Of two minds or one? A registered replication of Rydell et al. (2006)

Tobias Heycke¹, Sarah Gehrmann¹, Julia M. Haaf², & Christoph Stahl¹

¹ University of Cologne

² University of Missouri

To cite this article: Tobias Heycke, Sarah Gehrmann, Julia M. Haaf & Christoph Stahl (2018): Of two minds or one? A registered replication of Rydell et al. (2006), *Cognition and Emotion*, DOI: 10.1080/02699931.2018.1429389

Abstract

Evaluative conditioning (EC) is proposed as a mechanism of automatic preference acquisition in dual-process theories of attitudes (Gawronski & Bodenhausen, 2006; Rydell & McConnell, 2006). Evidence for the automaticity of EC comes from studies claiming EC effects for subliminally presented stimuli. An impression-formation study by Rydell and colleagues (2006) showed a selective influence of briefly presented primes on implicitly measured attitudes, whereas supraliminally presented behavioral information about the target person was reflected in explicit ratings. This finding is considered one of the strongest pieces of evidence for dual process theories (Sweldens, Corneille, & Yzerbyt, 2014), and it is therefore crucial to assess its reliability and robustness. The present study presents two registered replications of the Rydell et al. (2006) study. In contrast to the original findings, the implicit measures did not reflect the valence of the subliminal primes in both studies.

Keywords: evaluative conditioning, subliminal influence, implicit learning, replication

TH, SG, JH, & CS planned and conducted Experiment 1 and 2, analyzed the data, and wrote the article. Additional material and analysis scripts are provided at osf.io/c57sr. The study was funded by a grant awarded to the fourth author by Deutsche Forschungsgemeinschaft Grant 1269/3-1.

Correspondence concerning this article should be addressed to Tobias Heycke, Herbert-Lewin-Str. 2, 50931 Cologne, Germany, Tel. +49 221 4706339. E-mail: t.heycke@uni-koeln.de

We constantly rely on preferences to make evaluations and choices in our daily life: When browsing through online shopping sites, watching political debates, or when forming impressions of other persons. One way of acquiring preferences is through evaluative conditioning (EC). An EC effect is defined as a change in liking that is due to the pairing of a negative or positive stimulus (unconditioned stimulus, US) with a target object (conditioned stimulus, CS; De Houwer, 2007).

Hofmann, De Houwer, Perugini, Baeyens, and Crombez (2010) conducted a meta-analysis on EC and concluded that, “beyond any doubt[,] EC is a genuine phenomenon” (*p.* 410). Yet, the mechanisms that enable evaluative conditioning are not fully understood. Considerable progress has been made in this area, but there is still disagreement about the mental processes underlying evaluative conditioning (Sweldens, Corneille, & Yzerbyt, 2014). The current debate is between single- and dual-process views. Single-process views claim that there is only one path that leads to preference acquisition (De Houwer, 2009; Lovibond & Shanks, 2002; e.g., Mitchell, De Houwer, & Lovibond, 2009). According to single-process propositional views, evaluative conditioning results from the conscious formation of a proposition about the relation between a CS-US pair. Consequently, propositional views predict that awareness of the contingency between the CS and US would be necessary to establish an EC effect. Dual-process theories propose two distinct systems that may lead to preference acquisition: In addition to a conscious, propositional system, a second associative system is proposed in which a link between the stimuli can be formed automatically without contingency awareness. Thus, according to dual-process models, an EC effect could result from an associative system even when participants are unaware of the CS-US link. Influential examples for such accounts are the associative-propositional evaluation model (APE; Gawronski & Bodenhausen, 2006) and the systems of reasoning account (Rydell & McConnell, 2006). In addition to the distinction between propositional and associative systems, especially the latter model proposes a distinction between implicit and explicit attitudes, such that associations due to CS-US pairings are typically reflected in implicit evaluations, whereas propositions about CS-US relations typically form the basis of explicit attitudes (Gawronski & Bodenhausen, 2011; Rydell & McConnell, 2006).

One testable difference between dual- and single-process theories is therefore the question whether contingency awareness is a necessary condition for an EC effect. One way to experimentally assess the effect of contingency awareness on EC is via *subliminal* or suboptimal (i.e., brief and masked) presentation (Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006): If USs (or CSs) are presented suboptimally, they often do not enter conscious awareness, and the conscious formation of a proposition about the CS-US relationship is not possible (Dehaene et al., 2006). Following dual-process models, it should nevertheless be possible to form an unconscious association between CS and US through repeated pairing, because awareness of the relationship between stimuli is not needed for the formation of associations. These pairings would ultimately result in EC, which one should be able to detect with implicit evaluation measures (Gawronski & Bodenhausen, 2011). Finding EC with a subliminal conditioning procedure would be strong evidence for dual-process accounts.

Research on subliminal EC has so far yielded mixed results: On the one hand, there are several studies that found subliminal EC (e.g. De Houwer, Baeyens, & Eelen, 1994; Field & Moore, 2005; Fulcher & Hammerl, 2001; Houwer, Hendrickx, & Baeyens, 1997). Most of these studies, however had methodological issues (see Sweldens et al., 2014 for an overview).

On the other hand, numerous studies support the claim that awareness is necessary for EC (Dedonder, Corneille, Bertinchamps, & Yzerbyt, 2014; Pleyers, Corneille, Luminet, & Yzerbyt, 2007; Stahl, Haaf, & Corneille, 2016; Stahl, Unkelbach, & Corneille, 2009). Most of these studies, regardless of the results, have been criticized or disregarded by researchers in the field.

One study, however, is considered particularly convincing evidence for dual-process accounts (Sweldens et al., 2014). In this study, Rydell, McConnell, Mackie, and Strain (2006) showed a selective influence of briefly presented USs on implicit attitudes. In this experiment, the CS was a picture of a character named *Bob*, which was preceded by a briefly presented prime and followed by explicit, behavioral information about *Bob*. Prime and explicit information had opposite valence. The key finding of the study is that explicit attitudes were affected by behavioral information, whereas implicit attitudes reflected the valence of the briefly presented prime. Rydell et al. (2006) conclude that their participants held opposite implicit and explicit attitudes at the same time reflecting on the two types of information received. A prime-recognition check at the end of the study confirmed that the primes were not recognized above chance and Rydell et al. (2006) therefore concluded that the primes were indeed presented subliminally.

The finding is clearly interesting and theoretically relevant. The original study was cited more than 200 times and is generally viewed as methodologically strong (i.e., has fewer methodical issues than other studies claiming to have found subliminal EC; Dedonder et al., 2014). On the other hand, it is also based on a complex interaction obtained from a relatively small sample, inviting questions about its reliability. Furthermore, the study has not been independently replicated, as Sweldens et al. (2014) point out in their review (but see two similar studies by the original authors; Rydell & McConnell, 2006; Rydell, McConnell, & Mackie, 2008). Considering the important contribution to the discussion about EC without contingency awareness and dual-process theories in general (Gawronski & Bodenhausen, 2011), an independent replication would increase the confidence in the results, and thereby, the evidence for EC without contingency awareness.

In the present study, we conducted two replication attempts to investigate whether implicit attitudes can indeed form on the basis of briefly presented information (independently to contradicting explicit attitudes). Experiment 1 is a pre-registered replication using a German sample (see osf.io/cm2sj for the registration of Experiment 1 and osf.io/c57sr for the analysis scripts and material). Experiment 2 is a second pre-registered replication that aims at overcoming some of the potential limitations of Experiment 1: Specifically, to avoid potential cultural or language differences, we used the original material and sampled participants from a similar population as the original study. In addition, we pretested the prime presentation duration to ensure that the primes were indeed presented briefly enough so that they were not remembered at above-chance levels in the awareness check after the learning phase. The goal of both studies was to replicate the original study as closely as possible, while making sure that the experimental procedure would have the same psychological effects as the original procedure, and thus, would yield theoretically relevant results that allow for a valid test of the underlying theoretical claims by Rydell et al. (2006). On a continuum between exact/direct versus conceptual replications, therefore, Experiment 1 may be considered a conceptual replication (only) in the sense that the original material was translated into German (to make sure our German participants understood its meaning);

Experiment 2 may be considered a conceptual replication (only) in the sense that the original prime presentation duration was adapted to ensure chance-level prime awareness as reported in the original study.

Experiment 1

The goal of the first Experiment was to replicate the study of Rydell et al. (2006) as closely as possible. We contacted the first author and received the original material of the study. All instructions and the material was translated into German to conduct the study at the University of Cologne in Germany.

