $RR[0.07, 303.5]^1$).

2.2.3. Main Conditioning Task

The d' scores for the conditioning task were entered into a one-way t-tests against 0, indicating lack of discrimination between the cues, and consequently, lack of learning. B was computed for the difference, with the predictions of H1 (learning is present) modelled as a half-normal distribution centred on 0, with an SD equal to 0.7 (the expected effect size if learning is present, derived from Pessiglione et al., 2008).

The results indicate that participants were able to successfully learn, with the average d' significantly greater than zero ($M = 1.793$, $SE = 0.301$, $p < .001$, $B_{H(0,0.7)} = 2514517$, $RR_{B>6}[0.081, 633.5]$).

Figure 1. Learning (d') plotted against objective level of awareness (d') obtained in the supraliminal pilot study. Ribbon represents a 95% confidence interval. $N = 25$.

2.3. Pilot: Discussion

The purpose of the pilot was to establish a rough relation between the level of awareness as measured by the awareness measure by Pessiglione et al (2008), and the level of learning it can support. In

¹ All data analysis in this paper was conducted in JASP (JASP Team, 2022). All Bayes Factors were computed with the normal likelihood distribution, using the calculator available at <https://harry-tattan-birch.shinyapps.io/bayes-factor-calculator/>

the pilot, the mean awareness was $d' = 0.9$, and the mean learning was $d' = 1.793$. These are the crucial facts we need. If both these measures result from the influence of the same knowledge base, namely conscious perception, then as conscious perception goes to zero, both should also go to zero (Dienes, 2015). Thus, on a plot of awareness against learning (*Figure 1*), a line from the point given by the two means going to (0,0) gives a rough estimate of the relation that should be obtained between awareness and learning, assuming it is linear. While there are uncertainties in both the estimates and their linearity, we only need a rough estimate, as we will model uncertainty around this estimate, and robustness regions will be provided. Now we are in a position to proceed with the replication.

The theory that the Pessiglione et al. (2008) method produces unconscious learning involves two predictions: 1) participants will perform at chance on the awareness measure; and 2) participants will show conditioning. These are the crucial tests we will consider below. 1) might be regarded as an outcome neutral test in order for the paradigm to be relevant for showing unconscious learning. From the point of view of a replication, however, it constitutes a crucial test of whether the procedure does result in stimuli being subliminal.

3. Experiment 1: Direct replication

Experiment 1 aims to directly replicate unconscious instrumental conditioning found in Pessiglione et al. (2008). For this reason, all methods are in keeping with those employed in that original study. The original frequentist analyses are supplemented with Bayes Factors, in order to disambiguate potentially non-significant results as either indicating support for the null hypothesis, or indicating insensitive data. The Stage 1 registration can be found at <https://osf.io/gf8jp/>. The in-principle accepted Stage 1 manuscript can be found at <https://osf.io/cmdfs/>. The task code, materials, timestamped raw data files, data processing script, and summarised data files are available at <https://osf.io/ke6yi/>.

3.1. Method

3.1.1. Participants

61 participants (33 females) were recruited at the University of Sussex ($M_{AGE} = 21$; $SD_{AGE} = 3.93$; range = 18-41 years; 5 participants did not report their age). Sample size was determined with the Bayesian Stopping Rule, using previously obtained effect sizes as empirical priors, or cease at 200 participants if the result in the awareness check remains insensitive (see 3.2.2. *Planned Analyses* for detail). In keeping with the original study, participants were told they will be reimbursed with their earnings from the task, but at the end this was rounded to a fixed amount of £6. Ethical approval was granted by the School of Psychology ethics committee at the University of Sussex, and the study was conducted in accordance with the

Declaration of Helsinki. Data for three participants were partially missing due to a technical glitch, yielding a usable sample of 59 participants.

3.1.2. *Stimuli and Materials*

All stimuli and materials used were equivalent to those reported in the original study. The stimuli were 15 randomly chosen characters from Agathodaimon font, presented in white typeface on a black screen, with a size of 240 by 180 pixels. For each participant, 3 were randomly assigned to PDT1, 3 to PDT2, and 9 to the main task, with 1 rewarding, 1 punishing, and 1 neutral cue in each of the 3 blocks. Two black-and-white visual noise masks of the same size as the stimuli were generated by scrambling one character image into 30 by 30 pixel squares. The same two masks were used for all participants in the same fashion (one preceding and one following the target stimulus). The outcome images were a circled £1 coin image for reward, a crossed-out £1 image for punishment, and a greyed-out coin for neutral.

The task was programmed using Matlab 2018b (MathWorks, 2018), running Psychophysics Toolbox (Brainard, 1997), and presented on a Dell LCD monitor with a 60Hz refresh rate (manufactured in 2006 to approximate the screen technology used in the original experiment by Pessiglione et al). Responses were collected with a standard keyboard.

3.1.3. *Procedure*

3.1.3.1. Perceptual discrimination task 1.

Participants were seated with their chin on a chin rest placed at 50cm distance from the screen. Each session commenced with a PDT1 used to determine the individual cue display duration. The duration was either (33ms or 50ms); the largest for which they show chance-level ($\leq 50\%$) discrimination performance. In the task, participants were shown two cues, each forward-backward masked (67ms), separated by a 3s interval indicated with a fixation cross, following the method of Pessiglione et al (2008). Following the display, they were asked to report whether the cues presented were the same or different, using the arrow-keys. The task consisted of 2 blocks of 120 trials (with 60 same and 60 different trials in each, in a randomised order). The first block was conducted with 50ms display duration of each cue. If discrimination accuracy at this stage was at chance (assessed with a chi-squared test for each participant), the task was ended and a 50ms display duration adopted in the main task. If performance in the first block was above chance, the duration was decreased to 33ms for the second block, and performance assessed again. If it was at chance at the end of the second block, the duration of 33ms was adopted in the main task. Two participants who remained above chance at 33ms were not able to take part, yielding a final sample of 56.

3.1.3.2. Main conditioning task.

The task was identical to the original protocol and the pilot study, with the exception that the cues were presented subliminally, for the duration determined in the PDT1 (33 or 50ms). Participants were asked to choose between making a response by pressing a spacebar (Go), or refraining from a response (NoGo), to masked cues (see Fig.2). In each block, one cue was paired with reward, one with punishment, and one with the neutral outcome. Hence, participants could choose either to take a “risky” action (where they might win £1, lose £1, or have a neutral outcome depending on the proceeding cue) or to refrain from acting and thus ensure a neutral outcome.

Each trial began with a fixation cross (500ms), followed by a forward mask (67ms), one of the target cues (determined subliminal display duration), and backward mask (67ms). Subsequently, a question mark appeared on the screen, indicating that the response may be made. Regardless of the response (Go or NoGo), the response window remained open for 3000ms, and participants’ response was collected at the end – Go if the spacebar was being pressed, and NoGo if it was released. Finally, the choice made (Go! or No!) was displayed (500ms), followed by the outcome (reward, punishment, or neutral; 2000ms).

In order to counterbalance motor conditions, the ‘risky’ response was pseudo-randomised to be Go for half the participants, and NoGo for the other half. There were 3 blocks of 120 trials, with 40 trials of each type (rewarding, punishing, neutral). Within each block, the order of the trial types was randomised without constraints.

Figure 2. Trial sequence in Experiment 1 (top) and 2 (bottom). In both experiments, participants were presented with a stimulus (predictive of a rewarding, punishing, or neutral outcome if Go is executed) between two visual masks. Following each presentation, they were asked to make a Go or NoGo response. If a Go response was chosen, the outcome was presented, depending on the type of cue. If a NoGo response was chosen, the outcome was a greyed-out coin. In Experiment 2, each trial ended with an immediate awareness check, composed of a judgment of cue symmetry and confidence.

3.1.3.3. Preference task.

Following each of the three conditioning blocks, the three cues used were shown on the screen side by side, unmasked, in a randomised order. Participants were asked to rate them in order of preference, from most (3) to least (1) liked.

3.1.3.4. Perceptual discrimination task 2.

A PDT2 with 120 trials (60 same, 60 different) and 3 new stimuli was repeated at the end of the testing session. There were no adjustments to the display duration, which was kept at the level determined in PDT1. The task allows to determine whether or not participants' cue awareness remained at chance level.

3.2. Analysis and Results

3.2.1. Data pre-processing

Identical to pilot study.

