Accuracy of familiarity decisions to famous faces perceived without awareness depends on attitude to the target person and on response latency

Anna Stone* and Tim Valentine

Department of Psychology,
Goldsmiths College,
University of London,
New Cross,
London SE14 6NW
United Kingdom

Abstract

Stone and Valentine (2004) presented masked 17 ms faces in simultaneous pairs of one famous and one unfamiliar face. Accuracy in selecting the famous face was higher when the famous person was regarded as ‘‘good’’ or liked than when regarded as ‘‘evil’’ or disliked. Experiment 1 attempted to replicate this phenomenon, but produced a different pattern of results. Experiment 2 investigated alternative explanations and found evidence supporting only the effect of response latency: responses made soon after stimulus onset were more accurate to liked than to disliked faces, whereas responses made after a longer delay were equally accurate to disliked faces. It appears that the effect of negative valence was corrected within the space of a few hundred milliseconds. Experiment 3, using an affective priming paradigm, supported the concept that an early-arising effect of valence is corrected if it is misleading to the directed task.

Introduction

There is much evidence that facial expressions can be detected, and can influence psychophysiological and behavioural responses, without awareness of the expression (e.g., Dimberg & Ohman, 1996; Dimberg, Thunberg, & Elmehed, 2000; Johnsen & Hugdahl, 1991, 1993;Mogg & Bradley, 1999; Murphy & Zajonc, 1993; Niedenthal, 1990; Ohman, Esteves, & Soares, 1995; Robinson, 1998; Saban & Hugdahl, 1999; Whalen et al., 1998; Wong, Shevrin, & Williams, 1994). All of these studies presented masked faces for very brief exposure duration [target-to-mask stimulus onset asynchrony (SOA) of less than 35 ms]. Participants were at chance in two-alternative forced-choice tasks of identifying the facial expression, confirming the absence of awareness of the expression. The recognition without awareness of facial emotional expressions is often interpreted in terms of the importance to the individual of detecting the emotion of others. The question then arises of whether facial identities, like facial expressions, can be recognised without awareness of identity. Literature relevant to this question will be examined.

Banse (1999) presented the face or name of the participant or a relationship partner as the prime stimulus, for 10.5 ms with backward masking, followed after SOA of 42 ms by a Chinese letter (re Murphy & Zajonc, 1993). Targets were evaluated more positively when preceded by the partner's face/name than by the participant's own face/name: this was the predicted effect, based on the observation that a partner tends to be evaluated more positively than the self. The similarity in the results obtained from face and name primes was taken as implying that the person schemata had been activated. One drawback is that accuracy was well above chance in a two-alternative forced-choice of self or other, with a single face or name presented under the same masked conditions, so it is not clear that participants were entirely unaware of the face or name identity. The limited number of stimuli raises the possibility that even if participants were able to perceive only a vague outline of the masked face, this could have sufficed to enable a correct decision about which of the two persons had been presented. Banse (2001) used an affective priming paradigm to investigate whether famous and personally familiar faces might be recognised without awareness of identity. Primes were presented for the same duration used by Banse (1999) and again, it is not clear that participants were entirely unaware of facial identity.

Stone, Valentine, and Davis (2001) reported that responses to famous faces perceived without awareness of facial identity differed according to valence. Experiment 1 found that skin conductance responses to masked 17 ms faces were higher to the faces of famous persons subsequently evaluated ‘‘good’’ than to the faces of persons evaluated ‘‘evil,’’ but did not distinguish between famous and unfamiliar faces. (Responses tended to be higher to good faces than to unfamiliar faces, but tended to be lower to evil faces than to unfamiliar faces.) When faces were exposed for 220 ms, a duration that permits conscious recognition, there was an effect of familiarity but no effect of valence: skin conductance responses were higher to famous faces than to unfamiliar faces with no difference between ‘‘good’’ and ‘‘evil’’ faces. Responses were above chance accuracy in a two-alternative forced-choice of ‘‘good’’ or ‘‘evil’’ to masked 17 ms faces of whose identity participants were unaware (Experiment 3).

Stone and Valentine (in press) used a procedure based on Mogg and Bradley (1999). Masked 17 ms faces were presented in simultaneous pairs of a famous and an unfamiliar face, matched on physical characteristics, in LVF and RVF. These were followed by a dot-probe in either LVF or RVF to which participants made a speeded two-alternative forced-choice discrimination response. Orientation of attention towards the famous face would be demonstrated by faster or more accurate responses to the dot-probe when it appeared in the same VF as the famous face. Participants were subsequently asked to evaluate each famous person as ‘‘good’’ or ‘‘evil’’ on a 7-point scale from -3 (very evil) to +3 (very good). Fewer errors were made when the dotprobe was presented in the same VF as the famous face compared to the opposite VF, as long as the famous person was evaluated as neutral or good (evaluation from -1 to 3). A reverse effect was observed, with more errors to dot-probes presented in the VF of the famous face, when the famous person was evaluated as evil (-3 or -2). The within-item analysis, comparing performance between participants who had evaluated the same famous persons as good-neutral or evil, confirmed that the effect was due to participants reactions to the famous persons and not to any confounding factor. This effect was interpreted as the orientation of attention towards the faces of famous persons evaluated as ‘‘good’’ but not towards those evaluated as ‘‘evil.’’

In a separate awareness check task, the same masked 17 ms famous–unfamiliar face pairs were presented simultaneously while participants attempted to select the famous face. Overall, accuracy at chance supported participant claims of no awareness of facial familiarity or identity. At the same time, responses were more accurate to the faces of famous persons evaluated as ‘‘good’’ than to the faces of persons evaluated as ‘‘evil.’’ This was published as Experiment 1 in Stone and Valentine (2004). One limitation is that, without specification of the processes underlying the valence effect, the possibility cannot be ruled out that it arose from some particular (and unknown) characteristic of the specific stimulus set. For this reason it was considered advisable to repeat the experiment using a different set of stimuli and a new participant sample.

Experiment 1 of the present paper was an attempt to replicate the valence effect with a larger set of stimuli. As will be seen, Experiment 1 did not replicate the valence effect: response accuracy was not consistently lower for negative than positive faces. Experiment 2 contrasted three potential explanations for the results of Experiment 1 and found evidence supporting only one, the influence of response latency. Experiment 3 used an affective priming task to examine another prediction derived from the results of Experiment 1.

Experiment 1

Experiment 1 was an attempt to replicate the effect of valence in the awareness check task with a larger set of stimuli. Consequent to the change in the stimulus set, the terms positive and negative were used in the evaluation of the famous persons, replacing the terms good and evil that were used by Stone and Valentine (2004). The stimuli were selected on the expectation that each famous person would be evaluated as positive by some participants and negative by others. Balanced valence ratings would permit a rigorous within-item analysis, in which accuracy could be compared for the same famous–unfamiliar face pair between participants evaluating the famous person positive vs. negative, ruling out confounds arising from systematic differences between famous and unfamiliar faces, or among famous faces. If the same famous–unfamiliar face-pair results in different responses depending on participants‘ evaluations of the famous person, then this will support the claim that responses depend on participants’ attitudes and not on any confounding factor. The prediction was that responses would be more accurate to the faces of famous persons rated positive than negative.