Method

Design. As in the original study, we realized a 2 (valence condition: half of the participants received negative primes and positive behavior in the first block and positive primes and negative behavior in the second block, while for the other half the blocks were the other way around) $\times$ 2 (times of measurement: after block 1 and after block 2) $\times$ 2 (dependent variable: implicit or explicit attitude) design. The last two factors were manipulated within participants, the first factor between participants.

Procedure. For efficiency reasons, we decided to conduct the experiment solely at the computer and therefore computerized the measures that had been administered as paper-pencil tasks in the original study. We used PsychoPy as the experimental software (Peirce, 2007). As in Rydell et al. (2006), we conducted a modified version of an attitude-learning paradigm developed by Kerpelman and Himmelfarb (1971).

Participants were seated at a computer and instructed that they would be learning something about a person (called either Bob, Jan, Tim, Tom or Ben). Only the name Bob was used in the original study. While Bob may be a common American name, it is not very common in German, so we included the above mentioned additional names that are more common in Germany. One of those five names was randomly assigned to the target character for each participant anew.

Each trial started with a blank screen for 1000 ms. Afterwards a fixation cross was displayed in the middle of the screen for 200 ms, which was replaced by either a positive or negative prime word in white letters on a black background, which was presented for 23.5 ms. The prime was presented for 23.5 ms instead of 25 ms (as in the original study) due to a slight difference in CRT monitor parameters (i.e., 85 Hz instead of presumably 80 Hz in the original study). We tried to mimic the exact prime presentation conditions as in the original study, based on the information given by the first author of the original study.

All in all we used 10 positive and 10 negative primes. Each prime was presented 10 times, resulting in 200 trials altogether. Half of the participants received 10 $\times$ 10 positive primes in the first block and 10 $\times$ 10 negative primes in the second block. The other half received negative primes first and positive primes in the second block. The order of the 100 primes within each block was randomized for each participant anew. The prime was immediately followed by an image of the target person, which was shown without any additional information for 247 ms.

After the target person was shown alone, negative or positive behavioral descriptions -written in white letters- appeared underneath the image, until the participant pressed

either the “c” key or the “u” key to indicate whether he/she thought that the behavior was characteristic or uncharacteristic for the target person. After the judgment, the participant received feedback for 5000 ms about whether their judgment was correct or incorrect (i.e., the corresponding German word was shown on the screen). In the first block of the learning phase, 50 positive behaviors and 50 negative behaviors were randomly drawn out of a list of 100 positive and 100 negative behaviors respectively, and presented in random order (determined for each participant anew). In the second block, the remaining 50 positive and 50 negative behaviors were presented in random order. The valence of the primes was always opposite to the valence of the characteristic behaviors (i.e., when a negative prime was shown, the negative behavior was uncharacteristic and the positive behavior characteristic of the target character).

After the first block, participants performed the first set of attitude measures, consisting of explicit evaluations of the target person and an IAT. We used the same measures as in the original study (but in a computerized version). The order of implicit and explicit measures was counterbalanced across participants, with the same order of measures after the first and second block. In the explicit attitude measure, participants were asked to judge how likable they found the target person, on a 9-point scale, ranging from 1 (very unlikable) to 9 (very likable). After the participant indicated his/her response, five 9-point semantic differential scales (Bad-Good, Mean-Pleasant, Disagreeable-Agreeable, Uncaring-Caring, Cruel-Kind) were presented together on a computer screen and participants were asked to judge the target person. Subsequently the participants were requested to evaluate the target person on a feeling thermometer ranging from 0° to 100°. All scores were z-standardized and a mean explicit rating score was calculated.

To assess implicit attitudes the participants may hold toward the target person, an implicit association test (IAT; Greenwald, McGhee, & Schwartz, 1998) was administered. As in Rydell et al. (2006), we used a modified IAT version reported by McConnell and Leibold (2001) consisting of seven 20-trial blocks. Participants had to press the left key (“d”) or right key (“k”) as part of a classification task. In the first block, images of the target character as well as images of other persons had to be categorized as depicting the target character or other persons (counterbalanced over participants whether the target character was on the left or the right key). The target character name and “not” plus the target character name were displayed on the left or right bottom of the screen as category labels. As in the original study we selected one image—out of the six images in the original material at random, for each participant anew, to represent the target character (in the learning phase and in the IAT). The remaining five other images were then used in the IAT to represent other individuals that were not the target character. In the second block, positive and negative words had to be categorized as positive or negative (with the position of the corresponding key for positive/negative counterbalanced across participants). Now the labels “positive” and “negative” were displayed on the left and right bottom of the screen. In the third and fourth block, the first and second block were combined: Half of the participants had the target character together with negative words on one key and all other images and positive words on the other key (Combination 1), while the other half of participants classified the target character together with positive words on one key and all other images and negative words on the other key (Combination 2). The category labels from block 1 and block 2 were combined in this block (i.e., Bob + positive). In the fifth block, a classification of

the target character vs. other persons was again performed, but the location of the key for classifying the target character (vs. other persons) was switched compared to block 1 and block 3/4. In the final two blocks (i.e., blocks six and seven), participants again worked on a combined classification task (i.e., a combination of blocks 2 and 5): Participants who worked on Combination 1 in blocks 3 and 4 now worked on Combination 2 (and vice versa). The order of trials in the critical blocks 3/4 and 6/7 was randomized for each participant anew, controlling that there were no more than 5 repetition trials (same key pressed as in previous trial) and controlling that there were approximately the same amount of response-switch (switching between the d and k key) and response-repetition trials.

We processed the reaction times (RT) of the critical blocks 3/4 and 6/7 in the IAT as described in the original study (Rydell et al., 2006). All trials were used and responses faster than 300 ms recoded to 300 ms and responses slower than 3000 ms recoded to 3000 ms. Reaction times were transformed to reduce skewness (natural log) and a mean log reaction time was calculated for each participant in each critical block. An IAT effect was calculated subtracting the mean log RT of the block in which the target character and positive words (Combination 2) were presented from the mean log RT of the block in which the target character and negative words (Combination 1) were presented. The IAT effect was then z-standardized (across both IAT measures and all participants). Therefore a larger IAT effect indicated a more positive evaluation of the target character.

After the second experimental block and the second set of measures, a recognition test of the briefly presented primes was administered: 40 words were presented in a random order on the computer screen simultaneously (20 of them were the previously presented primes, the other 20 were words that had not been shown in any task previously). Participants were asked to select those 20 words which they thought had been shown as primes during the learning phase. At the end of the study, we asked demographic questions (Age, Study/Job, Gender, Suspected goal of study, Comments).

Materials. Two individuals translated the instructions and materials independently and a third person adjusted the translations, if necessary. The original material we received contained eight positively and eight negatively valenced words used in the IAT, but the article referred to ten positive and ten negative words.¹ We therefore added the four words “successful” (erfolgreich), “talented” (begabt), “hectic” (hektisch) and “disgusting” (widerlich) to complete the required ten items for each category.² We also noticed that some valenced IAT targets were quite similar to the briefly presented primes during the learning phase. The word „laughter“ (Lachen), which translated into German appears very similar to the prime word „smile“ (Lächeln) and the word „friend“ (Freund), which was used as a prime, were replaced in the IAT with „trust“ (Vertrauen) and „humor“ (Humor).³

In our experiment we addressed a potential issue with the original prime-recognition check: The words “help“, “candy“, “crash[ed]“ and “gun“ occurred as new words in the recognition task but also as behaviors that describe the character “Bob.”⁴ Participants

¹After the data of Experiment 1 was collected we were informed that a different set of target words was used in the original study.

²The first three words were taken from a study by Bluemke and Friese (2006), while the last word was explicitly mentioned in the original article by Rydell et al. (2006).

³Neither “laughter” nor “friend” were used as targets in the original Experiment, but were included in the material we initially received.

⁴Additionally the words “Love” and “Peace“ were used as new words in the recognition test, but were

might have selected those words because they remembered that they saw those words before, without remembering their source (i.e., whether they were primes or words in a different task, such as the IAT, or presented as part of the sentences describing Bob). If those words were selected frequently it might have resulted in an underestimation of participants' knowledge of the primes. The conclusion that participants were not able to recognize the primes during learning must therefore be interpreted with caution.

To eliminate potential issues regarding these words, we substituted the word “help” (Hilfe) by the word “beautiful” (schön), “crash” was replaced by “fear” (Angst) and “gun” was replaced by “pistol” (Pistole) in the recognition task. “Candy” was replaced in the behaviors with “magazines” (Zeitschriften).