3.2.2. First crucial test: Awareness check

Absence of awareness was determined by assessing discrimination performance on the second perceptual discrimination task, indexed by type I d' scores (corresponding to the objective threshold of awareness). At the group level, d' scores were entered into a one-way t -test against 0, which indicates no ability to discriminate between the stimuli (no sensitivity between signal versus noise). B was computed on the obtained mean d' , with the H_1 (awareness present) modelled as a half-normal distribution with a mean of 0 and a SD equal to the value derived from the pilot study, following the regression method outlined by Dienes (2015, p.211-213). The mean learning d' (1.793) and the corresponding mean awareness d' (0.9) from the supraliminal pilot were used to estimate the mean awareness d' expected from the obtained level of learning in the main subliminal conditioning task. This was done using the regression line drawn between the supraliminal mean values and the point of origin (no learning, no awareness). The expected d' if awareness is present in the unconscious task (H_1) can then be derived from the learning d' value we actually obtain in this experiment (full sample learning d' $M = 0.63$), resulting in an SD of 0.333 d' units for the model of H_1 for testing awareness. In line with the Bayesian Stopping Rule, data collection would continue until support for the H_0 at the group level is found ($B_{H(0,X)} < 1/6$), or cease at 200 in the event of the data remaining insensitive. The upper cap was determined with the Bayesian sample size estimation method (Dienes, 2015b), using the above regression method to derive the average awareness d' (0.37) expected from Pessiglione et al (2008) learning d' (0.7), and adjusting the awareness SE obtained in the pilot (0.2) in line with sample size increases. A robustness region is reported, as described in the Pilot.

At a group level, awareness d' was significantly above 0 ($M = 0.37$, $SE = 0.09$; $t(55) = 4.24$, $p < .001$, $B_{H(0,0.333)} = 2273.23$, $RR_{B>6}[0.037, 110 \text{ } d' \text{ units}]$; see Fig.3). This result indicates that the full sample was aware in the PDT2, an outcome contrary to that found at this stage in Pessiglione et al. (2008).

Following the original method, performance for every individual compared to chance (50% accuracy) was assessed with chi-square tests. Participants who showed significant above-chance performance (16) were excluded from further analysis, as well as those who explicitly reported seeing the stimuli on-screen (0).

Figure 3. **A:** Awareness (d') obtained at the group level in the perceptual discrimination task of Experiment 1 (direct replication). Plots present the data before any individual exclusions on the awareness criterion ($N = 56$). Asterisks indicate significance at: $*$ = $p < .05$, $**$ = $p < .005$, $***$ = $p < .001$. Tilde indicates a sensitive B (> 6), supporting the H1.

3.2.3. Second crucial test: Main conditioning task

Presence of learning in the conditioning task was assessed with d' scores. d' scores were entered into a one-way t-tests against 0, indicating lack of discrimination between the cues, and consequently, lack of learning. B was computed with H_1 modelled as a half-normal distribution with a mean of 0 and a SD equal to 0.7 (expected effect size if learning is present, derived from Pessiglione et al., 2008). Resulting $B_{H(0,0.7)} > 6$ can be taken as evidence of learning. $B_{H(0,0.7)} < 1/6$, can be taken as evidence for absence of learning. In the event of an insensitive result, data collection will cease at 170 participants (upper cap estimated in the same way as in section 3.2.2., using a learning d' of 0.7 as the expected effect size and learning SE of 0.3 obtained in the pilot). A robustness region will be reported, as described in the pilot.

As pre-registered, following the original analysis steps, after excluding individual participants who showed significant above-chance perception, learning d' was above 0 ($M = 0.37$, $SE = 0.11$; $t(39) = 3.24$, $p < .001$; $B_{H(0,0.7)} = 93.65$, $RR_{B>6} [0.06, 12 d' \text{ units}]$; see Figure 4). This demonstrates that the participants deemed unaware on the PDT2, following the criteria of the original paper, were able to learn the unconscious stimulus-action-outcome associations in the conditioning task.

Figure 4. A: Learning (d') obtained in the main conditioning task of Experiment 1 (direct replication) following the exclusion of participants aware on PDT2 ($N = 40$). **B:** Proportions of correct responses (Go to rewarding or NoGo to punishing cues), on each trial, averaged across the three blocks. Ribbons around the smoothed curves represent 95% CI. **C:** Proportions of Go responses to rewarding (red) vs punishing (blue) cues. Neutral cues are ignored since they provided no outcome. Participants were able to learn to refrain from making a Go response over the course of the block. Asterisks indicate significance at: * = $p < .05$, ** = $p < .005$, *** = $p < .001$. Tilde indicates a sensitive B supporting the H1.

3.2.4. Exploratory analysis: Awareness check

We also assessed the awareness level in the remaining sample following the individual awareness exclusions ($N=40$), which was not a part of the registered analysis. In this sample, awareness d' was not significantly different from 0 ($M = 0.05$, $SE = 0.03$; $t(39) = 1.55$, $p = 0.064$; $B_{H(0,0.333)} = 0.57$, $RR_{B<1/6}[0, 1.2]$; see Fig.5). This shows the awareness data after individual awareness-related exclusions were insensitive. This result is comparable to the findings of Pessiglione et al. (2008).

Figure 5. Awareness (d') obtained at the group level in the perceptual discrimination task of Experiment 1 (direct replication), after individual exclusions on the awareness criterion ($N = 40$). Note that this step was not part of the registered analysis.

Figure 6. **A:** Learning (d') plotted against awareness (d') from PDT2, before any individual exclusions on the awareness criterion ($N = 56$). **B:** Learning (d') plotted against awareness (d') from PDT2, after the exclusion of participants aware on PDT2 ($N = 40$). This depiction was not part of the registered analysis.

3.2.5. Exploratory analysis: Preference ratings

The preference ratings (from 1 [least liked] to 3 [most liked]) were entered into a repeated-measures ANOVA, with cue type (rewarding, punishing, neutral) as a factor. There was no main effect of cue type on preference ratings ($M_{\text{REW}} = 2.33$, $M_{\text{PUN}} = 1.93$, $M_{\text{NEU}} = 2.05$ ranking units; $F(2,342) = 1.408$, $p = .246$). Note that preference ratings were lost for the first two participants due to a programming error.

3.3. Conclusions of Experiment 1

Experiment 1 aimed to directly replicate unconscious instrumental conditioning found in Pessiglione et al. (2008), following the exact methods and analysis steps, supplemented with Bayes factors. Using type I d' as a proxy for learning (i.e. being able to discriminate between the stimuli), we found sensitive evidence in favour of learning in the sample deemed unaware according to the original criteria (i.e. the participants that remained after removing those showing individual above chance performance on the PDT2). This result replicates the effect found in the original paper.

However, we were not able to replicate the first crucial test, namely that awareness at the group level, assessed with PDT2, was absent. In the original paper, absence of awareness on the PDT2 was asserted through a non-significant difference between the obtained d' ($M = 0.08$, $SD = 0.2$, $SE = 0.04$) from 0, with no participants being excluded due to individual above-chance performance. In the present study, the full sample was found to be significantly above chance, with a Bayes factor providing strong support for the presence of awareness as indicated by this test. After individual exclusions of above-chance participants in an exploratory analysis (section 3.2.4.), awareness in the remaining sample was not significantly different from 0, with an insensitive Bayes factor (note that the aforementioned sensitive evidence for learning was obtained on this sample). We postulate that the assertion of absence of awareness from a non-significant t-test against 0 in the original paper resulted from a similar situation. Indeed, computing a Bayes factor on the provided awareness d' statistics from that paper also yields an insensitive result ($B_{H(0, 0.333)} = 1.28$).

We would like to note the caveat that our registered stopping rule (sensitive evidence for H_0 , or 200 participants if insensitive in the full sample) did not explicitly allow for the effect to go in the other direction, i.e. sensitive evidence for H_1 . While this was not an expected outcome, logically the requirement for sensitivity should have been bi-directional and as such, the very strong support for H_1 in our data warranted stopping data collection.

To conclude, while we were able to replicate the presence of learning in the conditioning task, we were not able to assert absence of awareness at the group level in the PDT2, a necessary precondition for claiming that unconscious instrumental conditioning took place. Consequently, in Experiment 2 we propose a series of methodological improvements to reassess this claim.