Method

Participants

Participants were 46 students at Goldsmiths College, London. Data were excluded from eight participants who failed to evaluate a minimum of 10% of the famous people as negative in the post-experimental evaluation, since it was suspected that these participants might not have been using the scale correctly. Data were excluded from five participants who selected more famous faces than expected by chance (binomial distribution, one-tailed, cut-off = 55.6%) since for these participants, the possibility of some awareness of facial familiarity could not be ruled out. The remaining 33 participants were 24female and nine male, aged between 18 and 50, mean 23.4, SD 7.9 years. All of the participants had been resident in the UK for at least 10 years by self-report to maximise the likelihood of knowledge of the famous faces.

Stimuli

Photographs of famous and unknown faces of a uniform quality were digitised to produce images of 16 greys, 150 x 200 pixels in size. The stimulus set comprised 126 pairs of one famous with one unfamiliar face matched on sex, race and approximate age. The faces in each pair showed a similar pose and facial expression. No data were collected to verify equivalence between the famous and unfamiliar faces on distinctiveness, attractiveness, or any other feature on which the stimuli might vary. The intention was to perform analyses within-items, with each famous person rated as positive by some participants and as negative by others, so that systematic variations between famous faces and their paired unfamiliar faces could not explain any observed experimental result. Names and examples of stimuli are given in Appendices A and B. The average luminance of the famous and unfamiliar faces was approximately 5.5 cd/m_2, measured on a Minolta CS100 colour chronometer at a distance of approximately 1 m from the screen, the distance at which participants were seated.

The mask was a collage of parts of unfamiliar faces, of the same size as the famous and unfamiliar faces.

Apparatus

A personal computer running MEL2 software was used to display the faces at a 640 x 480 screen resolution. Response times and accuracy of response were measured and recorded by the computer.

Design

The design comprised two independent within-item factors: valence of famous person (evaluated by experimental participants) and visual field containing the famous face (LVF vs. RVF). The dependent variables were speed and accuracy of response. The stimuli were presented in two blocks of 126 trials each, with each face pair appearing once in each block. Each famous face appeared in the LVF in one block and the RVF in the other block. In each block, there were equal numbers of famous faces in the LVF and the RVF. For each famous face, approximately equal numbers of participants saw it in the LVF and the RVF in each block. Thus, visual field and block were fully counterbalanced across participants. Within each block, the sequence of presentation was randomised by the computer for each participant. The duration of the forward and backward masks was 100 ms.

Stone and Valentine (2004) suggested that very few faces could be consciously recognised when presented for 17 ms in pairs with a mask similar to that used in the present series of experiments.

Procedure

Participants carried out the tasks individually in a darkened, air-conditioned room. The sequence of events on each trial was as follows: central fixation cross for 500 ms, forward masks in LVF and RVF for 100 ms, famous and unfamiliar face for 17 ms, backward masks for 100 ms, then the question ‘‘left or right’’ displayed until a response was received. The response was made by pressing one of two keys: to the left of the keyboard to indicate the LVF, and to the right of the keyboard to indicate the RVF. Each trial was initiated by the response to the previous trial after an inter-trial interval of 1 s. The response time was calculated from the offset of the backward mask. The two faces were approximately 4.5 cm by 6 cm and were presented at a distance of 9 cm apart, subtending a visual angle of approximately 4 degrees from fixation. The masks were presented in the same screen position as the faces.

Participants were informed that two faces would be flashed up very briefly, one on either side of the screen, preceded and followed by a mask comprised of a collage of parts of unfamiliar faces. Each pair of faces would contain one famous person and one unfamiliar person, and participants were asked to select on which side of the screen the famous face had appeared. Participants were told they would find it very difficult to see the real faces and this should be no cause for concern, but they should attend carefully to the screen. They were asked to guess if unable to see anything of the stimulus faces. Participants were asked to look at the central fixation cross before each trial and to respond as quickly as possible. At the end of the task, participants were asked whether they had been able to recognise any of the faces displayed during the experiment, and were strongly encouraged to guess.

Following this, participants were shown the famous faces used in the experiment, one at a time, and asked to identify each face, either by name or by a combination of biographical information that uniquely identified the individual person. Faces that were uniquely identified were shown again, one at a time, in a different random sequence, and evaluated on a 7-point scale from -3 (very negative) to +3 (very positive). Participants were asked to evaluate the valence of the person, not the face, considering any knowledge they had of the person, and to give their first impression. Finally, participants were debriefed and thanked for their participation.

Results and discussion

Where a participant could not uniquely identify a famous face in the post-experimental evaluation, the responses for this combination of participant and item were excluded from the analysis (19% of trials). No masked 17 ms faces were recognised by any participant. Trials were excluded where the response time was longer than 5000 ms. Some items had missing data, having been evaluated as either positive or negative by all participants, and so 106 items were included in the analysis. The participants' analysis was calculated over these items. The proportion of participants giving a negative evaluation to each of the 106 included items ranged from 0.06 to 0.94, mean 0.30, SD 0.2.

For each participant, mean accuracy was calculated over the faces rated as positive and separately over the faces rated as negative (faces evaluated as zero were classified as positive). So two values were calculated for each participant: accuracy-positive$_p$ (mean = 50.5, standard error = .7) and accuracy-negative$_p$ (mean = 50.5, SE = 1.1). Similarly, for each item, mean accuracy was calculated over all participants rating the famous person as positive and separately over all participants rating the person as negative. So two values were calculated for each item: accuracy-positive$_i$ (mean = 51.1, SE = 1.0) and accuracy-negative$_i$ (mean = 50.1, SE = 1.5). If a famous person was evaluated as positive and negative by unequal numbers of participants, different numbers of participants contributed to the calculation of accuracy-positive$_i$ and accuracy-negative$_i$.

Overall accuracy for all famous faces did not differ from chance, one-sample $t_i(105) = 0.88$, ns, and $t_p(32) = 1.15$, ns. Paired-samples t tests revealed no effect of valence, $t_i(105) = 0.60$, ns, and $t_p(32) = 0.06$, ns, in contrast to the effect reported by Stone and Valentine (2004).

One obvious procedural difference is that participants in Experiment 1 of Stone and Valentine (2004) had performed a separate attention orientation task before the awareness check, whereas participants in the present study had not. The attention orientation task had presented the same masked 17 ms face pairs so that participants had gained experience in perceiving the stimuli. This suggests that a comparison of blocks 1 and 2 might be interesting. ANOVA was performed with two within-participants and within-items factors of block (1 vs. 2) and valence. The main effect of block was non-significant, $F_i(1, 105) = 1.43$, ns, and $F_p(1, 32) = 4.11$, $p = .051$, but the interaction with valence was significant, $F_i(1, 105) = 4.78$, $p < .04$, and $F_p(1, 32) = 5.50$, $p < .03$. Paired-samples t tests revealed that for negative faces, responses tended to be more accurate in block 1 ($mean_i = 53.4$, SE = 2.2) than in block 2 ($mean_i = 46.8$, SE = 2.3), $t_i(105) = 1.96$, $p < .06$, and $t_p(32) = 2.62$, $p < .02$. For positive faces, accuracy did not differ between block 1 ($mean_i = 50.2$, SE = 1.3) and block 2 ($mean_i = 52.1$, SE = 1.5), $t_i(105) = .93$, ns, and $t_p(32) = .66$, ns.

Mean accuracy of responses to negative and positive faces was calculated for each chunk of five trials (to provide a more stable view of the data) over all participants and items. Accuracy was negatively correlated with chunk for negative faces, $r(51) = -0.33$, $p < .02$, but tended to be positively correlated with chunk for positive faces, $r(51) = 0.24$, $p < .09$, confirming that response accuracy for negative faces, but not positive faces, declined as the task progressed. There are three possible reasons why responses to negative faces were less accurate towards the end of the task than at the start and these will be considered in turn.