In a small pilot study ($N = 16$) each valent word was rated on a 101 point likert scale, ranging from -50 to 50. The results confirmed that all translated primes ($M_{pos} = 78.17$; $M_{neg} = 21.34$), new words in the prime-recognition task ($M_{pos} = 74.04$; $M_{neg} = 22.14$), and words used in the IAT ($M_{pos} = 83.54$; $M_{neg} = 21.31$) were appropriately positive (or negative) in the German language.

As Rydell et al. (2006), we also made sure that primes and new words in the recognition task were on average of approximately equal length and equal frequency in the German language (Primes: $M_{length} = 6.1$, $M_{frequency} = 12460.35$; New words: $M_{length} = 7.2$, $M_{frequency} = 11778.9$).

Changes compared to original Experiment. In addition to computerizing all parts of the study and translating the material (thereby adding multiple different German names for the target character) we resolved minor inconsistencies between the paper and the material we received: (1) We received 16 IAT attribute words (8 positive, 8 negative) and added 4 words (2 positive, 2 negative) for a total of 20 as mentioned in the method section of the original study. Of the positive IAT attribute words, two were substituted by new words because one was very similar to a prime from the learning phase after the translation to German and another was used as a prime in the learning phase. (2) Beyond the translation and resolving the inconsistency between the paper and the material received, the only adjustment was that we replaced distracter words from the recognition test that had in fact been part of the IAT and/or behavioral descriptions (and were therefore not new to participants) with truly new words, so as to obtain a valid estimation of prime recognition. As the memory measure was administered in the very end of the experiment, we can be confident that this change cannot influence the main pattern of results obtained in the IAT and evaluative ratings.

Hypotheses.

Explicit Evaluation. Under both single-propositional and dual-process theories we would expect a more favorable explicit rating of the target stimulus (i.e., Bob) when positive behavior was characteristic of him than when negative behavior was characteristic of him (Rydell et al., 2006). We thus expect a significant difference between the results of the explicit attitude measure at time 1 compared to time 2, in the direction described above. Because half of the participants received positive behavioral information in the first block and negative information in the second block, whereas the order was reversed for the other

also previously used as valent words in the IAT after block 1 and block 2. To avoid potential problems in the prime recognition check, we substituted the words “love” (Liebe) and “peace” (Frieden) with “nature” (Natur) and “cash” (Bargeld) in the prime-recognition task.

half, we expect a 2 (time of measurement) $\times 2$ (valence condition: negative primes and positive behavior first vs. positive primes and negative behavior first) interaction.

IAT effect. Based on the findings by Rydell et al. (2006), and as predicted by a dual-process account, we would expect that the implicit measure contradicts the explicit measure in each experimental block. Specifically, we would expect an IAT effect to reflect the valence of primes (and opposite to the behavior information), resulting in a significant difference between the results of the implicit attitude measure at time 1 compared to time 2.

A single-process view, in contrast, would predict that a subliminal prime has no influence on the implicit measure. We would therefore expect that the IAT reflects the supraliminal behavioral information about the target character as does the explicit rating.

Both of these opposing patterns would be reflected – similar to the explicit evaluation – in a 2×2 interaction of time of measurement and the valence condition (negative primes and positive behavior first vs. positive primes and negative behavior first).

Analysis Plan

The data were analyzed in two ways. First, the same frequentist analyses were conducted as in the original article. They were based on the initial fixed sample size for which the underlying sampling assumptions hold. Second, we planned to continue sampling, while sequentially computing Bayes factors, until the relative evidence was strong enough for either the null or alternative hypotheses. For the frequentist analysis we used $\alpha = .05$ as inference criterion. In all Bayes Factor ANOVAs we used the following method to estimate whether a given factor or interaction should be interpreted as having an influence on the dependent variable: The full model with all effects (i.e., main effects and interactions) was compared to a model without the effect of interest (i.e., a main effect or an interaction). A Bayes factor for the alternative hypothesis over the null hypothesis (i.e., $BF_{10} > 1$) for a given effect indicates that the full model (that includes the predictor term corresponding to the effect) is favored over the model without the effect given the data (i.e., the predictor term improved the model, suggesting that the effect had an influence on the dependent variable). We follow Rouder and Morey (2012) and interpret the Bayes factor as the evidence for one model (in this case the full model) *relative* to a competing model (in this case the restricted model).

In model comparison using Bayes factors, prior distributions on target parameters represent the hypotheses to be tested. We chose weakly informative priors for the analysis here, as recommended by Rouder, Morey, Speckman, and Province (2012). Previous EC studies on contingency awareness using the IAT found small to medium effect sizes (Rydell et al., 2006: $\eta^2 = .1$; Olson & Fazio, 2001: $d = 0.37$). We used these previous findings as guideline for our expectation for this study and chose our priors accordingly. For our ANOVA analyses we therefore used default multivariate Cauchy priors on the effect size parameters (Rouder et al., 2012), with a scaling parameter of $r = 0.5$. For all Bayesian t tests Cauchy priors with a scaling parameter of $r = \sqrt{2}/2$ were used (Rouder et al., 2012; Rouder, Speckman, Sun, Morey, & Iverson, 2009), which expresses our expectation to find small to medium effects. When informative prior distributions are used, the question can be asked how the chosen priors affect the results. To test the robustness of the results to changes of the prior we analyzed the data with other priors around the chosen scaling factor. We used smaller and larger prior scaling factors of 0.3 and $\sqrt{2}/2$, respectively, for

the ANOVA. The pattern of Bayes factors did not change substantially (all BF that were >1000 were still >1000 and all BF > 10 were still larger 10). We additionally calculated all t tests with prior scales of 0.5 and 1: all Bayes Factors originally reported to be larger 1000 were still larger than 1000, all Bayes Factors that were initially larger than 10 were still larger than 10 when changing the prior width; and all Bayes Factors that were smaller than 10 remained smaller than 10.⁵

We additionally report the median and the 95% highest density interval (HDI) of the posterior distribution of the estimated effect size d . The 95% HDI contains the true population effect size with a probability of 95%. We considered a Bayes factor > 10 or $< 1/10$ to be convincing evidence for the alternative hypothesis relative to the null, or for the null hypothesis relative to the alternative, respectively. We use BF_{10} to indicate a Bayes factor for the alternative Hypothesis and BF_{01} as a Bayes factor for the null hypothesis (dividing 1 by the calculated Bayes factor). BF_{10} or BF_{01} smaller than 3 will only be considered to be anecdotal relative evidence (Schönbrodt, Wagenmakers, Zehetleitner, & Perugini, 2015). We performed all analyses in R (Version 3.4.1; R Core Team, 2016) and the R-packages *afex* (Version 0.18.0; Singmann, Bolker, Westfall, & Aust, 2016), *BayesFactor* (Version 0.9.12.2; Morey & Rouder, 2015), *dplyr* (Version 0.7.2; Wickham & Francois, 2016), and *papaja* (Version 0.1.0.9492; Aust & Barth, 2016).

Sample size rationale. For the initial frequentist analyses, we planned to collect the same number of participants as in the original study ($N = 50$), even though an a priori power analysis with the smallest interaction effect reported in the original study ($f = 0.32$) and $\alpha = \beta = .05$ yielded a required sample of 34 participants. Since we did not want to sample fewer participants than in the original study, an experiment with $N = 50$, would allow us to detect an interaction effect of $f = 0.26$ or larger (with $\alpha = \beta = .05$). Starting after reaching this initial sample size (thereby lowering the possibility of a false positive finding in the Bayesian analyses; Schönbrodt et al., 2015), we planned to continue the data collection and computed Bayes factors. We planned to stop data collection once we had gathered sufficient information (BF_{10} or BF_{01} larger 10) regarding our hypotheses of main interest (namely the two-way interactions, for both dependent measures, of time of measurement $\times$ order of information received) or once we had reached a maximum of 125 participants (which is 2.5 times the original sample size as recommended by Simonsohn, 2015 for replication studies).

Participants. We initially collected the data of 53 participants and - based on our Bayesian analyses - did not collect any additional data. One participant aborted the experiment and one data file was corrupted; both of these events were known during data collection, and these missing participants were replaced by 2 additional participants. We used the data of 51 participants in the final analysis.⁶

⁵See table robustnessBF.pdf on osf.io/c57sr for a detailed overview of all Bayes Factors with different prior scaling factors.