4. Experiment 2

Experiment 2 was designed to be conducted only in the event of Experiment 1 replicating the effect found in Pessiglione et al (2008). In light of the methodological and theoretical advances and debates in the field of unconscious learning (e.g. Dienes, 2015; Mudrik et al., 2014; Newell & Shanks, 2014), Experiment 2 aims to replicate the result, introducing changes to the paradigm targeted at increasing the methodological rigour. The Stage 1 registration can be found at <https://osf.io/gf8jp/>. The in-principle accepted Stage 1 manuscript can be found at <https://psyarxiv.com/p9dgn/>. The task code, materials, timestamped raw data files, data processing script, and summarised data files are available at <https://osf.io/t23by/>.

Firstly, in the original study, the measures in the awareness check and the learning task pertain to two different aspects of decision-making. The perceptual discrimination tasks (serving as threshold-setting and as awareness check) required a same/different perceptual judgment of 2 cues, separated by a 3s interval. In contrast, the main conditioning task required an approach/avoid response after a single stimulus. Hence, the measure used in the perceptual discrimination task reflected a different decision process than was required in the conditioning task, violating the relevance and sensitivity criteria (Berry & Dienes, 1993; Newell & Shanks, 2014). As such, the threshold-setting task was amended to match the conditioning task more closely. In order to further enhance the sensitivity of the task and ensure that it reliably settles on sub-threshold conditions for the largest possible number of participants, stimulus contrast was reduced and the exposure time range was extended beyond the original limit of either 50 or 33ms. Stimulus exposure was individually titrated for each participant in a stepwise manner based on both the accuracy and reported confidence in visual perceptions.² This procedure was used extensively in previous research for effective identification of sub-threshold conditions (Scott et al., 2018; Skora et al., 2021; Skora, Livermore, et al., 2022).

Secondly, the separate awareness check was replaced with a trial-by-trial measure, allowing to access the information about participants' awareness in a more immediate fashion (Berry & Dienes, 1993; Newell & Shanks, 2014).

Thirdly, the original task design leaves open the possibility that participants might occasionally experience awareness of the stimuli in the learning task, which does not become apparent in the final PDT2.

² Note that the contrast reduction, time-range and mask duration extension, and individual threshold titration are minor changes introduced and approved following the original in-principle accepted Stage 1 plan, but before any data were collected. In the original Stage 1 plan for Experiment 2, the task could not sensitively settle on sub-threshold conditions, which is a necessary assumption for the conditioning task. Keeping the task in line with the original Stage 1 plan would have likely resulted in an overwhelmingly large proportion of aware participants, precluding any meaningful conclusions. The introduced minor changes maximise the likelihood of arriving at suitable sub-threshold conditions. Crucially, the purpose of the experiment and the analyses conducted remained unchanged from the approved Stage 1.

This might occur either where the same brief moments of awareness do not reoccur in the PDT2 or where they are too infrequent to significantly influence the overall objective accuracy measure. Reliably excluding individual trials is impossible when only objective discrimination measures are collected. With this in mind, the trial-by-trial awareness check also included confidence ratings, allowing to exclude trials where participants were subjectively aware. Because the initial staircase should sensitively settle on sub-threshold conditions, a trial-by-trial check should only elicit a small number of aware trials. While post-hoc trial exclusion of conscious trials can lead to regression to the mean (Shanks, 2017), its effect is negligible if the majority of trials are unconscious. We examined this assumption by modelling worst-case scenarios for the presence of conscious trials at different proportions of observed unconscious trials, at different error rates (see Supplementary material). This allows us to determine the maximum percentage of conscious trials which could inadvertently contribute to unconscious knowledge. Thus, if we observe 80% of unconscious trials (leaving room for error in the remaining 0-20%), the maximum proportion of conscious trials possibly contained within our observed unconscious trials is 1.59%. This would be the percentage of our conscious knowledge potentially accounting for the learning effect. Using the observed d' found in the pilot study of 1.8, we find that the maximum influence from conscious knowledge where 80% of responses are attributed to unconscious responding is $d' = 0.03$. We consider this negligible and as such will adopt a strategy whereby provided the proportions of responses attributed to conscious responding does not exceed 20%, our exclusion criteria will be applied (see section 4.2.2.).

Finally, the forward and backward masks were generated afresh on each trial by randomly scrambling a black-and-white noise image. The use of different masks on each trial reduces the likelihood of participants building erroneous associations from possible salient repetitive features of the masks (or of some stimulus-mask combinations). Mask duration was also extended from 67 to 300ms in order to ensure robust masking with the larger possible range of stimulus display durations.

As in Experiment 1, original frequentist analyses were supplemented with Bayes Factors.

4.1. Method

4.1.1. Participants

45 participants (33 females) were recruited at the University of Sussex ($M_{AGE} = 20.59$, $SD_{AGE} = 1.52$, range = 18-25 years; 8 participants did not report their age). Sample size was determined with the Bayesian Stopping Rule, using previously obtained effect sizes as empirical priors, or cease at 170 participants should the result remain insensitive (see 4.2.2. *Planned Analyses* for detail). In keeping with the original study, participants were told they will be reimbursed with their earnings from the task, but at the end this will be rounded up to a fixed amount of £6. Ethical approval was granted by the School of Psychology ethics

committee at the University of Sussex, and the study was conducted in accordance with the Declaration of Helsinki.

4.1.2. *Stimuli and Materials*

The stimuli were 11 characters from Agathodaimon font, chosen pseudo-randomly to ensure six symmetrical and seven asymmetrical characters. All were presented in light grey typeface (RGB: 141,141,141) on a darker grey background (RGB: 115,115,115), with a size of 240 by 180 pixels. For each participant, two characters were randomly assigned to PDT1 (with one symmetrical and one asymmetrical character). The nine remaining stimuli were pseudo-randomly assigned to the main task, with one rewarding, one punishing, and one neutral cue in each of the three blocks, such that each block contained both symmetrical and asymmetrical cues. Both the forward and backward masks were generated afresh on each trial by randomly scrambling a 240 by 180 pixels black-and-white noise image in blocks of 3x3 pixels.

4.1.3. *Procedure*

4.1.3.1. Perceptual discrimination task 1.

Each session commenced with a PDT, allowing to determine the individual cue display duration. In this task (in contrast to Experiment 1), participants were shown a fixation cross (500ms), followed by a single cue (display duration starting at 600ms), forward-backward masked (300ms). After each sequence, participants were asked to report whether the cue presented was symmetrical or asymmetrical, using the arrow-keys. Next, they were asked to report whether they had any confidence in their judgment, or if they were guessing, also using the arrow-keys. They were instructed to report 'some confidence' if they had any degree of confidence, and 'total guess' only if they felt they did not perceive the cue and were responding randomly. With every correct response made with confidence (taken as aware), the display duration of the cue was reduced by 50ms on the subsequent trial. When a duration of 100ms was reached, or the first guess was made, the display duration returned to the previous level (+50ms), and subsequently decreased in 16ms steps on the following trials (corresponding to single screen refresh duration in a 60Hz screen used here). This reduction continued until participants reported guessing again, at which point the display duration remained the same until a guess was reported on six consecutive trials, regardless of response accuracy. The cue display duration on those trials was applied as the individual threshold of conscious awareness in the main conditioning task. Participants who reach the minimum possible display duration (16ms) without guessing were not able to take part. The average established display duration was 80ms ($SD = 0.044ms$; mode = 84ms).

4.1.3.2. Main conditioning task.

The conditioning task was identical to Experiment 1, with the exception that an awareness check was added at the end of each trial. Following feedback presentation (reward/punishment), participants were asked to report if the masked cue was symmetrical or asymmetrical, using the arrow keys. Next, they were asked to report their confidence in that judgment on a binary scale (between ‘some confidence’ and ‘total guess’). There were 3 blocks of 120 trials, with 40 trials of each type (rewarding, punishing, neutral). Within each block, the order of the trial types was randomised without constraints. Prior to beginning, participants were explicitly instructed that the symmetry judgments are not related to the rewarding/punishing outcomes. They were also shown a different pair of example cues to illustrate what is meant by symmetry.

4.1.3.3. Preference task.

Identical to Experiment 1.

4.2. Analysis

4.2.1. Data pre-processing

Identical to Experiment 1.

4.2.2. Crucial test: Main conditioning task

Individual trials where participants made a correct symmetry judgment with confidence were marked as ‘aware’ trials and excluded ($M = 23.76$ trials, mode = 1 trial)³. In cases where exclusions exceeded 20% (72) of trials, the entire participant was excluded from analysis (N excluded participants = 4, final N = 41). The remaining trials were analysed with type I d' in a manner identical to Experiment 1. B was computed with H_1 modelled as a half-normal distribution with a mean of 0 and a SD equal to 0.7 d' units (expected effect size if learning is present, derived from the original study). Resulting $B_{H(0,0.7)} > 6$ can be taken as evidence of learning. $B_{H(0,0.7)} < 1/6$, can be taken as evidence for absence of learning. In the event of an insensitive result, data collection was determined to cease at 170 participants (upper cap estimated in the same way as in section 3.2.3., using a learning d' of 0.7 as the expected effect size and learning SE of 0.3 d' units obtained in the pilot). A robustness region is reported.