Practice

The practice explanation proposes that participants became better able to extract information from masked faces as the task progressed, so that the valence associated with each face was more strongly activated, leading to the decline in accuracy of responses to negative faces. The practice explanation would be consistent with Experiment 1 of Stone and Valentine (2004) in which a separate attention orientation task before the awareness check had presented the same masked famous-unfamiliar face pairs. There is some empirical evidence to support the practice account. Dagenbach, Carr, and Wilhelmsen (1989, Experiment 1) measured participants' exposure thresholds for chance performance in a presence/absence decision on briefly exposed masked words, before starting the experimental trials. Participants were tested again on the same presence/absence decision, at the same threshold exposure duration, after the experimental trials were complete. They found that for a sizeable group of participants (15 of 52, approximately 29%) accuracy in the presence/absence decision was above chance in the post-experimental re-test when it had been at chance in the pre-experimental threshold setting procedure, suggesting that some participants' ability to detect briefly exposed masked stimuli had improved as a result of practice. In the present experiment, increasing ability to derive valence from masked 17 ms famous faces could explain why the accuracy of responses to negative faces declined throughout the task.

Strategy

Comments made by participants during debriefing suggested that they had used different strategies for attempting to perceive the masked faces at the start of the task and towards the end. Many reported that to begin with, they had tried very hard to see anything they could, but this was so difficult they had given up and simply guessed their responses. It makes sense that deliberate effort would be reduced towards the end of a task in which there was no feedback or evidence that effort improved accuracy.

Relevant to this possibility, Snodgrass, Shevrin, and Kopka (1993) reported a study investigating the different strategies that participants might adopt for engaging with a task involving the perception of briefly exposed masked stimuli. Participants were asked to decide which of four words, two with pleasant meanings and two with unpleasant meanings, had been presented on each trial, using one of two strategies. The Pop strategy asked participants to ‘‘look where the word is presented and say whatever word pops into your mind.’’ The Look strategy asked participants to ‘‘look very hard where the word is presented. . . for anything you can see. . . use these cues when making your decision’’ (p. 196). Strategy interacted with word meaning: pleasant words tended to be identified more accurately than unpleasant words in the Pop strategy, whereas the converse was observed in the Look strategy.

The Pop strategy gave results similar to the pattern observed in Stone and Valentine (2004) and the non-significant trend in the present Experiment 1, block 2, that is, less accurate responses to evil-negative faces than to good-positive faces. The Look strategy gave results similar to those obtained in Experiment 1, block 1, a tendency to more accurate responses to negative faces. This raises the possibility that something similar to the Look strategy was used in Experiment 1, block 1, and something similar to the Pop strategy in Experiment 1, block 2, which would be consistent with the comments made by participants during debriefing. There are other reports of strategy affecting responses in a task involving perception of briefly exposed masked stimuli (e.g., Dagenbach et al., 1989; Kahan, 2000). It seems that the failure to observe the predicted valence effect consistently in Experiment 1 may have been due to participants’ strategies.

Response latency

In addition to practice and strategy, a third possibility is that response latency measured from stimulus face onset may be a factor, since participants tended to respond faster as the task progressed. A substantial literature on affective priming is relevant to this possibility and will be briefly described.

Affective priming refers to the phenomenon in which responses to a valenced target stimulus are facilitated (inhibited) by the prior presentation of a prime stimulus of the same (opposing) valence to the target. Affective priming appears to be a reliable phenomenon, having been observed using words, objects and faces as primes; using words and pictures as targets; using evaluative decision, lexical decision, degraded word identification, and word pronunciation as the task; and using response time and accuracy as the dependent variable (see Fazio, 2001, for a review). Affective priming has been observed with masked primes of whose identity participants were not aware (e.g., Banse, 2001; Bargh, Litt, Pratto, & Spielman, 1989; see Fazio, 2001, for a review). There is substantial evidence that attitudes towards primes are activated automatically, without deliberate evaluation intent (e.g., Bargh, Chaiken, Raymond, & Hymes, 1996; Bargh et al., 1989; Hermans, De Houwer, & Eelen, 1994, 2001). Attitudes appear to be activated very rapidly, with rime to target SOA between 0 and 300 ms (e.g., Banse, 1999, 2001; Bargh et al., 1996; Bargh et al., 1989; De Houwer & Hermans, 1994; De Houwer, Hermans, & Eelen, 1998; Glaser & Banaji, 1999; Hermans et al., 1994, Hermans, De Houwer, & Eelen, 2001; Klauer, Rossnagel, & Musch, 1997; Murphy & Zajonc, 1993; Musch & Klauer, 2001).

There is also evidence that the effect of the attitude evoked by the prime is corrected, where it is misleading with regard to the required response to the target. Several studies have reported significant effects at short SOA of up to 300 ms that are absent at longer SOA, which suggests that the correction process requires some duration (e.g., De Houwer et al., 1998; Glaser & Banaji, 1999; Hermans et al., 1994, 2001; Klauer et al., 1997; see Fazio, 2001, for a review).

In the present Experiment 1, it is possible that lower accuracy of responses to negative than positive faces appeared in block 2 because responses were faster (block 1 mean response time = 851 ms, block 2 mean = 622 ms), and so were formulated during a period when the negative valence of the famous person influenced participants to select the unfamiliar face. With the longer response latency in Experiment 1, block 1, the influence of negative valence had been corrected before the responses were selected, and so response accuracy was not lower for negative faces than for positive faces.

Re-analysis of Experiment 1

The block design of Experiment 1 permits a comparison of shorter latency responses with longer latency responses in blocks 1 and 2. If only block has a significant effect, this would argue against the response latency account. If only response latency has a significant effect, this would make the practice account less likely. It is less easy to investigate the strategy account since participants were not asked on a trial-by-trial basis which strategy they had used. In addition, strategy may have been confounded with practice such that the Look strategy was preferred in block 1 and the Pop strategy in block 2, as suggested by participants during debriefing, or confounded with response latency, such that the Look strategy was used on longer latency trials and the Pop strategy on shorter latency trials.

Responses were divided into short latency and long latency around the grand median response time of 604ms. The mean accuracy of short and long latency responses was calculated for each participant, and separately for each item, by block and valence. The data were analysed in separate ANOVA for participants and items, with factors of valence, block (1 vs. 2) and response latency (short vs. long). Two participants and 34items had missing data and were excluded from the ANOVA. See Table 1 and Figure 1.

Table 1: Mean accuracy (and standard error) of fast and slow responses in the first and second block of Experiment 1, for positive and negative items.