⁶In addition, during the final day of data collection, two participants instead of only one were collected accidentally, resulting in 51 (instead of the planned 50) complete data sets after testing a total of 53 participants. Results do not change when we excluded the 51st participant (largest deviation of $d = 0.05$; largest deviation of $\eta_G^2 = .005$). Two participants had to rate the target character also before the first learning phase, due to a programming error. Excluding these participants did not substantially alter the results (largest deviation of $d = 0.07$; largest deviation of $\eta_G^2 = .01$).

Confirmatory Analysis

We first report the set of planned and pre-registered analyses to test our hypotheses, before we turn to additional exploratory analyses.

Recognition Test. A one-tailed t test indicated that the primes were recognized above chance in the recognition task at the end of the study, $M = 12.53$, with chance level at 10, $t(50) = 5.22$, $p < .001$; $BF_{10} = 10,227.03$, $d = 0.70$, 95% HDI [0.39, 1.00].

Explicit and Implicit Evaluation. As in the analysis by Rydell et al. (2006) we first ran a $2 \times 2 \times 2$ ANOVA with the type of the dependent variable (IAT, Explicit), the time of measurement (time 1, time 2) and the valence condition (negative primes and positive behavior first vs. positive primes and negative behavior first) with the first two factors manipulated within participants. As in the original study, we found statistical evidence for a three way interaction of the variables mentioned above, $F(1, 49) = 49.16$, $MSE = 0.36$, $p < .001$, $\hat{\eta}_G^2 = .149$, $BF_{10} = 1.14 \times 10^7$.

Explicit Evaluation. We subsequently tested for a 2 (time of measurement) $\times$ 2 (valence condition) ANOVA in the explicit rating data only. As in the original study, we indeed found a significant interaction of the time of measurement and whether positive behavior was presented first and negative second or vice versa, $F(1, 49) = 274.76$, $MSE = 0.26$, $p < .001$, $\hat{\eta}_G^2 = .757$, $BF_{10} = 4.07 \times 10^{30}$.

As can be seen in Figure 1 and Table A1, the rating of the character was according to the explicit information given to the participants. When given positive information first, the rating of the target character was more positive than when given negative information, $t(27) = 11.52$, $p < .001$; $BF_{10} = 1.37 \times 10^9$, $d = 2.09$, 95% HDI [1.41, 2.79]. Similar, when negative behavior was learned in the first block, the rating of the target was more negative than when positive information was learned in the second block, $t(22) = 12.14$, $p < .001$; $BF_{10} = 2.8 \times 10^8$, $d = 2.41$, 95% HDI [1.57, 3.28].

Implicit Evaluation. The analysis of the IAT data also revealed a significant 2-way interaction of time of measurement and condition, $F(1, 49) = 14.24$, $MSE = 0.44$, $p < .001$, $\hat{\eta}_G^2 = .075$, $BF_{10} = 66.31$.

As can be seen in Figure 2 and Table A2, when negative primes were presented first and positive primes afterwards, we can observe some indications that the IAT effect became more negative $t(27) = -2.54$, $p = .017$; $BF_{10} = 2.92$, $d = -0.44$, 95% HDI [-0.82, -0.06]. This pattern is inconsistent with the prime valence but consistent both with the verbal behaviors presented to the participants, as well as the explicit ratings. Similarly, when positive primes were presented first and negative primes afterwards, the IAT effect became more positive $t(22) = -2.77$, $p = .011$; $BF_{10} = 4.50$, $d = -0.52$, 95% HDI [-0.95, -0.09]. This pattern is again consistent with the explicit evaluation and inconsistent with the notion that the IAT effect reflects prime valence.

Exploratory analyses. Our counterbalancing factors had no influence on the pattern of results in the implicit nor explicit evaluations (but we cannot confidently conclude an absence of such effects because the Bayes factors were mostly not conclusive). Additionally, no correlation between the number of remembered primes and the IAT effect score was found, $r = .15$, 95% CI [-0.13, .41], $t(49) = 1.05$, $p = .301$, $BF_{01} = 2.28$.

Composite effect. We additionally computed a composite IAT effect score for both groups together (i.e., negative primes first and positive primes first) and compared it to zero.

Figure 1. Standardized explicit evaluation effects of Experiment 1, split by the valence of the explicit information and the time of the explicit measurement. Error bars represent 95% within-subject confidence intervals. Pos: positive information, neg: negative information.

We subtracted the mean log IAT effect when the target character was shown with a negative prime from the mean log IAT effect when the character was shown with a positive prime.

Values larger 0 would therefore indicate that the target character’s IAT score was in line with the valence of the briefly presented primes; values smaller than 0 would indicate that it was in line with the behavioral information received. We observed a negative value for the composite score ($M = -0.08$), and a one-sided t test of the hypothesis derived from the results of Rydell and colleagues (2006) that the effect is *larger* than zero was consequently not significant, $t(50) = -3.78$, $p > .999$. The Bayes factor yielded relative evidence for the null hypothesis over the alternative hypothesis, $BF_{01} = 28.58$, $d = -0.50$, 95% HDI $[-0.79, -0.21]$, which clearly indicates that the target character was not implicitly evaluated in line with the valence of the briefly presented primes.

Similarly, we computed a difference score for the explicit rating. We subtracted the z-standardized mean explicit rating when negative behavior was learned from the mean rating when positive behavior was learned. A positive value therefore indicated an explicit rating in line with the behavioral information, while a negative value would indicate an explicit rating in line with the briefly presented primes. Unsurprisingly, a one-sided t test revealed that the target character was rated in line with the behavioral information received, $t(50) = 16.74$, $p < .001$; $BF_{10} = 2.16 \times 10^{19}$, $d = 2.29$, 95% HDI $[1.76, 2.83]$.

Figure 2. Standardized IAT effects of Experiment 1, split by the valence of the primes and the time of IAT measurement. Error bars represent 95% within-subject confidence intervals. Pos: positive prime, neg: negative prime.

Discussion Experiment 1

We replicated the pattern of results of the original study for the explicit evaluation. When participants learned that the target character was a good person, he was evaluated more positively than when participants learned that he was a bad person. Considering that participants had 100 trials to learn about the target character in each block, we would consider this is a trivial learning effect consistent with both single- and dual-process accounts.

Contrary to the original study, we did not replicate the pattern of effects for the implicit measures. While the IAT effects in the original study reflected a change in line with the briefly presented primes, the pattern observed in this experiment was opposite to it (i.e., in line with the explicit information about the target character). We therefore conclude that we did not find evidence that the briefly presented primes had an effect on implicit evaluations.

Limitations of Experiment 1. There are two aspects that might account for the difference in result patterns between the original study and the present replication: First, the translation of the material into German may have affected the results. For instance, the German prime words were slightly longer (6.05 letters) than the original primes (4.9 letters), so one might speculate that the primes could not be fully processed in the brief presentation

time. However, this interpretation does not fit with the prime-recognition result, which showed a higher rate of prime recognition than in the original study. We would therefore conclude that prime processing was not reduced in this replication. However, other language or cultural factors may have moderated the influence of the briefly presented primes.

Second, perhaps more importantly, our results suggest that prime processing was not fully subliminal. As in the study by Rydell et al. (2006), we measured the recognition of the briefly presented primes with a recognition test in the end of the study. Recall that one half of the primes were shown in the first block only, the other half in the second block only. Since participants additionally received no information about the primes, the prime-recognition test in the end of the study is a weak proxy for the actual visibility of the primes during the learning phases (i.e., decay over time and interference from subsequent stimuli may have weakened prime memory). We additionally substituted multiple words from the list of new words used in the original study with words that were not previously used (see above). We found – contrary to the original study – that the primes were recognized above chance. We do not know whether this difference is purely due to a more strict measure of memory (i.e., due to our using truly new words as distracters in the recognition list given to the participants, instead of including words that were also used in the behavioral statements), or due to a higher visibility of the primes. Note that we aimed at closely replicating the presentation conditions of the original study (black background, white letters, font type and size), used the same images as backward masks, and even realized a slightly shorter prime presentation time compared to the original study (i.e., 23.5 ms instead of 25 ms). Together with the slightly increased word length, it could be argued that, if anything, it should have been more difficult to perceive the briefly presented primes. However, in contrast to the original study, the result of the prime-recognition check suggested that the primes were not subliminal, but were supraliminally processed (at least in part).