The d' scores from the conditioning task after excluding aware trials were entered into a one-way t-test against 0, indicating lack of discrimination between the cues. The results indicate that participants were not able to discriminate between the subliminally-presented stimuli ($M = 0.03$, $SE = 0.04$ d' units; $t(40) = 0.77$, $p = .447$, $B_{H(0,0.7)} = 0.11$, $RR_{B < 1/6}[0.5, \text{Inf } d' \text{ units}]$; see Figure 6), with B indicating strong evidence for the absence of learning.

³ For comparison, incorrect and confident trials yielded $M = 14.87$, mode = 3 trials.

Figure 7. A: Learning (d') obtained in the main unconscious conditioning task of Experiment 2 ($N = 41$). **B:** Proportions of correct responses (Go to rewarding or NoGo to punishing cues), averaged across the three blocks (unaware trials only). **C:** Proportions of Go responses to rewarding (red) vs punishing (blue) cues. Neutral cues are ignored since they provided no outcome. Participants were not able to learn to respond Go more often to rewarding than punishing cues (unaware trials only). Cross indicates a sensitive B supporting the H_0 .

4.2.3. Exploratory analysis: Main conditioning task, guess trials

We also assessed the learning effect in the unaware trials, this time considering all confident trials as aware ($M = 38$ trials, mode = 9 trials), regardless of their accuracy (in contrast to considering only correct and confident trials as aware in the crucial analysis above). Those confident trials were excluded. The same participant exclusions as in the crucial analysis above applied. The results indicate that participants were not able to discriminate between the subliminally-presented stimuli ($M = 0.02$, $SE = 0.04$ d' units; $t(40) = 0.58$, $p = .283$, $B_{H(0,0.7)} = 0.09$, $RR_{B<1/6}[0.4, \text{Inf } d' \text{ units}]$), with B indicating strong evidence for the absence of learning. This result is aligned with the findings of the crucial analysis.

4.2.4. *Exploratory analysis: Main conditioning task, all trials*

We also assessed the learning effect without applying the awareness exclusion criteria. Without applying participant-wise exclusions (i.e. all participants, all unaware trials), the learning d' was not significantly different from 0, with an insensitive B according to our criteria ($M = 0.09$, $SE = 0.09$ d' units; $t(44) = 1.11$, $p = .272$, $B_{H(0,0.7)} = 0.35$, $RR[0, 1.5$ d' units]). Without applying any exclusions at all (i.e. all participants, all aware and unaware trials), the learning d' was also not significantly different from 0, also with an insensitive B (learning d' $M = 0.14$, $SE = 0.09$ d' units; $t(44) = 1.60$, $p = .117$, $B_{H(0,0.7)} = 0.79$, $RR_{1/6>B>6}[0, 3.3$ d' units]).

4.2.5. *Exploratory analysis: Main conditioning trials, aware trials*

We also assessed the learning effect on the trials classified as aware (correct and confident) on the trial-by-trial awareness check. However, due to the small proportion of aware trials ($M = 23.76$, mode = 1 trial), d' was only computable or meaningful for a small number of participants. Without applying any further exclusions (i.e. all aware trials, all participants with a computable d'), the aware learning d' was numerically higher than for unaware trials, but not significantly different from 0, with an insensitive B, likely due to the small resulting sample size ($M = 0.23$, $SE = 0.31$ d' units; $t(25) = 0.74$, $p = 0.464$, $B_{H(0,0.7)} = 0.77$, $RR_{1/6>B>6}[0, 3.7$ d' units]). However, linear regression (on the sample without awareness-related trial and participant exclusions), revealed that the number of aware trials significantly predicted learning d' ($R^2 = 0.62$, $F(1,43) = 70.62$, $p < .001$; $\beta = 0.009$ d' units/aware trial, $SE = 0.001$, $t(44) = 8.40$, $p < .001$; $B_{H(0,0.006)} = 42620030835485600$, $RR_{B>6}[0.00018, 3.42$ d' units]). B was computed for the raw regression slope using the ratio-of-means heuristic (Dienes, 2019), computing the expected slope for the relation between the number of aware trials and learning d' from the ratio of their means (M aware trials = 23.76, M learning $d' = 0.14$). The resulting ratio (0.006) was entered into B calculation as the scaling factor.

While we classified only correct and confident trials as aware in the above analysis, for completeness we also assessed the learning effect on all confident trials (whether correct or incorrect). Without applying any further exclusions (i.e. all confident trials, all participants with a computable d'), learning d' was again numerically higher than for unaware trials ($M = 0.28$, $SE = 0.20$ d' units), but not significantly different from 0, with an insensitive B, again likely due to the small sample size ($t(35) = 1.42$, $p = 0.083$, $B_{H(0,0.7)} = 1.24$, $RR[0, 3.1$ d' units]).

4.2.6. *Exploratory analysis: Preference ratings*

The preference ratings (from 1 [least liked] to 3 [most liked]) given by included participants were entered into a repeated-measures ANOVA, with cue type (rewarding, punishing, neutral) as a factor. There was a significant main effect of cue type on the ratings ($M_{REW} = 2.14$, $M_{PUN} = 1.99$, $M_{NEU} = 1.86$ ranking units;

$F(2,366) = 3.583$, $p = 0.029$). Post-hoc comparison of means (Tukey-adjusted for multiple comparisons) revealed that only the difference between rewarding and neutral cues was significant and sensitive ($M_{\text{diff}} = 0.28$, $SE = 0.10$ ranking units; $t = -2.675$, $p = 0.021$, $B_{H(0,0.45)} = 18.09$), in contrast to the rewarding-punishing difference ($M_{\text{diff}} = 0.15$, $SE = 0.10$ ranking units; $t = -1.416$, $p = 0.335$, $B_{H(0,0.3)} = 1.60$) and the punishing-neutral difference ($M_{\text{diff}} = 0.13$, $SE = 0.10$; $t = -1.259$, $p = 0.419$, $B_{H(0,0.15)} = 1.71$). The expected values for calculating Bs for the differences were estimated from the bar chart provided in Pessiglione et al. (2008).

4.3. Conclusions of Experiment 2

Experiment 2 aimed to replicate the result of Experiment 1, introducing changes to the paradigm targeted at increasing its methodological rigour, and satisfying the relevance, sensitivity, and immediacy criteria. Those changes included amending the type of decision required in PDT1 to match that required in the conditioning task, introducing an individually-titrated threshold of conscious awareness, and replacing the separate awareness check in PDT2 with a trial-by-trial check, which allowed robust exclusion of all subjectively aware trials from analysis.

Again using type I d' as a proxy for learning (i.e. being able to discriminate between the stimuli), we found sensitive evidence for the absence of learning, when all subjectively aware trials were excluded. This result challenges the finding of learning with an insensitive awareness effect found in the original paper (Pessiglione et al., 2008), and in the same way, clarifies our direct replication (Experiment 1) – showing, in both cases, that apparent evidence for learning was likely due to inadequate exclusion of aware trials.

In the exploratory analysis of the preference ratings, the rewarding stimulus was rated significantly higher than the neutral cue, suggesting that participants may have acquired some knowledge of the stimulus values over the course of learning. However, successful instrumental learning should result in the largest difference between ratings of rewarding and punishing stimuli (in addition to the difference between the valenced and neutral cues), a comparison which failed to reach sensitivity, preventing any clear conclusions. This result could also be attributed to the fact that while participants who were aware on over 20% of trials were excluded, the remaining sample still included incidental moments of stimulus awareness, which may have affected the overall preference ratings.

5. General Discussion

The present study revisited the question of unconscious instrumental conditioning by attempting to replicate the seminal study by Pessiglione et al. (2008) both directly (Experiment 1), and with amendments targeted at increasing the methodological and analytical rigour (Experiment 2). We find that learning took place when replicating the original methods directly (albeit we could not assert that the full sample was unaware in the separate awareness check, PDT2). However, following the enhancement of the

sensitivity of the threshold-setting perceptual discrimination task, and the introduction of a trial-by-trial awareness check permitting exclusion of individual trials, we found evidence for the absence of learning. All conclusions were based on informed Bayes factors.