	Positive			Negative		
	Fast	Slow	Total	Fast	Slow	Total
Block 1	50.2 (2.9)	49.3 (2.1)	49.7 (2.6)	49.7 (4.0)	56.1 (3.3)	52.9 (2.5)
Block 2	51.8 (2.3)	52.0 (3.0)	51.9 (2.0)	45.5 (2.9)	53.0 (3.8)	49.2 (2.5)
Total	51.0 (1.8)	50.6 (1.8)	50.8 (1.3)	47.6 (2.6)	54.6 (2.5)	51.1 (1.9)

A

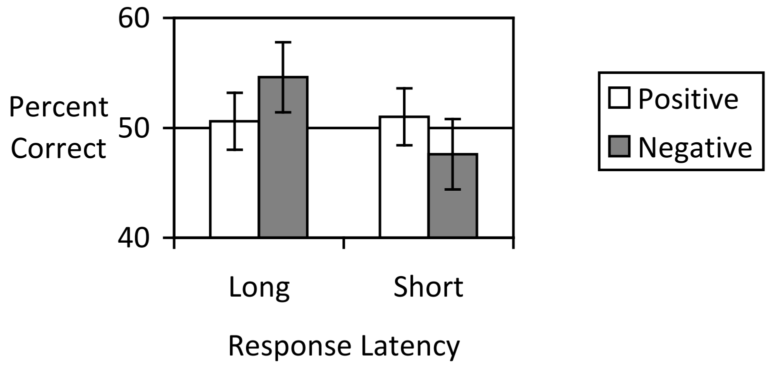

B

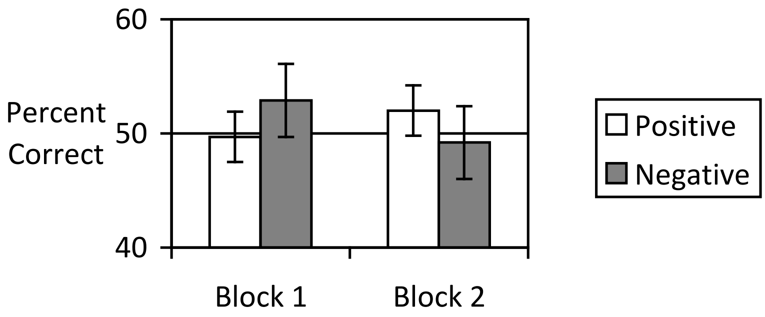

Figure 1: accuracy of responses by valence and response latency (A) and by valence and block (B) in Experiment 1. Bars represent standard errors.

The two-way interaction of valence with response latency approached significance, $F_i(1, 71) = 3.25$, $p < 0.08$, and $F_p(1, 30) = 8.20$, $p < 0.01$. See Fig. 1A. Paired-samples t tests investigated the interaction, including the maximum number of items and participants for whom data was complete (collapsing over block). For negative faces, shorter latency responses ($mean_i = 46.1$, SE = 2.7) were less accurate than longer latency responses ($mean_i = 57.1$, SE = 2.2), $t_i(100) = 3.25$, $p < 0.005$, and $t_p(32) = 2.53$, $p < 0.02$. For positive faces, there was no difference in accuracy between shorter latency responses ($mean_i = 50.9$, SE = 1.4) and longer latency responses ($mean_i = 50.0$, SE = 1.5), $t_i(100) = 0.46$, ns, and $t_p(32) = 1.26$, ns. Shorter latency responses tended to be less accurate to negative than positive faces, $t_i(100) = 1.69$, ns, and $t_p(32) = 1.60$, ns, whereas longer latency responses were more accurate to negative faces, $t_i(100) = 2.53$, $p < 0.02$, and $t_p(32) = 2.87$, $p < 0.01$. This is consistent with the response latency account, but could also be consistent with the strategy account if the Look strategy tended to be used on longer latency responses and the Pop strategy on shorter latency responses.

The two-way interaction of half with valence was non-significant, $F_i(1, 71) = 2.08$, ns, and $F_p(1, 30) = 2.78$, ns. See Fig. 1B. Paired-samples t tests have already been reported. Responses tended to be more accurate to negative faces in block 1, and more accurate for positive faces in block 2. This is consistent with the practice account, but could also be consistent with the strategy account if Look strategy was used predominantly in block 1 while the Pop strategy was used predominantly in block 2.

This inconclusive pattern of results does not enable a decision between the three accounts. Experiment 2 was designed to investigate the accounts directly by manipulating strategy, practice and response latency as independent factors.

Over-correction of valence?

Another observation is worth mentioning: In Experiment 1, longer latency responses were more accurate to negative than positive faces. This could have arisen if the correction for the biasing effect of negative valence went too far and resulted in over-correction. Such over-correction has been observed in studies of affective priming. For example, Glaser and Banaji (1999), in a series of six experiments, reported reverse affective priming, that is, slower responses to targets of valence congruent with the prime than to targets of incongruent valence. They attributed this reverse priming to an over-correction for the biasing effect of the primes (e.g., Greenwald & Banaji, 1995; Stapel, Koomen, & Zeelenberg, 1998; Strack, Schwarz, Bless, Kubler, & Wanke, 1993). Based on the observation of reverse priming with prime-target SOA as short as 150 ms, they theorised that the over-correction was an automatic effect, not requiring deliberate intent by the participants. Klauer et al. (1997, Experiment 2) also reported reverse affective priming in an evaluative decision task, but only at a longer SOA of 1200 ms, and not at shorter SOA of 0 and 100 ms. It seems that the correction for the biasing influence of the primes takes some short time to develop. Glaser and Banaji's (1999) reverse priming was obtained at SOA of 150 and 300 ms, falling within the 100 ms SOA and the 1200 ms SOA of Klauer et al. (1997, Experiment 2) that observed congruent and reverse priming, respectively.

The Klauer et al. (1997) study also manipulated the proportion of prime-target pairs of consistent evaluation (the consistency proportion). The weakest reverse priming was observed with consistency proportion of 0.75, in which condition the primes had predictive value for the evaluative decision to the targets. Correction for the biasing effects of the primes would have been less useful with such a high consistency proportion, so the weaker reverse priming in this condition fits the correction theory. Under the same logic, the strongest reverse priming should have been observed when the consistency proportion was 0.25, where the prime has least predictive value for the evaluative decision to the targets. In fact, the strongest reverse priming was observed with consistency proportion of 0.50, although the difference between this and the 0.25 consistency proportion was very small. Klauer et al. (1997) suggested that inconsistent pairs would have been more easily observed with 0.50 consistency proportion, since in this condition, consistent and inconsistent pairs would follow each other frequently during the experimental priming task.

In Stone and Valentine (2004), the effect of negative valence of the famous person was to cause participants to select the paired unfamiliar face, contrary to the task instruction. There was a similar tendency, though non-significant, in the shorter latency trials of Experiment 1 (see Section 2.2.4). Thus, the effect of valence was misleading to the required response, which implies that the effect of valence should have been corrected, given sufficient time. More accurate responses to negative faces than to positive faces on longer latency trials in Experiment 1 is consistent with the proposal that over-correction, requiring some duration, may occur.

This over-correction hypothesis was investigated by introducing a new affective priming task, described as Experiment 3.

Experiment 2

This experiment investigated the response latency, practice and strategy accounts that were offered as possible explanations for the failure to observe consistently lower response accuracy to negative faces than positive faces in

Experiment 1. The post-experimental evaluation scale was altered to use the terms liked and disliked instead of positive and negative, on the expectation that the new terms would focus participants' attention on their personal attitude towards the target famous persons.

Response latency was investigated by varying the duration of the backward mask between 500 and 100 ms. The duration of 500 ms should have ensured that responses were selected after the effect of valence had been corrected, and so disliked faces should be selected at least as often as liked faces, consistent with block 1 of Experiment 1. The duration of 100 ms, combined with instructions to respond without thinking too long, should have ensured that responses were decided before the correction took place, and so responses to disliked faces should be less accurate than responses to liked faces. This would be consistent with block 2 of Experiment 1. Half the participants in each backward mask condition were instructed to use the Pop strategy and the other half the Look strategy. If strategy is the major determinant of responses then the Pop strategy should result in higher accuracy for liked faces than disliked faces throughout, while the Look strategy should result in the opposite effect. If the practice account is correct, then disliked faces should be selected at least as often as liked faces in the first half, but less often than liked faces in the second half.

Method

Only the changes from Experiment 1 will be described.