Note that prime presentation may also have been supraliminal in the original study: The differences in prime-recognition performance may be explained by the improvement made to the prime recognition test (i.e., by using truly new words). However, we cannot rule out the possibility that the primes were too easy to read in the present Experiment 1, and were therefore processed in a propositional fashion instead of the automatic associative process. Following that logic, it would not be surprising that we did not observe an effect of the primes on the IAT, considering both primes and verbal information were supraliminal: The clearly visible behavioral information participants learned about the target person, should have a much larger effect on preferences than a briefly presented word with no clear relation to the target person. In order to allow for automatic associative effects on implicit attitudes, it might be necessary to ensure that prime processing is truly subliminal in the sense that prime recognition is at chance as reported in the original study. Keeping the original presentation duration would therefore lead to functionally different presentation conditions (i.e., above-chance awareness as opposed to chance-level awareness in the original study), and this would render the results of the replication study all but uninterpretable. In order to allow for a valid interpretation of the results of our replication attempt, we therefore aimed to realize presentation conditions that ensured chance-level awareness of the primes (see Stroebe & Strack, 2014).

Experiment 2

The goal of the second Experiment was to conduct a second replication of the original study, addressing the potential limitations outlined above. We used the same procedure as in Experiment 1, but used the original materials and a sample of participants from the same population of US undergraduate students. Primes in Experiment 2 were presented for the longest possible duration for which prime memory (measured as in the original study) did not exceed chance level in a pretest.

Pretest

Participants completed a single learning block with 10×10 positive (negative) primes, receiving the same instructions as in Experiment 2. They were then asked about their explicit evaluations of the target person (this was done to erase any prime knowledge from working memory) before completing an adapted version of the prime-recognition task that contained only the positive (negative) words. We successively tested prime presentation durations starting with 23.5 ms (continuing with 20 ms, 17 ms, 13 ms etc.), until the primes could no longer be identified above chance. We used sequential Bayesian analysis (with the same priors and specifications as in Experiment 1, starting each analysis after 10 participants) and stopped data collection for any given presentation time as soon as the Bayes Factor for the test of prime memory performance against chance level has reached $BF_{10} > 5$ or $BF_{01} > 5$ (with a maximum N of 25 participants per condition). We decided a priori to use the longest presentation time for which prime identification did not exceed chance level (i.e., until we obtain a Bayes factor in support of the null hypothesis).⁷

With a prime presentation time of 23.5 ms, the primes were visible above chance, $BF_{10} = 5.41$, $d = 0.48$, 95% HDI [0.05, 0.92], $N = 22$. This also tended to be the case with a presentation time of 20 ms, $BF_{10} = 2.16$, $d = 0.36$, 95% HDI [-0.04, 0.76], $N = 24$. A test of a prime presentation duration of 17 ms yielded convincing evidence that prime stimuli were remembered at chance level, $BF_{01} = 5.78$, $d = -0.08$, 95% HDI [-0.48, 0.32], $N = 21$.

Procedure

We used the same experimental scripts as in Experiment 1, but with the original English instructions and stimulus material (including only the name Bob for the target character). In contrast to Experiment 1, we used the same IAT target words as in the original study. As in Experiment 1, we used new words (i.e., words that were not used during the learning phase) for the recognition task, to ensure a more exact approximation of the visibility of the primes.⁸ Most importantly, we changed the presentation time of the prime to 17 ms.

⁷The data of the pretest and Experiment 2 were collected under a Born Open Data protocol (Rouder, 2016) in which they were automatically logged, uploaded, and made freely available as they were created (<https://github.com/PerceptionCognitionLab/data1/tree/master/repRydell>).

⁸In the original study, some of the prime words, as well as some of the new words in the prime recognition task, also appeared as part of the behavioral statements during the learning phase. We slightly adjusted the behavioral statements (e.g., substituted single words in the behavioral description with similar but different words or rearranged the sentence and omitted the word in question) to remove this possible confound. Two independent raters confirmed that the behavioral statements were still clearly identifiable as positive or negative.

Analysis plan

As in Experiment 1, we conducted frequentist and Bayesian analyses after a fixed initial sample size and used sequential Bayesian testing for all further analyses. We used the same annotation, inference criteria, and priors as in Experiment 1.

Sample size rationale. As in Experiment 1, we initially collected the data of 50 participants. We additionally decided to continue recruiting participants using sequential Bayesian analysis (analyzing the data after every day of data collection) until the two Bayes factors of interest (see below) were larger than 10 (or smaller than $1/10$), or until a maximum of 125 participants was reached. No new participants would be recruited after reaching this criterion, but those already signed up at this point were allowed to take part in the study. We decided to exclude participants who (a) aborted the experiment or (b) experienced or reported unforeseeable circumstances (e.g., noises, difficulty with the instructions; based on our experience we did not expect any of those cases to occur).

Data collection stopping rule and Confirmatory Analysis. As in the exploratory analyses of Experiment 1 we planned to calculate a composite effect both for the IAT and the explicit evaluation. We planned to run one-sided t tests and compare the score to 0 (see above) to determine whether (a) the explicit evaluation is in line with the behavioral information and (b) whether the IAT effects are in line with the prime valence. We planned to stop data collection once we have gathered sufficient information on these two contrasts (i.e., $BF_{10} > 10$ or $BF_{01} > 10$).

Additional Analyses. We planned to additionally conduct the same confirmatory analyses as in Experiment 1, but without using these as stopping rules for the data collection.

Participants. We collected the data of 57 participants and—based on our Bayesian analyses—did not collect any additional participants. We collected more than the planned 50 participants, as we allowed participants who had already signed up to participate as well (see above).

Confirmatory Analysis

We will first report on the performance of the participants in the prime recognition test and compare it to the chance level, then analyse the composite effect score testing the two crucial hypotheses (i.e., explicit rating in line with the behavioral information and IAT effect in line with the valence of the primes), and we will additionally report the same confirmatory analyses as in Experiment 1 and the original study.

Recognition test. A one-tailed t test indicated that the primes were not recognized above chance in the recognition task at the end of the study, $M = 10.07$, with chance level at 10, $t(56) = 0.46$, $p = .325$; $BF_{01} = 4.66$, $d = 0.06$, 95% HDI $[-0.19, 0.31]$. This was contrary to the results of the first Experiment but in line with the results of the original study.

Stopping rule analysis: Composite effect. We again computed a composite effect score for the explicit ratings and the IAT data. Values larger than zero in the explicit rating indicated ratings in line with the valence learned from the behavioral statements, values larger than zero from the IAT data indicated IAT effects in line with the valence of the briefly presented primes.

Unsurprisingly, a one-sided t test of the explicit ratings revealed that the target character was rated in line with the behavioral information received, $t(56) = 20.31$, $p < .001$; $BF_{10} = 3.83 \times 10^{24}$, $d = 2.64$, 95% HDI $[2.09, 3.21]$.

Importantly, as in Experiment 1 we found a negative value for the composite score of the IAT ($M = -0.11$), indicating that the IAT effect was not in line with the valence of the primes. The pre-registered one-sided t test of the composite score of the IAT data, testing whether the effect is *larger* than zero, was consequently not significant, $t(56) = -4.41$, $p > .999$. The Bayesian analysis also yielded relative evidence for the absence of the predicted effect, $BF_{01} = 33.85$, $d = -0.56$, 95% HDI $[-0.84, -0.28]$.

Overall confirmatory analysis. As in Experiment 1 and in the original analysis we additionally ran a $2 \times 2 \times 2$ ANOVA with the type of the dependent variable (IAT, Explicit), the time of measurement (after block 1, after block 2) and the valence condition (negative primes and positive behavior first vs. positive primes and negative behavior first) with the first two factors manipulated within participants. We again found statistical evidence for a three way interaction of the variables mentioned above, $F(1, 55) = 54.69$, $MSE = 0.35$, $p < .001$, $\hat{\eta}_G^2 = .149$, $BF_{10} = 3.74 \times 10^7$.

Figure 3. Standardized explicit evaluation effects of Experiment 2, split by the valence of the explicit information and the time of the explicit measurement. Error bars represent 95% within-subject confidence intervals. Pos: positive information, neg: negative information.

Explicit Evaluation. We therefore computed a 2 (time of measurement) $\times 2$ (valence condition) ANOVA of the explicit rating data only. A significant interaction of the time of measurement and whether positive behavior was presented first and negative second or vice versa was found, $F(1, 55) = 403.79$, $MSE = 0.22$, $p < .001$, $\hat{\eta}_G^2 = .828$, $BF_{10} = 1.83 \times 10^{42}$.