This difference can be attributed chiefly to the ability to detect individual trials where participants showed stimulus awareness by discriminating stimulus symmetry correctly and with confidence, a measure introduced in Experiment 2. Absence of such a granular method of assessing awareness in Experiment 1 meant that occasional trials with subjective awareness of the stimulus would have remained undetected. This, in turn, increased the chances that learning was not fully unconscious, such that the observed effect could have arisen through a mixture of conscious and unconscious trials. As is evident in Experiment 2, even with a sensitive threshold-setting task, most participants exhibited some degree of awareness on occasion. This may reflect visual adaptation and normal fluctuations in visual sensitivity. The ability to identify those trials in an immediate (i.e., trial-by-trial) manner and to exclude them ensured that only genuinely unaware trials were analysed, uncontaminated by individual aware trials. Critically, when only unaware trials were analysed, learning did not take place.

This result supports the theoretical perspectives converging on the view that flexible, goal-oriented behaviour, supported by long-range, recurrent information processing patterns, requires conscious perceptual access to relevant stimuli (Baars, 2002; Dehaene & Changeux, 2011; Dehaene, Charles, King, & Marti, 2014; Lamme, 2006; van Gaal, de Lange, & Cohen, 2012). While simpler forms of associative learning (including classical conditioning) may be feasible in the absence of conscious perception of relevant stimuli, instrumental conditioning is considerably more complex. It requires the integration of information from separate modalities involved with visual processing of the stimulus, extracting the expected stimulus value, deploying a selective action or refraining from action altogether, relating the reinforcement to the stimulus-dependent action, and comparing the actual outcome to the expected outcome, in order to update the expected stimulus value and store it in memory for subsequent interactions. It may therefore not be so surprising that conscious perception appears to be crucial for the flexible and long-lasting information processing strategies underpinning instrumental conditioning. An unconsciously presented stimulus might simply not have the capacity to evoke such a broad range of activity (although it remains an open, and worthwhile, question at which stage the process breaks down in the unconscious case).

Our result also adds to the growing body of work demonstrating that complex forms of learning, including instrumental conditioning, cannot operate without conscious perception of the stimulus (Mertens & Engelhard, 2020; Reber et al., 2018; Travers et al., 2018). Since Stage 1 acceptance of this registered report, we have also produced corroborating evidence for both trace and delay instrumental conditioning scenarios occurring only with conscious perception in our paradigms (Skora et al., 2021; Skora, Livermore, et al., 2022; Skora, Scott, et al., 2022). Unconscious stimuli are also less likely to drive long-term

behavioural adaptations, in comparison to consciously perceived stimuli (e.g. in conflict adaptation or post-error slowing; de Lange et al., 2011; Kunde et al., 2012; van Gaal et al., 2012). Together, this evidence supports the notion that learning instrumental associations requires perceptual conscious access. However, learning complex structures, including those involving valence, may occur implicitly when the stimuli are consciously perceived (Jurchiř, 2022; Jurchiř et al., 2020; Waroquier et al., 2020).

Such a conclusion has implications for the debate about the function(s) of consciousness, supporting the position that adaptive behaviour – even simple, entirely deterministic instrumental behaviour of the kind presented here – requires perceptual consciousness in order to successfully operate. This is in line with theoretical proposals that consciousness is closely linked to action, providing a frame of reference for agents' interactions with the world in an embodied fashion (Clark, 2016; Land, 2012; Merker, 2005; Seth et al., 2016). On some theories, consciousness may be necessary to support flexible, longer-term decision-making going beyond simple stimulus-stimulus or stimulus-response associations, thus permitting building complex, counterfactual models of the agent in its world. Complex learning, labelled unlimited associative learning (UAL), has also been considered an evolutionary marker of consciousness (Birch et al., 2020, 2021; Ginsburg & Jablonka, 2019). UAL is defined as a capacity of agents to learn about themselves and their environments in an open-ended, unlimited (within lifetime) fashion. UAL, including learning novel responses and trace conditioning, in addition to learning from compound features, second-order conditioning, and flexible re-learning, is proposed to rely on conscious access, and constitute a positive marker of consciousness. While the general claim may be constrained to perceptual consciousness (given evidence for implicit learning; see e.g. (Dienes & Seth, in press), the results of the present experiment indeed support such adaptive approaches to understanding the function of consciousness (Cleeremans & Tallon-Baudry, 2022).

One might worry that d' is a limited proxy for assessing learning, since it collapses performance across an entire block and therefore discounts potential within-block incremental improvements in the learning process (as would be expected during conditioning). However, inspection of the learning curves and the proportion of Go responses in response to rewarding and punishing stimuli strongly suggests that participants failed to improve their responses as the block progressed (in Experiment 2, where awareness was appropriately controlled, Fig.6 B, C). Due to this, and considering also the absence of discrimination between the cues overall, we deemed it redundant to conduct finer-grained analyses of the learning process, such as reinforcement learning modelling as conducted in the original Pessiglione et al. (2008) paper.

The conditioning paradigm investigated here constitutes a variant of trace conditioning: the stimulus, response, and outcome are temporally separated from each other (although the ongoing response bridged the gap between stimulus and the outcome). Could this temporal separation have

impaired the integration of stimulus information needed to allow the stimulus to become predictable of a specific action-dependent outcome? Perhaps bringing the events (stimulus, response, outcome) to the point of overlap, as in delay conditioning, would maximise the likelihood of successful instrumental learning. However, recently we demonstrated that participants were not able to learn in a comparable task even in a delay conditioning scenario with primary (appetitive and aversive) outcomes (Skora et al., 2021).

It is of course impossible to entirely exclude the possibility that modifications to a task will inadvertently disrupt learning, or make any learning more difficult to detect. In our Experiment 2, the symmetry and confidence judgments, needed for the trial-by-trial awareness check, introduced a longer spacing between the trials, which could increase the cognitive load during the learning process. It is conceivable that both of these factors (increased temporal spacing, increased cognitive load) could have impeded learning.

The way masking is implemented can also affect stimulus processing and, consequently, learning. Using the same masks on every trial (like in Experiment 1) can introduce systematic mask-cue interactions (for instance, a mask can bring out features of a given stimulus more than of another, making it more distinguishable or prompting participants to build associations from repetitive combinations). To avoid these potential confounds, we used a different mask on each trial in Experiment 2. However, it is also conceivable that new visual features of each new mask on every trial could blur cue-outcome associations, or cause participants to search for links between masks and outcomes (despite being explicitly instructed not to). We consider those issues unlikely, since the same masking method has been successfully used to show simpler associative learning (e.g. Scott et al., 2018).

Finally, while the trial-by-trial awareness check in Experiment 2 constituted a methodological improvement, it still leaves open a possibility that the knowledge gained from the conscious (albeit excluded) trials affected behaviour on the unconscious trials. We guarded against this by excluding participants with a high proportion (over 20%) of aware trials. For those with fewer than 20% of such trials, we demonstrate that the potential effect of conscious knowledge is negligible (see the registered Supplementary material).

Altogether, even though any task modification may carry the possibility of unintended consequences, the modifications we made were guided solely by the goal of improving detection of (un)awareness. Still, it remains a worthwhile pursuit to investigate unconscious instrumental learning in different kinds of paradigms, for instance with different types of decisions than Go/NoGo, different stimulus suppression methods, or different stimuli. This would be necessary to make general claims about the feasibility of instrumental learning in absence of awareness, or at varying degrees of awareness.

To conclude, our study revisited the question of the feasibility of unconscious instrumental conditioning by attempting to replicate a previous study by Pessiglione et al. (2008) both directly, and in a

second experiment with improvements to methodology and analysis. We found evidence for learning when replicating the original methods directly. However, following the enhancement of the sensitivity of the threshold-setting task, and the introduction of immediate and relevant trial-by-trial awareness checks allowing for exclusion of individual trials, we found evidence for the absence of learning in response to subliminally presented stimuli. This result provides robust evidence that instrumental conditioning cannot be achieved without stimulus awareness in a simple and therefore potentially highly general paradigm, in line with other emerging evidence that complex forms of learning may rely on conscious perception of the relevant stimuli. Altogether, our results support the theoretical view that perceptual consciousness may be necessary for complex, flexible processes, especially where selective action and behavioural adaptation are required, and they contribute to mapping out the role of consciousness in adaptive behaviour.