Participants

Participants were 64 undergraduate students at Goldsmiths College, London. All had watched UK television for at least 5 years, by self-report, to maximise the likelihood of knowledge of the famous faces. Data were excluded from eight participants who lacked familiarity with the famous faces or failed to comply with experimental instructions. Data were excluded from 18 participants who selected more famous faces than expected by chance (binomial distribution, one-tailed, cutoff = 55%, a = .05) since for these participants the possibility of some awareness of facial familiarity cannot be ruled out. Of these 18 participants, 15 were in the 500 ms backward mask condition and 3 in the 100 ms backward mask condition. This is consistent with the response latency account: disliked faces would be selected less accurately with 100 ms than with 500 ms backward mask, so overall accuracy would be higher at 500 ms.

The remaining 38 participants were 35 female and three male, aged between 18 and 36, mean 22.0, SD 5.6 years. There were 23 participants in the 500 ms backward mask condition and 15 in the 100 ms backward mask condition. More participants performed the 500 ms backward mask condition, in which no effect of valence was predicted, than the 100 ms mask condition, in which lower accuracy of responses to disliked faces than to liked faces was predicted, in order that a null effect in the 100 ms backward mask condition could not be attributed to lack of power in the statistical analysis.

Stimuli

These comprised facial photographs of 80 celebrities, of whom 40 had been included in the items analysis of Experiment 1, and the remaining 40 were new stimuli. The new celebrities were selected on the expectation that they would be liked by some participants and disliked by others (this entailed avoiding popular comedians, criminals and military dictators).

Design

There were two within-participant and within-item factors, valence and half, and two between participant and within-item factors, strategy (Look vs. Pop) and mask duration (500 vs. 100 ms). The dependent variable was accuracy of response, and a correct response was scored by selecting the famous face in a pair.

Each face pair was presented once in each of four blocks giving 320 trials altogether. The position of the famous face was counterbalanced as follows. The faces were randomly divided into two sets. Half the participants in each strategy and backward mask condition saw one set in the LVF in blocks 1 and 3 and the RVF in blocks 2 and 4, and the other set in the RVF in blocks 1 and 3 and the LVF in blocks 2 and 4. For the other half of the participants in each condition, these positions were reversed. Thus, face position was fully counterbalanced over the factors of block, strategy and mask duration. This design ensured equal numbers of famous faces in LVF and RVF in each block and minimised the likelihood that the same face would be presented twice in succession.

Procedure

The major change from previous experiments was the addition of the specific instructions. In the Look condition, participants were asked to ‘‘try hard to see anything you can of the faces that flash up. Most people find that they can see nothing at first, but as the experiment progresses, they are able to see bits of the faces. Some people can see the outline of a face, or the hair, or maybe the eyes or the mouth. Please try hard to see anything you can, and use what you see to make a guess about which of the 2 faces is famous.’’ In the Pop condition, participants were asked to ‘‘relax, and not make any effort to see the faces. Just look at the cross in the centre, relax, let the faces flash up, and let the answer pop into your head.’’ In both the Look and Pop conditions, participants were told, ‘‘It is very important that you follow this strategy because the purpose of the experiment is to contrast the effects of different strategies.’’ In the 100 ms backward mask condition, participants were asked to respond on each trial without thinking too long, and this instruction was emphasised after the practice trials had been completed.

Results and discussion

Where a participant was unable to uniquely identify a famous face in the post-experimental evaluation, the trials for this combination of participant and item were excluded from the analysis (16% of trials). If a face was correctly identified during the experimental trials then all four trials for this face were excluded (1.3% of trials). Trials were excluded where the response time was longer than 5000 ms or shorter than 200 ms from face offset (<1% of trials). For each participant, liked and disliked faces were selected according to the participant’s own ratings. For each item, participants who liked and disliked the item were selected according to the same ratings. Faces evaluated as zero were classified as liked in order to distinguish between disliked persons and the rest. Many items had missing data, having been evaluated as liked or disliked by all participants in a condition (mask duration · strategy), so the item analysis consisted of 29 items. The participant analysis was calculated over these items. The proportion of participants evaluating each of the included famous persons as disliked ranged from 0.18 to 0.78, mean = 0.44, so the included items were roughly balanced between liked and disliked.

ANOVA was performed by participants and by items, with valence and half as within-participant and within-item factors, and strategy and mask duration as between-participant and within-item factors. The main effect of mask duration was non-significant, $F_i(1, 28) = 2.13$, ns, and $F_p(1, 34) = 3.27$, $p < 0.08$, but the interaction of mask duration with valence was significant by items, $F_i(1, 28) = 4.85$, $p < 0.04$, and approached significance by participants, $F_p(1, 34) = 3.00$, $p < 0.10$. See Figure 2 and Table 2.

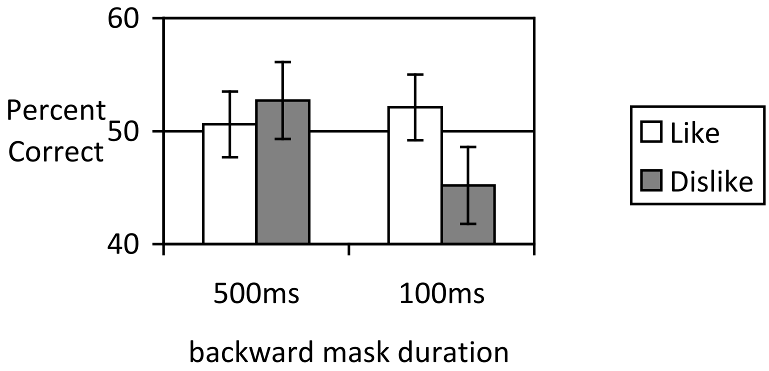

Figure 2: accuracy in the 500ms (long response latency) and 100ms (short response latency) backward mask conditions of Experiment 2. Bars represent standard errors.

The interaction was investigated using t tests with a set at 0.025 (one-tailed) for each individual comparison. With mask duration of 500 ms, responses to disliked faces and liked faces were equally accurate, $t_i(28) = 0.97$, ns, and $t_p(22) = 0.64$, ns. With mask duration of 100 ms, responses to disliked faces were less accurate than responses to liked faces, $t_i(28) = 2.05$, $p = 0.025$, and $t_p(14) = 2.28$, $p < 0.02$. Responses to disliked faces were less accurate at 100 ms than at 500 ms mask duration, significant by items and marginal by participants, $t_i(28) = 2.32$, $p < 0.02$, and $t_p(36) = 1.90$, $p = 0.033$. Responses to liked faces were equally accurate at 500 and at 100 ms, $t_i(28) = 0.61$, ns, and $t_p(36) = 0.16$, ns.

The two-way interaction of half with mask duration approached significance by participants and items, $F_i(1, 28) = 3.52$, $p < 0.08$, and $F_p(1, 34) = 3.43$, $p < 0.08$. Inspection of Table 2 reveals a tendency for accuracy to increase from the first half to the second half at 500 ms mask duration, but to decrease at 100 ms mask duration. The reason for this interaction is not obvious and so it will not be interpreted. The two-way interaction of mask duration with strategy was non-significant, by items, $F_i(1, 28) = 2.37$, ns, and only marginally significant by participants, $F_p(1, 34) = 3.83$, $p < 0.06$, and so will not be interpreted. No other effects reached significance, all F_i and $F_p < 1.9$, $p > 0.18$. In particular, neither strategy nor task half interacted with valence, all F_i and $F_p < 1$.

Table 2: Mean accuracy (and standard error) of responses to disliked and liked items at 100 and 500 ms backward mask duration, under the Look and Pop strategies, in the first and second half of Experiment 2.