As in the previous studies, the rating of Bob was according to the explicit information given to the participants (see Figure 3 and Table B1). When learning that Bob was a good person in the first block, the rating of the target character was more positive than when learning Bob was a bad person, $t(30) = 15.18$, $p < .001$; $BF_{10} = 4.66 \times 10^{12}$, $d = 2.63$, 95% HDI [1.86, 3.41]. Similar, when negative behavior was learned about Bob in the first block, the explicit rating of Bob was more negative than when positive information was learned, $t(25) = 13.40$, $p < .001$; $BF_{10} = 1.16 \times 10^{10}$, $d = 2.52$, 95% HDI [1.70, 3.35].

These results again show a simple learning effect, but also demonstrate that the participants paid attention to the instructions and the information presented.

Figure 4. Standardized IAT effects of Experiment 2, split by the valence of the primes and the time of IAT measurement. Error bars represent 95% within-subject confidence intervals. Pos: positive prime, neg: negative prime.

Implicit Evaluation. An ANOVA of the IAT data also revealed a significant 2-way interaction of time of measurement and valence condition, $F(1, 55) = 18.96$, $MSE = 0.54$, $p < .001$, $\hat{\eta}_G^2 = .102$, $BF_{10} = 535.23$.⁹

As shown in Figure 4 and Table B1, when negative primes were presented first and positive primes afterwards, the IAT effect became more negative $t(30) = -4.77$, $p < .001$; $BF_{10} = 529.65$, $d = -0.80$, 95% HDI [-1.21, -0.39]. This pattern is inconsistent with the prime

⁹Additionally, we found no indications for an influence of the order of dependent measures or the key assignment during the IAT on the IAT results (as in Experiment 1 the Bayes factors were mostly not conclusive).

valence but consistent with the behavior descriptions presented to the participants, with the explicit ratings, as well as with the results of Experiment 1. Even though descriptively true, and contrary to Experiment 1, there was no statistical evidence that the IAT effect became more positive when negative primes were shown after positive primes, $t(25) = -1.45$, $p = .161$; $BF_{01} = 1.91$, $d = -0.25$, 95% HDI [-0.63, 0.12]. One should note, however, that we ran two tailed t tests in these analyses; the direction of the effect was, for both valence order conditions, not in line with the prime valence but instead in line with the valence of the explicitly learned information. Taken together, these analyses—similar to the composite effect analyses—showed that there was no evidence for the pattern described in the original study (i.e., no evidence for IAT effects in line with the valence of the briefly presented primes).

Discussion Experiment 2

In Experiment 2, we again set out to test whether implicit attitudes could be selectively influenced by briefly presented primes, while explicit attitudes were influenced by explicitly learned behavioral information. In contrast to Experiment 1, we used the original instructions and prime words, and collected data from a similar population as in the original study (i.e., US undergraduate students). We also used a shorter presentation time of the briefly presented primes (i.e., 17 ms compared to 25 ms). We speculated that we did not find an effect of the primes in Experiment 1 because they were sometimes visible and therefore may have been processed consciously, thus competing with the explicit statements instead of affecting implicit attitudes. In Experiment 2, we found that primes were not recognized above chance in the prime recognition task. This result is contrary to the above-chance recognition in Experiment 1, but in line with the original results. If the primes in Experiment 1 indeed did not affect the IAT effect because they competed with explicit information, we should be able to detect an effect of the primes on the IAT in Experiment 2, where primes were not processed consciously. In a composite effects analysis, however, we found compelling evidence that the IAT effect was not in line with the valence of the briefly presented primes. All analyses of the explicit rating support the previous findings that the explicit ratings were in line with the valence learned about Bob throughout 100 trials. These findings indicate that participants paid attention to the instructions and stimuli during the learning phase.

One noteworthy observation from the unstandardized and untransformed IAT effects from both Experiment 1 and Experiment 2 (see table A2 and table B2) is that IAT effects appear to generally reflect a positive evaluation of the target character. As one might expect, IAT effects were positive (i.e., greater than 0) when a positive explicit behavior was learned. Interestingly, IAT effects were also greater than zero in the second IAT after negative explicit behavior had been learned in the second block (Valence Condition 1). Importantly, when negative explicit information was learned first, IAT effect scores in the first IAT do not differ from zero (and are greater than zero in the second IAT). Thus, there appears to be an asymmetrical effect of negative versus positive learned behavioral information on the IAT, which should be investigated by future studies and replication attempts.

Given the results of Experiment 2, we were not able to replicate the dissociating between implicit and explicit measures described in the original study, despite close adherence to the original methods and procedures. Compared to the original study, we used the same instructions, IAT target words, and primes. We slightly modified the behavioral descriptions

to eliminate overlap between the words used in the descriptions and those used as primes or new words in the recognition task. The most relevant modification to the original study protocol was the reduction of the presentation duration of the briefly presented primes: We presented primes for only 17 ms in order to achieve a functional equivalence in the perception of primes compared to the original experiment. Nevertheless, we were unable to reproduce the implicit attitude findings of the original study and instead found *evidence for the absence* of the effect.

General Discussion

The main goal of this set of Experiments was to investigate the potential influence of subliminal stimuli on (implicit) attitudes. In two experiments, we did not observe such an influence. We decided to reduce the visibility of the briefly presented stimuli in Experiment 2 compared to Experiment 1, in order to achieve at-chance prime recognition as described in the original study. However, even with prime recognition at chance level, we were unable to find an effect of the primes on the IAT effect.

A key finding of the original work was that participants held two distinct attitudes at the same time. Replicating this finding—regardless of the actual visibility of the primes—would be an important contribution to the literature. However, we were unable to reproduce this “dual attitude” effect using the same presentation time of prime words as in the original study in Experiment 1, or using a shorter presentation time in Experiment 2.

One noteworthy finding of Experiment 1 and the pilots prior to Experiment 2 was that prime stimuli seemed to be visible when presented for the same amount of time as in the original study. We can only speculate whether the prime recognition tests in our Experiments were more sensitive (as we used truly new words versus words that had previously appeared in other parts of the study), or whether differences in the presentation method of the words (e.g., font size, font color) caused prime words to be recognized above chance in our studies with a presentation time of 23.5 ms. Future studies should take into account that merely using the same presentation time as reported in another study does not guarantee that stimuli are perceived in a similar way; ideally, trial-by-trial visibility checks should be used to determine actual visibility (Stahl et al., 2016).

Limitations

In Experiment 2, we reduced the presentation time of the primes to ensure a prime presentation below the visibility threshold (i.e., to present the prime words subliminally as described in the original study). It seems possible that the reduced presentation time—while successfully reducing visibility—was insufficient for a subliminal processing of the prime words. Future studies could investigate the possibility of a longer (i.e., 25 ms) presentation time, while ensuring that the prime words are indeed not visible to the participants. This could be achieved by a parafoveal and/or masked presentation of the prime word, as implemented by Rydell and colleagues in a subsequent replication study of this effect (Rydell et al., 2008). Including a forward mask and a parafoveal presentation could have the effect that prime words are not visible to the participants, but a longer presentation time could result in an unconscious processing of the primes and possibly therefore in an effect of prime words on implicit attitude measures. It should be noted however, that another study on

Evaluative Conditioning, using a masked and parafoveal presentation method, was not able to detect EC effects on implicit or explicit measures (Dedonder et al., 2014).

One additional noteworthy characteristic of the method used in the original experiment—which has not been discussed previously—is the length of feedback during the learning phase. At 5000 ms, this feedback is relatively long; this was sufficiently noteworthy that some participants mentioned it after the study. While we did not change the length of feedback—in order to keep the method as close as possible to the original study—it seems possible that participants’ attention may not remain focused on the screen for the entire duration of the feedback. Importantly, prime processing may suffer from this reduction of attention because the brief prime words were presented just after the feedback, with only a short fixation cross preceding the prime. As attention towards the subliminally presented stimulus might be a necessary precondition for their cognitive processing (Dehaene et al., 2006), this characteristic of the learning phase could be further investigated—and perhaps improved upon—in future studies. Such improvements (i.e., ways to ensure full attention on the briefly presented primes) might help strengthen the effects of primes on attitudes, but they would simultaneously induce greater prime visibility, thereby rendering it less likely to obtain the original finding.

Implications & Outlook

While the original study is cited as an important finding demonstrating the influence of briefly presented primes on participants’ implicit attitudes (Sweldens et al., 2014), our results show that the finding is not easy to replicate. However, at this point we would refrain from drawing strong conclusions from the present two replication studies, for two reasons: First, while we ran two replication in two different labs, we used the same experimental script (i.e., same font, same colors, etc.). Specific—unknown—settings might have interfered with the original effect; and replications by other labs, perhaps using a variety of sensible presentation parameters, would be highly desirable. Second, there are two replications of the original finding (albeit with slightly altered methods; Rydell & McConnell, 2006; Rydell et al., 2008). Those studies—while not independent—are important replications of the effect and should be taken into account when running additional replication attempts. Nevertheless, should future studies confirm the interpretation that the original findings are not robust, this will have important theoretical implications as discussed next.