6. Acknowledgments

This work was supported by the Sackler Centre for Consciousness Science, funded by the Dr. Mortimer and Theresa Sackler Foundation. We thank the Reviewers for invaluable suggestions at both stages of this pre-registered report.

7. Declaration of interests

Declarations of interests: none reported.

References:

- Axelrod, V., & Rees, G. (2014). Conscious awareness is required for holistic face processing. *Consciousness and Cognition*, 27, 233–245. <https://doi.org/10.1016/j.concog.2014.05.004>
- Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, 6(1), 47–52. [https://doi.org/10.1016/S1364-6613\(00\)01819-2](https://doi.org/10.1016/S1364-6613(00)01819-2)
- Baars, B. J., Ramsøy, T. Z., & Laureys, S. (2003). Brain, conscious experience and the observing self. *Trends in Neurosciences*, 26(12), 671–675. <https://doi.org/10.1016/j.tins.2003.09.015>
- Berry, D. C., & Dienes, Z. (1993). *Implicit Learning: Theoretical and Empirical Issues*. Lawrence Erlbaum.
- Birch, J., Ginsburg, S., & Jablonka, E. (2021). The learning-consciousness connection. *Biology & Philosophy*, 36, 36–49. <https://doi.org/10.1007/s10539-021-09802-5>
- Birch, J., Ginsburg, S., & Jablonka, E. (2020). Unlimited Associative Learning and the origins of consciousness: a primer and some predictions. *Biology and Philosophy*, 35(6), 56. <https://doi.org/10.1007/s10539-020-09772-0>
- Blake, R., & Fox, R. (1974). Adaptation to invisible gratings and the site of binocular rivalry suppression. *Nature*, 249(5456), 488–490. <https://doi.org/10.1038/249488a0>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436.
- Breitmeyer, B. G. (2015). Psychophysical “blinding” methods reveal a functional hierarchy of unconscious visual processing. *Consciousness and Cognition*, 35, 234–250. <https://doi.org/10.1016/j.concog.2015.01.012>
- Clark, A. (2016). *Surfing Uncertainty: Prediction, action and the embodied mind*. Oxford University Press.
- Clark, R. E., & Squire, L. R. (1998). Classical conditioning and brain systems: The role of awareness. *Science*, 280, 77–81.
- Cleeremans, A., & Tallon-Baudry, C. (2022). Consciousness matters: phenomenal experience has functional value. *Neuroscience of Consciousness*, 1, 1–11. <https://doi.org/10.1093/nc/niac007>
- de Lange, F. P., van Gaal, S., Lamme, V. A. F., & Dehaene, S. (2011). How awareness changes the relative weights of evidence during human decision-making. *PLoS Biology*, 9(11). <https://doi.org/10.1371/journal.pbio.1001203>
- Dehaene, S., & Changeux, J.-P. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, 70(2), 200–227. <https://doi.org/10.1016/j.NEURON.2011.03.018>
- Dehaene, S., Charles, L., King, J.-R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, 25, 76–84. <https://doi.org/10.1016/j.conb.2013.12.005>
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79, 1–37. www.elsevier.com/locate/cognit
- Dehaene, S., Naccache, L., Cohen, L., Bihan, D. Le, Mangin, J.-F., Poline, J.-B., & Rivière, D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, 4(7), 752–758. <https://doi.org/10.1038/89551>
- Dehaene, S., Sergent, C., & Changeux, J.-P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences*, 100(14), 8520–8525. <https://doi.org/10.1073/pnas.1332574100>
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5(781), 1–17. <https://doi.org/10.3389/fpsyg.2014.00781>
- Dienes, Z. (2015a). *How many participants might I need? - YouTube*. https://www.youtube.com/watch?v=10Lsm_o_GRg
- Dienes, Z. (2015b). How Bayesian statistics are needed to determine whether mental states are unconscious. In *Behavioral Methods in Consciousness Research* (pp. 199–220). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199688890.003.0012>
- Dienes, Z. (2016). How Bayes factors change scientific practice. *Journal of Mathematical Psychology*, 72, 78–89. <https://doi.org/10.1016/j.jmp.2015.10.003>
- Dienes, Z. (2019). How do I know what my theory predicts? *Advances in Methods and Practices in Psychological Science*, 1–14. <https://doi.org/10.1177/2515245919876960>
- Dienes, Z., & McLatchie, N. (2017). Four reasons to prefer Bayesian analyses over significance testing.

- Psychonomic Bulletin & Review*, 25(1), 207–218. <https://doi.org/10.3758/s13423-017-1266-z>
- Dienes, Z., & Seth, A. K. (n.d.). Conscious And Unconscious Mental States Zoltan Dienes and Anil K. Seth In: O. Braddick (Ed.), *Oxford Research Encyclopedia of Psychology*, Oxford University Press. <https://oxfordre.com/psychology>. In *Oxford Research Encyclopedia of Psychology* (pp. 1–39). Oxford University Press.
- Dupoux, E., Gardelle, V. de, & Kouider, S. (2008). Subliminal speech perception and auditory streaming. *Cognition*, 109(2), 267–273. <https://doi.org/10.1016/j.cognition.2008.06.012>
- Faivre, N., Mudrik, L., Schwartz, N., & Koch, C. (2014). Multisensory integration in complete unawareness: Evidence from audiovisual congruency priming. *Psychological Science*, 25(11). <https://doi.org/10.1177/0956797614547916>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/bf03193146>
- Ginsburg, S., & Jablonka, E. (2019). *Evolution of the sensitive soul*. MIT Press.
- Jurchiş, R. (2022). Unconscious knowledge of rewards guides instrumental behaviors via conscious judgments. *Cognitive, Affective and Behavioral Neuroscience*, 1, 1–14. <https://doi.org/10.3758/S13415-022-01044-6/FIGURES/5>
- Jurchiş, R., Costea, A., Dienes, Z., Miclea, M., & Opre, A. (2020). Evaluative conditioning of artificial grammars: Evidence that subjectively-unconscious structures bias affective evaluations of novel stimuli. *Journal of Experimental Psychology: General*, 149(9), 1800–1809. <https://doi.org/10.1037/xge0000734>
- Knight, D. C., Nguyen, H. T., & Bandettini, P. A. (2003). Expression of conditional fear with and without awareness. *Proceedings of the National Academy of Sciences*, 100, 15270–15283.
- Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362(1481), 857–875. <https://doi.org/10.1098/rstb.2007.2093>
- Kouider, S., & Dupoux, E. (2001). A functional disconnection between spoken and visual word recognition: evidence from unconscious priming. *Cognition*, 82(1), B35–B49. [https://doi.org/10.1016/S0010-0277\(01\)00152-4](https://doi.org/10.1016/S0010-0277(01)00152-4)
- Kunde, W., Reuss, H., & Kiesel, A. (2012). Consciousness and cognitive control. *Advances in Cognitive Psychology*, 8(1), 9–18. <https://doi.org/10.5709/acp-0097-x>
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10(11), 494–501. <https://doi.org/10.1016/j.tics.2006.09.001>
- Land, M. F. (2012). Minireview The Operation of the Visual System in Relation to Action. *Current Biology*, 22, 811–817. <https://doi.org/10.1016/j.cub.2012.06.049>
- Lin, Z., & He, S. (2009). Seeing the invisible: The scope and limits of unconscious processing in binocular rivalry. *Progress in Neurobiology*, 87(4), 195–211. <https://doi.org/10.1016/J.PNEUROBIO.2008.09.002>
- Lindley, D. V. (1957). A Statistical Paradox. *Biometrika*, 44(1/2), 187. <https://doi.org/10.2307/2333251>
- Mastropasqua, T., & Turatto, M. (2015). Attention is necessary for subliminal instrumental conditioning. *Scientific Reports*, 5(1290), 1–7. <https://doi.org/10.1038/srep12920>
- Melloni, L., Molina, C., Pena, M., Torres, D., Singer, W., & Rodriguez, E. (2007). Synchronization of neural activity across cortical areas correlates with conscious perception. *Journal of Neuroscience*, 27(11), 2858–2865. <https://doi.org/10.1523/JNEUROSCI.4623-06.2007>
- Merker, B. (2005). The liabilities of mobility: A selection pressure for the transition to consciousness in animal evolution. *Consciousness and Cognition*, 14(1), 89–114. [https://doi.org/10.1016/S1053-8100\(03\)00002-3](https://doi.org/10.1016/S1053-8100(03)00002-3)
- Mertens, G., & Engelhard, I. M. (2020). A systematic review and meta-analysis of the evidence for unaware fear conditioning. *Neuroscience & Biobehavioral Reviews*, 108, 254–268. <https://doi.org/10.1016/j.neubiorev.2019.11.012>
- Mudrik, L., Faivre, N., & Koch, C. (2014). Information integration without awareness. *Trends in Cognitive*