Disliked						Liked					
500ms			100ms			500ms			100ms		
Pop n=13	Look n=10	Total n=23	Pop n=8	Look n=7	Total n=15	Pop n=13	Look n=10	Total n=23	Pop n=8	Look n=7	Total n=15
Half 1											
50.4 (3.4)	51.5 (3.7)	51.0 (2.3)	49.2 (4.8)	47.6 (5.8)	48.4 (4.1)	47.7 (3.3)	52.7 (3.9)	50.2 (2.9)	56.8 (3.9)	54.2 (4.4)	55.5 (3.1)
Half 2											
49.6 (3.2)	59.3 (4.3)	54.5 (2.7)	41.7 (4.9)	42.3 (4.3)	42.0 (3.1)	49.5 (3.4)	52.6 (4.0)	51.0 (3.0)	50.1 (3.3)	47.5 (4.8)	48.8 (3.0)
Total											
50.0 (2.7)	55.4 (2.8)	52.7 (1.8)	45.4 (4.0)	45.0 (4.0)	45.2 (3.0)	48.6 (2.8)	52.6 (2.1)	50.6 (2.1)	53.4 (2.6)	50.8 (3.0)	52.1 (2.1)

These results support the prediction based on the response latency account. With 100 ms backward mask it was predicted that responses would be decided before the effect of negative valence was corrected, so that responses to disliked faces would be less accurate than responses to liked faces, and this effect was observed. In contrast, with 500 ms backward mask, it was predicted that the effect of negative valence would have been corrected before responses were decided, and indeed responses to disliked and liked faces were equally accurate. The ANOVA was significant in the items analysis, which is the more important analysis since it rules out potential confounds arising from physical differences between famous and unfamiliar faces, or between liked and disliked famous faces. The ANOVA was marginally significant in the participants analysis, offering weak support for the main items analysis.

The factors of strategy and task half showed no reliable effects, in particular failing to interact with valence, providing no support for either the strategy or the practice explanations. Accuracy of responses to disliked faces was not lower under the Pop strategy than under the Look strategy. Accuracy of responses to disliked faces did not decline with practice, revealed by the null effect of task half.

A potential confound must be addressed. Response latency was manipulated by varying the duration of the backward mask, so it is possible that the mask duration rather than the response latency was responsible for the pattern of results. Perhaps the 100 ms backward mask was less effective than the 500 ms backward mask, and a less effective mask may have allowed valence to become activated, so that valence influenced participant responses with 100 ms backward mask but not with 500 ms backward mask. This must be regarded as unlikely for three distinct reasons. First, the backward mask duration was 100 ms throughout Experiment 1 and yet there was an interaction of response latency with valence, suggesting that response latency rather than backward mask duration is the key factor. Second, Esteves and Ohman (1993) reported that the duration of the backward mask (between 30 and 120 ms) had no effect on the likelihood of conscious perception of a masked facial emotional expression. Third, in the present experiment, participants were more likely to be aware of the identity of a masked 17 ms face at 500 ms mask duration (item mean percentage conscious recognition = 2.7%) than with 100 ms mask duration (item mean = 0.2%), Wilcoxon signed-ranks $z_i(29) = 2.69$, $p < 0.01$, and Mann–Whitney $z_p(37) = 2.42$, $p < 0.05$, which contradicts the proposal that masking is more effective at 500 ms than 100 ms mask duration.

Experiment 2 supported the response latency account but not the practice account or the strategy account. It therefore appears that neither practice nor strategy significantly moderated the effect of lower response accuracy to disliked than liked faces, and the only relevant factor was response latency from face onset. This is consistent with the affective priming literature in which effects of valence are swiftly activated and swiftly corrected. It is supposed that the effect of negative valence was corrected because it led participants to make incorrect responses by selecting the paired unfamiliar face instead of the famous face.

Experiment 3

A new affective priming task was introduced to investigate the possibility that the higher response accuracy to negative faces than positive faces that was observed on longer latency trials in Experiment 1 was due to over-correction for the negative valence of the famous face. The affective priming task was given to a subset of the participants in Experiment 2 and was always performed after the original task and before participant evaluations of the famous persons.

Prime stimuli in Experiment 3 were masked 17 ms faces of famous persons. Targets were clearly visible words of pleasant or unpleasant meaning, and the required response was a pleasant–unpleasant decision to the target word. Normal affective priming suggests facilitation of responses to target words of the same valence as the prime face, that is, liked faces should result in faster responses to pleasant words and disliked faces in faster responses to unpleasant words. However, if the biasing effect of the prime face were corrected, there might be no interaction of prime and target valence. Further, if the biasing effect of the prime face were overcorrected, then the result would be faster responses to targets of valence incongruent with the prime face. It is necessary to examine closely the conditions that might be required for correction to occur.

The first condition is that the valence evoked by the prime face should be recognised as misleading with respect to the required response to the target. The required target response was an evaluative decision (pleasant or unpleasant), and the prime and target valence were congruent and incongruent on approximately equal numbers of trials, so the prime valence was in fact misleading with respect to the required response to the target. The question is whether this would be recognised. In this respect, a complication is introduced by the different strategies used in Experiment 2. Given that Experiment 2 always preceded Experiment 3 in the same experimental session, and evidence that the effect of strategy can carry forward into a subsequent task (e.g., Carr & Dagenbach, 1990; Dagenbach et al., 1989; Kahan, 2000), it is necessary to examine whether the Look and Pop strategies predict equivalent correction for the biasing effect of the prime face.

The Look strategy in the original task instructed participants to look hard at the stimulus faces and use any partial perception as the basis for their responses. This instruction seems likely to have focused deliberate attention on the masked faces and it may be supposed that this focus of attention carried forward into Experiment 3 (e.g., Carr & Dagenbach, 1990; Dagenbach et al., 1989; Kahan, 2000). So in Experiment 3, the combination of deliberate attention to the masked faces, the instruction to use them as a basis for the response, and equal numbers of valence-congruent and valence-incongruent trials, should have led to the recognition that prime valence was misleading for the required target response. This should have prompted an automatic correction process. In contrast, the Pop strategy asked participants to relax, not make any effort to see the faces, and let the answer pop into their heads. In Experiment 3, this instruction seems likely to have removed the deliberate focus of attention from using the masked faces as a basis for the response, in which case, the irrelevance of the masked faces to the instructed task would have been less obvious. Hence, correction for the effect of prime valence was less likely to occur. This line of reasoning predicts that reverse affective priming (faster responses to target words of valence incongruent with the prime face) should be stronger under the Look strategy than under the Pop strategy.

The second condition for correction of the biasing effect of prime valence is sufficient time for the correction to occur. The theorised correction occurred in the original task of Experiment 1 on long latency trials but not on short latency trials. However, faces were presented in simultaneous pairs in Experiment 1, and were presented singly and centrally in Experiment 3. It is quite possible that correction could occur more quickly under the easier conditions of Experiment 3. Given that Glaser and Banaji (1999) reported reverse affective priming with SOA of only 150 ms, it is possible that correction could occur even at the shorter prime-target SOA of 100 ms in the present Experiment 3. Therefore, no prediction was made for the effect of backward mask duration.

Method

Participants

Participants were 44 undergraduate students at Goldsmiths College, London, comprising all except the first 20 from Experiment 2. Data from 11 participants were not analysed for reasons detailed below in the Results section. The remaining 33 participants were 26 female and seven male, aged between 18 and 41, mean 22.9, SD 6.7 years. There were 17 participants in the 100 ms backward mask condition and 16 participants in the 500 ms backward mask condition; 18 participants in the Look strategy condition and 15 participants in the Pop condition.