One theoretical aspect of the original study worth discussing is the duality of attitudes and their measurement, as proposed in the systems of reasoning account (Rydell & McConnell, 2006). The original finding suggests that explicitly learned attitudes map onto explicit rating measures, whereas implicitly acquired attitudes map onto implicit measures. In other studies in support of the systems of reasoning account, it was shown that explicit attitudes adapted quickly to new (explicit) information, while implicit attitudes did not adapt quickly, when new (explicit) information was encountered (Rydell & McConnell, 2006). This effect was taken as an indication for the duality of fast-changing explicit attitudes versus slow-changing implicit attitudes. However, as the authors themselves suggest, the delay in adaptation to new information exhibited by implicit measures might be due to their reduced sensitivity (i.e., the IAT may simply be less sensitive to the information), instead of supporting the existence of dual attitudes. According to the systems of reasoning account (Rydell & McConnell, 2006), demonstrating an influence of subliminal primes on implicit

attitudes (while influencing the explicit attitudes using supraliminal behavior descriptions) was therefore important to lend credibility to the proposed duality of attitudes. We were not able to demonstrate this important effect in our two replication studies. If additional studies support the present finding, the differences in implicit and explicit attitudes found previously might indeed turn out to be best explained by a difference in the measures, not in mental constructs.

Interestingly, in the above-mentioned studies, it was shown that 100 trials of explicit counter-attitudinal information had an effect on implicit attitudes (Rydell & McConnell, 2006). With these results in mind, it would be possible that the explicit information presented to participants in our replication studies might similarly have interfered with the opposite-valence information provided by the briefly presented primes. Additional findings also contradict the proposition of an unequivocal mapping of explicit and implicit measures and attitudes. First, several EC studies reported unconscious EC effects on explicit ratings (Hütter, Sweldens, Stahl, Unkelbach, & Klauer, 2012; Olson & Fazio, 2001). Second, recent work also showed that implicit attitudes can be accessed and reported in a conscious manner (Hahn, Judd, Hirsh, & Blair, 2014). Third, additional work showed that even indirect measures of EC effects relied on contingency awareness between CS and US during the learning phase (Pleyers et al., 2007; Stahl et al., 2009). Taken together, the assumption of dual attitude models, that “people can hold different implicit attitudes and explicit attitudes about an attitude object at the same time” (Rydell & McConnell, 2006, p. 1007) is challenged both by recent work and the present replication attempts.

Another theoretical aspect of this replication attempt is the role of contingency awareness in EC. The results of the present replication studies are in line with other studies claiming that contingency awareness is necessary for evaluative conditioning (Dedonder et al., 2014; Pleyers et al., 2007; Stahl et al., 2016, 2009). The present studies did not provide evidence for EC in the absence of contingency awareness; yet the existence of such an effect has been proposed by dual-process models of attitude acquisition (e.g., Gawronski & Bodenhausen, 2006). As the original study has previously been cited as (methodologically) strong evidence for EC with subliminal stimuli—and therefore for EC without contingency awareness—the results of these replication attempts question the robustness of this effect and therefore lend support for the notion that EC requires contingency awareness. Clearly, the present work realizes only one of many possible ways to investigate contingency awareness in EC; other possible realizations should be taken into account by future studies addressing the necessity of contingency awareness in EC.

To conclude, we set out to replicate the finding that briefly presented primes had a selective influence on participants’ implicit attitudes. In two experiments, we were not able to replicate the original finding but instead found evidence for its absence. If confirmed by future work, this suggests a lack of robustness of the original finding that would considerably weaken the empirical support for dual-attitude theories.

References

- Aust, F., & Barth, M. (2016). *Papaja: Create apa manuscripts with rmarkdown*. Retrieved from <https://github.com/crsh/papaja>
- Bluemke, M., & Friese, M. (2006). Do features of stimuli influence IAT effects? *Journal of Experimental Social Psychology*, 42(2), 163–176. doi:10.1016/j.jesp.2005.03.004
- De Houwer, J. (2007). A conceptual and theoretical analysis of evaluative conditioning. *The Spanish Journal of Psychology*, 10(02), 230–241.
- De Houwer, J. (2009). The propositional approach to associative learning as an alternative for association formation models. *Learning & Behavior*, 37(1), 1–20. doi:10.3758/LB.37.1.1
- De Houwer, J., Baeyens, F., & Eelen, P. (1994). Verbal evaluative conditioning with undetected US presentations. *Behaviour Research and Therapy*, 32(6), 629–633. doi:10.1016/0005-7967(94)90017-5
- Dedonder, J., Corneille, O., Bertinchamps, D., & Yzerbyt, V. (2014). Overcoming Correlational Pitfalls: Experimental Evidence Suggests That Evaluative Conditioning Occurs for Explicit But Not Implicit Encoding of CS-US Pairings. *Social Psychological and Personality Science*, 5(2), 250–257. doi:10.1177/1948550613490969
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, 10(5), 204–211. doi:10.1016/j.tics.2006.03.007
- Field, A. P., & Moore, A. C. (2005). Dissociating the effects of attention and contingency awareness on evaluative conditioning effects in the visual paradigm. *Cognition & Emotion*, 19(2), 217–243. doi:10.1080/02699930441000292
- Fulcher, E. P., & Hammerl, M. (2001). When All Is Revealed: A Dissociation between Evaluative Learning and Contingency Awareness. *Consciousness and Cognition*, 10(4), 524–549. doi:10.1006/ccog.2001.0525
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132(5), 692–731. doi:10.1037/0033-2909.132.5.692
- Gawronski, B., & Bodenhausen, G. V. (2011). The Associative–Propositional Evaluation Model. In *Advances in Experimental Social Psychology* (Vol. 44, pp. 59–127). Elsevier. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/B9780123855220000020>
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality*

- and Social Psychology*, 74(6), 1464–1480. doi:[10.1037/0022-3514.74.6.1464](https://doi.org/10.1037/0022-3514.74.6.1464)
- Hahn, A., Judd, C. M., Hirsh, H. K., & Blair, I. V. (2014). Awareness of Implicit Attitudes. *Journal of Experimental Psychology: General*, 143(3), 1369–1392. doi:[10.1037/a0035028](https://doi.org/10.1037/a0035028)
- Hofmann, W., De Houwer, J., Perugini, M., Baeyens, F., & Crombez, G. (2010). Evaluative conditioning in humans: A meta-analysis. *Psychological Bulletin*, 136(3), 390–421. doi:[10.1037/a0018916](https://doi.org/10.1037/a0018916)
- Houwer, J. de, Hendrickx, H., & Baeyens, F. (1997). Evaluative Learning with “Subliminally” Presented Stimuli. *Consciousness and Cognition*, 6(1), 87–107. doi:[10.1006/ccog.1996.0281](https://doi.org/10.1006/ccog.1996.0281)
- Hütter, M., Sweldens, S., Stahl, C., Unkelbach, C., & Klauer, K. C. (2012). Dissociating contingency awareness and conditioned attitudes: Evidence of contingency-unaware evaluative conditioning. *Journal of Experimental Psychology: General*, 141(3), 539–557. doi:[10.1037/a0026477](https://doi.org/10.1037/a0026477)
- Kerpelman, J. P., & Himmelfarb, S. (1971). Partial reinforcement effects in attitude acquisition and counterconditioning. *Journal of Personality and Social Psychology*, 19(3), 301–305. doi:[10.1037/h0031447](https://doi.org/10.1037/h0031447)
- Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, 28(1), 3–26. doi:[10.1037//0097-7403.28.1.3](https://doi.org/10.1037//0097-7403.28.1.3)
- McConnell, A. R., & Leibold, J. M. (2001). Relations among the Implicit Association Test, Discriminatory Behavior, and Explicit Measures of Racial Attitudes. *Journal of Experimental Social Psychology*, 37(5), 435–442. doi:[10.1006/jesp.2000.1470](https://doi.org/10.1006/jesp.2000.1470)
- Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, 32(02), 183. doi:[10.1017/S0140525X09000855](https://doi.org/10.1017/S0140525X09000855)
- Morey, R. D., & Rouder, J. N. (2015). *BayesFactor: Computation of bayes factors for common designs*. Retrieved from <https://CRAN.R-project.org/package=BayesFactor>
- Olson, M. A., & Fazio, R. H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science*, 12(5), 413–417.
- Peirce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods*, 162(1-2), 8–13. doi:[10.1016/j.jneumeth.2006.11.017](https://doi.org/10.1016/j.jneumeth.2006.11.017)
- Pleyers, G., Corneille, O., Luminet, O., & Yzerbyt, V. (2007). Aware and (Dis)Liking: Item-Based Analyses Reveal That Valence Acquisition via Evaluative Conditioning Emerges Only When There Is Contingency Awareness. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(1), 130–144. doi:[10.1037/0278-](https://doi.org/10.1037/0278-)