- Sciences*, 18, 488–496. <https://doi.org/10.1016/j.tics.2014.04.009>
- Newell, B. R., & Shanks, D. R. (2013). Unconscious Influences on Decision Making : A Critical review. *Behavioural and Brain Sciences*, 37(1), 1–19. <https://doi.org/10.1111/j.1365-2141.1970.tb01443.x>
- Olsson, A., & Phelps, E. A. (2004). Learned Fear of “Unseen” Faces After Pavlovian, Observational, and Instructed Fear. *Psychological Science*, 15, 822–828.
<http://journals.sagepub.com/doi/pdf/10.1111/j.0956-7976.2004.00762.x>
- Peirce, C. S., & Jastrow, J. (1886). On small differences of sensation. *Memoirs of the National Academy of Sciences*, 3, 73–83.
- Pessiglione, M., Petrovic, P., Daunizeau, J., Palminteri, S., Dolan, R. J., & Frith, C. D. (2008). Subliminal instrumental conditioning demonstrated in the human brain. *Neuron*, 59, 561–567.
<https://doi.org/10.1016/j.neuron.2008.07.005>
- Raio, C. M., Carmel, D., Carrasco, M., & Phelps, E. A. (2012). Nonconscious fear is quickly acquired but swiftly forgotten. *Current Biology*, 22(12), R477–R479. <https://doi.org/10.1016/j.cub.2012.04.023>
- Reber, T. P., Samimizad, B., & Mormann, F. (2018). Cue discriminability predicts instrumental conditioning. *Consciousness and Cognition*, 61, 49–60.
<https://doi.org/10.1016/J.CONCOG.2018.03.013>
- Rebuschat, P. (2013). Measuring Implicit and Explicit Knowledge in Second Language Research. *Language Learning*, 63(3), 595–626. <https://doi.org/10.1111/lang.12010>
- Rosenthal, C. R., Kennard, C., & Soto, D. (2010). Visuospatial sequence learning without seeing. *PLoS ONE*, 5(7), e11906. <https://doi.org/10.1371/journal.pone.0011906>
- Rothkirch, M., & Hesselmann, G. (2017). What we talk about when we talk about unconscious processing - A plea for best practices. *Frontiers in Psychology*, 8(835), 1–6.
<https://doi.org/10.3389/fpsyg.2017.00835>
- Sand, A., & Nilsson, M. E. (2016). Subliminal or not? Comparing null-hypothesis and Bayesian methods for testing subliminal priming. *Consciousness and Cognition*, 44, 29–40.
<https://doi.org/10.1016/J.CONCOG.2016.06.012>
- Scott, R. B., Samaha, J., Chrisley, R., & Dienes, Z. (2018). Prevailing theories of consciousness are challenged by novel cross-modal associations acquired between subliminal stimuli. *Cognition*, 175, 169–185. <https://doi.org/10.1016/j.cognition.2018.02.008>
- Seitz, A. R., Kim, D., & Watanabe, T. (2009). Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron*, 61, 700–707. <https://doi.org/10.1016/j.neuron.2009.01.016>
- Seth, A. K., & Bayne, T. (2022). Theories of consciousness. *Nature Reviews Neuroscience*.
<https://doi.org/10.1038/S41583-022-00587-4>
- Seth, A. K., Verschure, P. F. M. J., Blanke, O., Butz, M. V., Ford, J. M., Frith, C. D., Jacob, P., Kyselo, M., McGann, M., Menary, R., Morsella, E., & Regan, J. K. O. (2016). Action-oriented understanding of consciousness and the structure of experience. In A. K. Engel, K. J. Friston, & D. Kragic (Eds.), *The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science* (pp. 261–281). MIT Press.
- Shanks, D. R. (2017). Regressive research: The pitfalls of post hoc data selection in the study of unconscious mental processes. *Psychonomic Bulletin & Review*, 24(3), 752–775.
<https://doi.org/10.3758/s13423-016-1170-y>
- Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioural and Brain Sciences*, 17, 367–447.
- Skora, L., Livermore, J. J. A., Nisini, F., & Scott, R. B. (2022). Awareness is required for autonomic performance monitoring in instrumental learning: Evidence from cardiac activity. *Psychophysiology*, e14047. <https://doi.org/10.1111/PSYP.14047>
- Skora, L., Scott, R. B., & Jocham, G. (2022). Stimulus awareness is necessary for both instrumental learning and instrumental responding to previously learned stimuli. *PsyArXiv*.
- Skora, L., Yeomans, M. R., Crombag, H. S., & Scott, R. B. (2021). Evidence that instrumental conditioning requires conscious awareness in humans. *Cognition*, 208, 104546.
<https://doi.org/10.1016/j.cognition.2020.104546>
- Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research*

- Methods, Instruments, & Computers*, 31(1), 137–149.
- Sterzer, P., Stein, T., Ludwig, K., Rothkirch, M., Hesselmann, G., Dubois, J., Yang, E., & Yuval-Greenberg, S. (2014). Neural processing of visual information under interocular suppression: a critical review. *Frontiers in Psychology*, 5(453), 1–12. <https://doi.org/10.3389/fpsyg.2014.00453>
- Timmermans, B., & Cleeremans, A. (2015). How can we measure awareness? An overview of current methods. In *Behavioral Methods in Consciousness Research* (pp. 21–46). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199688890.003.0003>
- Travers, E., Frith, C. D., & Shea, N. (2018). Learning rapidly about the relevance of visual cues requires conscious awareness. *Quarterly Journal of Experimental Psychology*, 71(8), 1698–1713. <https://doi.org/10.1080/17470218.2017.1373834>
- Vadillo, M. A., Konstantinidis, E., & Shanks, D. R. (2016). Underpowered samples, false negatives, and unconscious learning. *Psychonomic Bulletin & Review*, 23(1), 87–102. <https://doi.org/10.3758/s13423-015-0892-6>
- Van Den Bussche, E., Van Den Noortgate, W., & Reynvoet, B. (2009). Mechanisms of Masked Priming: A Meta-Analysis. *Psychological Bulletin*, 135, 452–477. <https://doi.org/10.1037/a0015329>
- van Gaal, S., de Lange, F. P., & Cohen, M. X. (2012). The role of consciousness in cognitive control and decision making. *Frontiers in Human Neuroscience*, 6(121), 1-. <https://doi.org/10.3389/fnhum.2012.00121>
- Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., Selker, R., Gronau, Q. F., Smíra, M., Epskamp, S., Matzke, D., Rouder, J. N., & Morey, R. D. (2017). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. In S. O. Lilienfeld & I. Waldman (Eds.), *Psychological Science Under Scrutiny: Recent Challenges and Proposed* (pp. 123–138). John Wiley & Sons. <https://doi.org/10.3758/s13423-017-1343-3>
- Wang, P., & Nikolić, D. (2011). An LCD monitor with sufficiently precise timing for research in vision. *Frontiers in Human Neuroscience*, 5(85), 1–10. <https://doi.org/10.3389/fnhum.2011.00085>
- Waroquier, L., Abadie, M., & Dienes, Z. (2020). Distinguishing the role of conscious and unconscious knowledge in evaluative conditioning. *Cognition*, 205, 104460. <https://doi.org/10.1016/j.cognition.2020.104460>

Supplementary materials:

Calculating the maximum potential influence of regression towards the mean resulting from post-hoc trial or participant exclusion.

Our paradigm proposes to use a trial-by-trial test of awareness in order to permit the exclusion of any trials where participants report having conscious knowledge of the stimuli. This approach has a number of key advantages over alternative approaches such as examining the mean accuracy in a separate test block. Firstly, it maximises the sensitivity of analysis by reducing unnecessary whole participant exclusions i.e. individual conscious trials can be excluded rather than all of a participant's responses excluded based on their mean accuracy. Secondly, the ability to exclude conscious trials in this way permits us to target an exposure duration that is as close as possible to the conscious threshold rather than having to deliberately aim substantially below that threshold in order to avoid the possibility of a small number of conscious trials contaminating the analysis. Finally, if a person's threshold is subject to variation over time, only a trial-by-trial analysis as proposed can reliably ensure that the effects of such momentary variations in conscious awareness are mitigated.