Stimuli

Ten of the faces used in Experiment 2 were selected as primes on the basis that they were correctly identified by over 90% of the first group of participants in Experiment 2, with roughly equal numbers of participants evaluating each face as liked and disliked. Names are presented in Appendix A. Another five unfamiliar faces were chosen to act as primes on filler trials. The mask was the same as used in Experiment 2.

The targets were four words: rose, pleasure, pain, and fighting. The two positive and two negative targets were balanced on word length and word frequency. The targets were those used by Snodgrass et al. (1993).

Design

There were three independent factors. Backward mask duration (100 vs. 500 ms) and strategy (Look vs. Pop) were between-participants and within-items, and congruence (face-valence congruent or incongruent with word-valence) was within-participants and within-items. Backward mask duration for each participant was the same as Experiment 2. Strategy was implemented by the instructions given to participants in Experiment 2 that always preceded Experiment 3. Each famous face was rated as liked or disliked by each participant after the affective priming task, by the procedure described in Experiment 2.

Each prime face was presented four times, once for each target word, making a total of 60 trials. Trials were presented in a different random sequence for each participant. The required response to the target was a two-alternative forced-choice of pleasant or unpleasant. The keys assigned to the two response options were counterbalanced across participants.

Procedure

Participants carried out the affective priming task following the original awareness check task, and before giving their evaluations of the famous persons. Four practice trials preceded the 60 experimental trials. The participant initiated the sequence of trials by pressing a key. The presentation of stimuli on each trial was as follows: forward mask for 100 or 500 ms, prime face for 17 ms, backward mask for 100 or 500 ms, and target word presented until the participant responded. Each subsequent trial was initiated by the response to the previous trial after an interval of 1 s.

Participants were informed that a series of words, some pleasant and some unpleasant, would be displayed one at a time on the screen. They were asked to respond by pressing one of two keys depending on whether the target word was pleasant or unpleasant and to respond as quickly as possible. The assignment of keys to responses was counterbalanced across participants. They were asked to report immediately any faces that were consciously recognised, even if recognition was very uncertain and even if the name could not be recalled. After completion of the experimental trials, participants gave their evaluations of the famous persons, as described in Experiment 2. Finally, participants were debriefed and thanked for their participation.

Results and discussion

Trials were excluded if the response to the target word was incorrect, if the response was slower than 1500 ms, if the participant was aware of the identity of the prime face during the experimental trials, or if the participant could not identify the prime face in the post-experimental rating. Eleven participants had only a single face rated liked or disliked of whose identity they were unaware during the experimental trials, and were therefore excluded from the analysis. Other data were considered valid and were included in the analysis.

ANOVA was performed by participants and by items with three factors of backward mask (500 vs. 100 ms), strategy (Look vs. Pop) and congruence (face-valence congruent or incongruent with word-valence). The dependent variable was mean response time for correct responses to target words. The approach of collapsing the two factors of word-valence and face-valence into a single factor of congruence is a common approach (e.g., De Houwer & Hermans, 1994; De Houwer et al., 1998; Glaser & Banaji, 1999; Hermans et al., 2001; Klauer et al., 1997; Musch & Klauer, 2001). This approach serves to enhance the clarity of the results and discussion by focusing on the experimental hypothesis. Exposition of any main effects of face-valence and word-valence that might be present would have detracted from the major hypothesis of the experiment.

The main effect of backward mask duration was significant, $F_i(1, 9) = 124$, $p < 0.001$, and $F_p(1, 29) = 24.9$, $p < 0.001$, showing slower responses in the 500 ms backward mask condition than in the 100 ms condition. The two-way interaction of strategy with congruence was significant, $F_i(1, 9) = 6.91$, $p < 0.03$, and $F_p(1, 29) = 7.03$, $p < 0.02$. All other effects were non-significant, all F_i and $F_p < 1$. The interaction was investigated with paired-samples t tests for each strategy. Under the Pop strategy, there was a tendency to faster responses in the congruent condition, $t_i(9) = 1.77$, ns, and $t_p(14) = 1.80$, ns, that is, normal affective priming. Under the Look strategy, there was a tendency to faster responses in the incongruent condition, $t_i(9) = 2.09$, $p < 0.07$, and $t_p(17) = 2.17$, $p < 0.05$, that is, reverse affective priming. See Table 3 and Figure 3.

Table 3: mean response time (and standard error) for correct responses to valence-congruent and valence-incongruent target words, under the Pop and Look strategies, with 100 and 500 ms backward mask, in Experiment 3.

	Look		Pop	
	Congruent	Incongruent	Congruent	Incongruent
100 ms mask	639 (19)	602 (15)	609 (14)	643 (9)
500 ms mask	813 (32)	785 (23)	803 (21)	825 (17)
Mean	726 (21)	693 (14)	706 (13)	734 (11)

The pattern of results supports the prediction that reverse affective priming would be stronger under the Look strategy than under the Pop strategy. This is consistent with the concept that the Look strategy would focus attention on the misleading nature of the prime's valence, and so engender a process of correction for the biasing effect of the prime, while the Pop strategy would focus less attention on the prime, and so engender a weaker (or no) correction process. The observation of this result in the item analysis rules out confounds based on variation in some visual characteristics of the faces. Differential priming of valenced target words according to whether the same prime face was liked or disliked (by different participants) confirms that the effect was due to recognition of the unique face identity.

Figure 3: mean response time in the valence-congruent and valence-incongruent conditions of Experiment 3 (affective priming), under the Look and Pop strategies. Bars represent standard errors.

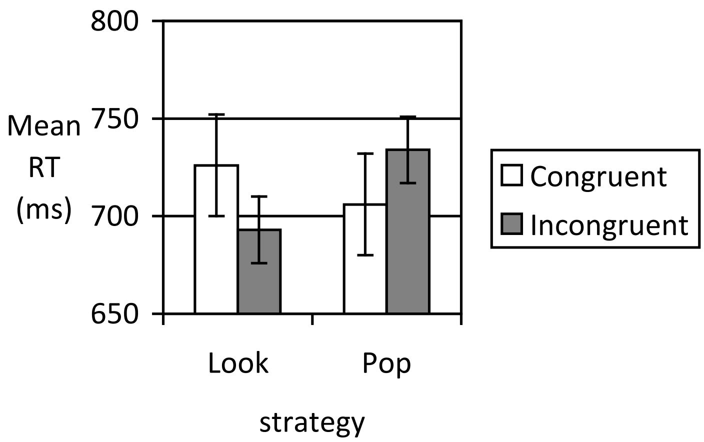

The interaction of the direction of affective priming (normal or reverse) with strategy was unaffected by the duration of the backward mask. It appears that the theorised correction for the biasing effect of prime valence that occurred under the Look strategy was as effective at the shorter backward mask duration (SOA of 117 ms) as at the longer backward mask duration (SOA of 517 ms). This might seem inconsistent with the results of Experiment 2, that reported an apparent correction for the biasing effect of negative valence at 500 ms backward mask duration but not at 100 ms mask duration, but it should be noted that Experiment 2 presented faces in simultaneous pairs while Experiment 3 presented faces singly and centrally. It is reasonable to suppose that processing of two faces simultaneously would be slowed in comparison with processing of a single face, so that the theorised over-correction could have occurred with the shorter backward mask duration in Experiment 3 but not in Experiment 2. Another issue is that the over-correction for the biasing effect of prime valence must have occurred before the response on each trial was decided, but it is not clear at exactly what point this happened. In Experiment 3, the response to the target word was formulated some time after the word was presented, but the lag is unknown. The lag would have allowed some extra time for the biasing effect of prime valence to be corrected.