7393.33.1.130

- R Core Team. (2016). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>
- Rouder, J. N. (2016). The what, why, and how of born-open data. *Behavior Research Methods*, 48(3), 1062–1069. doi:[10.3758/s13428-015-0630-z](https://doi.org/10.3758/s13428-015-0630-z)
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5), 356–374. doi:[10.1016/j.jmp.2012.08.001](https://doi.org/10.1016/j.jmp.2012.08.001)
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16(2), 225–237. doi:[10.3758/PBR.16.2.225](https://doi.org/10.3758/PBR.16.2.225)
- Rydell, R. J., & McConnell, A. R. (2006). Understanding implicit and explicit attitude change: A systems of reasoning analysis. *Journal of Personality and Social Psychology*, 91(6), 995–1008. doi:[10.1037/0022-3514.91.6.995](https://doi.org/10.1037/0022-3514.91.6.995)
- Rydell, R. J., McConnell, A. R., & Mackie, D. M. (2008). Consequences of discrepant explicit and implicit attitudes: Cognitive dissonance and increased information processing. *Journal of Experimental Social Psychology*, 44(6), 1526–1532. doi:[10.1016/j.jesp.2008.07.006](https://doi.org/10.1016/j.jesp.2008.07.006)
- Rydell, R. J., McConnell, A. R., Mackie, D. M., & Strain, L. M. (2006). Of Two Minds Forming and Changing Valence-Inconsistent Implicit and Explicit Attitudes. *Psychological Science*, 17(11), 954–958. doi:[10.1111/j.1467-9280.2006.01811.x](https://doi.org/10.1111/j.1467-9280.2006.01811.x)
- Schönbrodt, F. D., Wagenmakers, E.-J., Zehetleitner, M., & Perugini, M. (2015). Sequential Hypothesis Testing with Bayes Factors: Efficiently Testing Mean Differences. *SSRN Electronic Journal*. doi:[10.2139/ssrn.2604513](https://doi.org/10.2139/ssrn.2604513)
- Simonsohn, U. (2015). Small Telescopes: Detectability and the Evaluation of Replication Results. *Psychological Science*, 26(5), 559–569. doi:[10.1177/0956797614567341](https://doi.org/10.1177/0956797614567341)
- Singmann, H., Bolker, B., Westfall, J., & Aust, F. (2016). *Afex: Analysis of factorial experiments*. Retrieved from <https://CRAN.R-project.org/package=afex>
- Stahl, C., Haaf, J., & Corneille, O. (2016). Subliminal evaluative conditioning? Above-chance CS identification may be necessary and insufficient for attitude learning. *Journal of Experimental Psychology: General*, 145(9), 1107–1131. doi:[10.1037/xge0000191](https://doi.org/10.1037/xge0000191)
- Stahl, C., Unkelbach, C., & Corneille, O. (2009). On the respective contributions of awareness of unconditioned stimulus valence and unconditioned stimulus identity in attitude formation through evaluative conditioning. *Journal of Personality and Social Psychology*, 97(3), 404–420. doi:[10.1037/a0016196](https://doi.org/10.1037/a0016196)
- Stroebe, W., & Strack, F. (2014). The Alleged Crisis and the Illusion of Exact Replication.

- Perspectives on Psychological Science*, 9(1), 59–71. doi:[10.1177/1745691613514450](https://doi.org/10.1177/1745691613514450)
- Sweldens, S., Corneille, O., & Yzerbyt, V. (2014). The Role of Awareness in Attitude Formation Through Evaluative Conditioning. *Personality and Social Psychology Review*, 18(2), 187–209. doi:[10.1177/1088868314527832](https://doi.org/10.1177/1088868314527832)
- Wickham, H., & Francois, R. (2016). *Dplyr: A grammar of data manipulation*. Retrieved from <https://CRAN.R-project.org/package=dplyr>

Appendix A

Descriptive Statistics of Evaluative Measures and IAT reaction times in Experiment 1

Table A1

Unstandardized explicit ratings of the target character after Block 1 and Block 2 in Experiment 1 with either positive behavior being characteristic of the target character and negative behavior uncharacteristic ('positive') or negative behavior being characteristic of the target character and positive behavior uncharacteristic ('negative').

	Rating 1	Rating 2	Rating 3	Rating 4	Rating 5	Rating 6	Rating 7
Block 1, positive							
Mean	7.96	8.21	7.96	7.61	8.29	7.82	84.96
SD	1.29	1.20	2.08	1.81	1.70	2.26	17.41
Block 2, negative							
Mean	2.54	2.61	2.25	2.32	2.50	2.57	21.68
SD	2.15	1.93	2.03	1.98	2.22	1.99	24.68
Block 1, negative							
Mean	1.43	1.61	1.74	1.65	1.48	2.00	12.04
SD	1.47	1.50	1.81	1.64	1.47	1.73	18.63
Block 2, positive							
Mean	7.17	7.48	7.83	7.22	8.00	7.39	78.26
SD	1.99	1.93	1.80	2.00	1.78	1.88	20.60

Note. Rating 1 = 'Unlikable-Very Likable'; Rating 2 = 'Bad-Good'; Rating 3 = 'Mean-Pleasant'; Rating 4 = 'Disagreeable-Agreeable'; Rating 5 = 'Uncaring-Caring'; Rating 6 = 'Cruel-Kind' (all 9 point scales); Rating 7 = 'feeling thermometer' (0 to 100).

Table A2

Unstandardized and untransformed IAT effects (i.e., Mean RT[target character and negative] - Mean RT[target character and positive]) in Experiment 1 with higher values indicating a more positive attitude towards the target character (in ms).

	Time 1 IAT	Time 2 IAT	Time 1 IAT	Time 2 IAT
Valence Behavior	positive	negative	negative	positive
Mean	191.94	109.73	-36.39	53.44
SD	188.25	118.96	167.1	111.25

Note. The valence of the primes was always opposite to the valence of the behavior.

Appendix B

Descriptive Statistics of Evaluative Measures and IAT reaction times in Experiment 2

Table B1

Unstandardized explicit ratings of the target character after Block 1 and Block 2 in Experiment 2 with either positive behavior being characteristic of the target character and negative behavior uncharacteristic ('positive') or negative behavior being characteristic of the target character and positive behavior uncharacteristic ('negative').

	Rating 1	Rating 2	Rating 3	Rating 4	Rating 5	Rating 6	Rating 7
Block 1, positive							
Mean	8.39	8.03	8.10	8.00	8.32	8.39	84.65
SD	1.56	1.70	1.68	1.77	1.56	1.52	23.11
Block 2, negative							
Mean	2.06	1.81	1.71	1.71	1.81	1.74	11.13
SD	2.13	1.22	1.40	1.32	1.58	1.37	15.26
Block 1, negative							
Mean	1.58	1.58	1.46	1.58	1.50	1.58	8.38
SD	1.24	0.90	0.90	0.95	0.95	0.86	10.95
Block 2, positive							
Mean	7.42	7.15	7.31	7.23	7.77	7.69	77.73
SD	1.81	1.99	2.07	2.05	2.12	2.13	22.44

Note. Rating 1 = 'Unlikable-Very Likable'; Rating 2 = 'Bad-Good'; Rating 3 = 'Mean-Pleasant'; Rating 4 = 'Disagreeable-Agreeable'; Rating 5 = 'Uncaring-Caring'; Rating 6 = 'Cruel-Kind' (all 9 point scales); Rating 7 = 'feeling thermometer' (0 to 100).

Table B2

Unstandardized and untransformed IAT effects (i.e., Mean RT[target character and negative] - Mean RT[target character and positive]) in Experiment 2 with higher values indicating a more positive attitude towards the target character (in ms).

	Time 1 IAT	Time 2 IAT	Time 1 IAT	Time 2 IAT
Valence Behavior	positive	negative	negative	positive
Mean	227.6	59.98	43.15	94.33
SD	177.63	107.24	158.89	166.33

Note. The valence of the primes was always opposite to the valence of the behavior.