While employing a trial-by-trial awareness check has obvious advantages, any post-hoc exclusion of conscious trials or conscious subjects from the full sample brings a potential influence of regression to the mean (Shanks, 2017). In essence, while excluding the reported conscious trials will give the impression of leaving a set of purely unconscious trials for analysis, any noise in the measure of awareness (participants' reports of conscious perception) will mean that the selected subset may contain some proportion of conscious trials that were (due to measurement error) mislabelled as unconscious. Here we conduct an analysis that establishes the maximum potential influence that regression to the mean could have on our analyses for a given proportion of trials or participants excluded and a given error rate.

Our paradigm tests participants' ability to discriminate consciously whether a given shape is symmetrical or asymmetrical. For each symmetry-judgement, participants are required to report either that "I am guessing" or "I have some confidence". Any accuracy apparent in responses attributed to guessing is assumed to be unconscious and any apparent in responses attributed confidence is taken to be conscious. Because the awareness judgement is a two-alternative forced choice, the conscious plus unconscious attributions amount to 100% of responses. Crucially however, the observed proportions of each will differ from the true proportions according to the degree of error in the awareness report i.e. the number of trials where the true state is unconscious and yet reported with some confidence, or the true state is conscious but reported as guessing. For simplicity, we assume that there is a fixed amount of error in this measure and that it is the same for both conscious and unconscious attributions. It follows that the observed proportion of unconscious trials (those attributed to guessing) will be the proportion of trials that

are actually unconscious * the proportion of judgements without error + the proportion of trials that are conscious * the proportion of judgements made with error (misattributions). Note, that the maximum error is bounded by the proportion of responses attributed to each of the unconscious and conscious categories, e.g. if 90% of responses are attributed to unconscious judgement then the maximum error is 10% as this would account for all the conscious responses (10%). Adopting the following abbreviations, the formulae for observed proportion of unconscious trials can be written as below:

obsUC = The observed proportion of unconscious trials (trials attributed to guessing)

trueUC = The true proportion of unconscious trials

error = The proportion of judgments misattributed due to error

$$obsUC = (1 - error) * trueUC + error * (1 - trueUC)$$

This simplifies:

$$obsUC = trueUC + error * (1 - 2 * trueUC)$$

$$obsUC = error + trueUC * (1 - 2 * error)$$

From this, we can derive a formula for *trueUC*:

$$trueUC = \frac{(obsUC - error)}{(1 - 2 * error)}$$

And a formula for the proportion of conscious contamination of observed unconscious:

$$conscious\ contamination = \frac{error * (1 - trueUC)}{obsUC}$$

For example, if 90% of responses are attributed to unconscious responding (*obsUC* = 90%) and 10% to conscious it follows that the maximum error is 10% i.e. if 100% of responses are truly unconscious then 10% error would account for all the observed conscious responses. The true error can be anywhere between 0% and 10%. Applying the above formulae to each possible error rate (0-10%), we find that an error of 5% gives the maximum conscious contamination, namely 0.31%. That is, where 90% of observed trials are attributed to unconscious responding and all conscious trials are excluded the maximum extent to which regression to the mean could influence the observed learning is 0.31% of the conscious learning effect. In the conscious pilot, the *d'* for learning was 1.8. Hence, the maximum

effect of regression to the mean where 90% of responses are attributed to unconscious responding is $1.8 * 0.31\%$, $d' = 0.01$.

Supplementary table 1 shows the percentage of conscious contamination for each level of error for observed unconscious responding proportions of 95%, 90%, 85%, 80%, 75%, and 70%. Consulting this table and using the observed d' in the pilot study of 1.8 we find that the maximum influence from conscious knowledge where 80% of responses are attributed to unconscious responding is $d' = 0.03$. We feel that this is negligible and as such will adopt a strategy whereby provided the proportion of responses attributed to conscious responding does not exceed 20%, our exclusion criteria will be applied. Where participants make greater than 20% conscious attributions those individuals will be excluded as conscious. Precisely the same logic assessing the potential regression to the mean effect at the participant level applies. Hence, provided not more than 20% of participants make greater than 20% conscious attributions, we can be confident that the observed learning effect is not influenced by regression to the mean by more than a d' of 0.03. In the unlikely event that greater than 20% of participants make greater than 20% of conscious attributions then we will adopt an interval null in the Bayes analysis to account for this. This will modify the Bayes analysis only to the extent that the point null hypothesis will be replaced by an interval null modelled as a uniform from zero to the maximum influence of conscious contamination.

Table 1. Percentage of conscious contamination for each level of error for observed unconscious responding proportions of 95%, 90%, 85%, 80%, 75%, and 70%, computed with the formulae described in text. Frames highlight the maximum proportion of conscious contamination for the corresponding proportion of unconscious responding.

Observed unconscious	95%																				
error	0%	1%	2%	3%	4%	5%															
true percentage unconscious	95.0%	95.9%	96.9%	97.9%	98.9%	100.0%															
conscious contamination	0.00%	0.04%	0.07%	0.07%	0.05%	0.00%															
Observed unconscious	90%																				
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%										
true percentage unconscious	90.0%	90.8%	91.7%	92.6%	93.5%	94.4%	95.5%	96.5%	97.6%	98.8%	100.0%										
conscious contamination	0.00%	0.10%	0.19%	0.25%	0.29%	0.31%	0.30%	0.27%	0.21%	0.12%	0.00%										
Observed unconscious	85%																				
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%	13%	14%	15%					
true percentage unconscious	85.0%	85.7%	86.5%	87.2%	88.0%	88.9%	89.8%	90.7%	91.7%	92.7%	93.8%	94.9%	96.1%	97.3%	98.6%	100.0%					
conscious contamination	0.00%	0.17%	0.32%	0.45%	0.56%	0.65%	0.72%	0.77%	0.78%	0.77%	0.74%	0.66%	0.56%	0.41%	0.23%	0.00%					
Observed unconscious	80%																				
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%	13%	14%	15%	16%	17%	18%	19%	20%
true percentage unconscious	80.0%	80.6%	81.3%	81.9%	82.6%	83.3%	84.1%	84.9%	85.7%	86.6%	87.5%	88.5%	89.5%	90.5%	91.7%	92.9%	94.1%	95.5%	96.9%	98.4%	100.0%
conscious contamination	0.00%	0.24%	0.47%	0.68%	0.87%	1.04%	1.19%	1.32%	1.43%	1.51%	1.56%	1.59%	1.58%	1.54%	1.46%	1.34%	1.18%	0.97%	0.70%	0.38%	0.00%

Table 1. *continued*

Observed unconscious	75%																					
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%	13%	14%	15%	16%	17%	18%	19%	20%	
true percentage unconscious	75.0%	75.5%	76.0%	76.6%	77.2%	77.8%	78.4%	79.1%	79.8%	80.5%	81.3%	82.1%	82.9%	83.8%	84.7%	85.7%	86.8%	87.9%	89.1%	90.3%	91.7%	
conscious contamination	0.00%	0.33%	0.64%	0.94%	1.22%	1.48%	1.73%	1.95%	2.16%	2.34%	2.50%	2.63%	2.74%	2.81%	2.85%	2.86%	2.82%	2.75%	2.63%	2.45%	2.22%	
error	21%	22%	23%	24%	25%																	
true percentage unconscious	93.1%	94.6%	96.3%	98.1%	100.0%																	
conscious contamination	1.93%	1.57%	1.14%	0.62%	0.00%																	
Observed unconscious	70%																					
error	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%	13%	14%	15%	16%	17%	18%	19%	20%	
true percentage unconscious	70.0%	70.4%	70.8%	71.3%	71.7%	72.2%	72.7%	73.3%	73.8%	74.4%	75.0%	75.6%	76.3%	77.0%	77.8%	78.6%	79.4%	80.3%	81.3%	82.3%	83.3%	
conscious contamination	0.00%	0.42%	0.83%	1.23%	1.61%	1.98%	2.34%	2.67%	2.99%	3.29%	3.57%	3.83%	4.06%	4.27%	4.44%	4.59%	4.71%	4.78%	4.82%	4.82%	4.76%	
error	21%	22%	23%	24%	25%	26%	27%	28%	29%	30%												
true percentage unconscious	84.5%	85.7%	87.0%	88.5%	90.0%	91.7%	93.5%	95.5%	97.6%	100.0%												
conscious contamination	4.66%	4.49%	4.26%	3.96%	3.57%	3.10%	2.52%	1.82%	0.99%	0.00%												