General discussion

Experiment 1 was an attempt to replicate the effect observed by Stone and Valentine (2004) that detection of familiarity from masked 17 ms faces was more accurate when the famous persons were subsequently evaluated positively than when they were evaluated negatively. A different pattern of results was observed, with the accuracy of responses to negative faces declining throughout the task. Feedback from participants during debriefing and theoretical considerations suggested three potential explanations: practice, strategy and response latency.

The results of Experiment 2 supported only the response latency account. With short response latency (100 ms backward mask duration), responses were less accurate to disliked faces than to liked faces, while at long response latency (500 ms backward mask duration), responses were equally accurate to liked and disliked faces. Responses to disliked faces were less accurate at short than at long response latency. There was no effect of practice or strategy. The lower accuracy of responses to negative faces at short response latency was attributed to the biasing effect of negative valence that caused participants to select the paired unfamiliar face, contrary to task instruction. With longer latency, the biasing effect of negative valence was corrected, so that responses to the faces of negative persons were as accurate as responses to positive faces.

A speculative explanation for the effect of response latency can be offered, as follows. It is necessary to consider first how facial familiarity is detected. Farah, O’Reilly, and Vecera (1993) suggested that facial familiarity might be detected when any semantic information associated with the face becomes activated. There is an alternative view: the Burton, Bruce, and Johnston (1990) model of face recognition states that familiarity is detected when activation at the Person Identity Node reaches a threshold. This node is an a-modal representation of the concept of the person, and does not itself store any semantic information, but is connected to other nodes representing semantics. The concept that familiarity is detected when the Person Identity Node reaches a threshold was based on the observation of a ‘‘familiar-only’’ response that occurs when a face can be deemed familiar but the participant is not aware of any semantic information. However, there is the possibility that several items of semantic information might be activated, each one below threshold for awareness, but together adding up to sufficient activation to declare the face familiar. So, even when the familiar-only response occurred, familiarity may still have been detected from the activation of semantic information. Also, the participant might have access to affective information concerning their attitude towards the person, and this might be sufficient to declare the face familiar. On balance, it seems likely that detection of facial familiarity depends on the (maybe unaware) retrieval of semantic or affective information associated with the person.

Given that facial familiarity may be detected from the existence of any semantic or affective information associated with the person, it is relevant to consider the relative timescale of retrieval of these two types of information. Zajonc’s affective primacy hypothesis (Murphy & Zajonc, 1993; Zajonc, 1980; see Zajonc, 2001, for a review) states that affective information becomes available before semantics. Some support for this hypothesis stems from Bargh et al. (1989) who reported that the valence of masked stimuli could be detected with a shorter exposure duration than semantic meaning, suggesting that valence may be more strongly connected and more readily activated than semantic information. If this is the case, then valence may well become available before semantics. It should also be noted that affective priming has been observed with zero or negative stimulus onset asynchrony (SOA) between prime and target (e.g., De Houwer & Hermans, 1994; Hermans et al., 2001; Klauer et al., 1997; Musch & Klauer, 2001) whereas semantic priming requires a positive SOA. All of these lines of evidence suggest that affective information about a stimulus is activated more swiftly than semantic information.

Applied to the present series of experiments, the affective primacy hypothesis has the implication that valence exerted a stronger influence on shorter latency responses than on longer latency responses. On shorter latency responses, negative valence associated with the famous person led participants to select the unfamiliar face, and semantics were not sufficiently activated to contradict this decision. On the longer latency trials, semantic information pointed at the famous face, creating a conflict between the valence information and the semantic information. Assuming that a decision mechanism placed more weight on the semantic information, this could explain why responses to negative-disliked faces were at least as accurate as responses to positive-liked faces at longer response latency. The assumption that a decision mechanism placed more weight on semantics than affect seems reasonable - given the perception that a face generates a negative valence, but at the same time the face is known to have, e.g., a certain occupation, it is obvious that the face belongs to a famous person.

It remains to be explained why negative valence associated with a famous person should lead to the selection of the paired unfamiliar face. One possible mechanism is based on the mis-application of the mere exposure effect (Zajonc, 1980; see Zajonc, 2001, for a review), which states that a familiar stimulus is preferred over an unfamiliar stimulus, all other things being equal. Participants may have conceived a preference for one face in each pair, assumed that preference indicated familiarity, and so selected the preferred face as likely to be the famous face. Alternatively, participants might have simply selected the face they preferred without any assumption that preference denotes familiarity. Either way, when the famous person was negative-disliked, the unfamiliar face was preferred and so was selected.

Preference could have been detected in many ways. In an fMRI study conducted by Pizzagalli, Koenig, Regard, and Lehmann (1998), participants were asked to observe previously unfamiliar faces for 450 ms each, and subsequently rate the faces as liked or disliked. Liked and disliked faces activated different neuronal populations, with the centre of activation for disliked faces being more to the right than the centre of activation for liked faces. Pizzagalli et al. (2002) reported that the N170 ERP component that has been frequently associated with face processing was larger in amplitude for liked than disliked faces. Another possibility is that liked and disliked faces had different effects on the amygdala, e.g., Whalen et al. (1998) observed that fearful facial expressions perceived without awareness of the expression resulted in higher amygdale activation than happy facial expressions. These are just three possibilities: any mechanism for detecting preference between two faces could have generated the observed effect in the present experiments.

The observation of more accurate responses to negative than positive faces in the longer latency trials of Experiment 1 suggests that the biasing effect of negative valence might not merely be corrected when semantics were also available, but actually over-corrected. This could occur if the familiarity decision mechanism, assumed to place more weight on semantics than valence, had corrected the valence signal to make it consistent with the semantic signal. If this correction was strong then a negative face might generate a stronger preference than a positive face on the longer latency trials. Combined with the semantic information pointing at the famous face, an over-corrected preference signal could have generated the observed effect of more accurate familiarity detection from negative than positive faces.

Experiment 3 examined the possibility of over-correction for the biasing effect of valence in an affective priming paradigm, using masked 17 ms famous faces as primes and clearly visible pleasant or unpleasant words as targets. Results supported the concept of over-correction when participants' strategy in attempting to perceive the masked faces caused recognition of the misleading effect of prime valence for the required response to the target.

Several directions for future research are suggested by the experiments reported here. The speed with which the valence associated with a famous person becomes active, and may be subsequently corrected, could be estimated by including various backward mask durations. The conditions of strategic set under which correction is applied could also be systematically examined and may shed light on the mechanism of correction. It would be of interest to invite prosopagnosic participants to perform modified versions of these experiments: given ample evidence of covert face recognition in the form of activation of semantic information, it seems likely that prosopagnosics would be able to activate emotional information associated with famous faces of whose identity there was no awareness.

Other factors may have influenced the accuracy of participant choices, for example, face familiarity, attractiveness and distinctiveness, and these could be examined by collecting ratings from experimental participants. The experiments reported here could all be repeated with other masked stimuli; one obvious choice would be names rather than faces of famous and unfamiliar persons.

Several conclusions may be drawn. The most basic is that famous faces presented very briefly and masked can be identified without awareness of facial familiarity. A more interesting conclusion is that negative participant attitude towards the famous person has different effects on a familiarity decision at different response latencies: familiarity is less likely to be detected a shorter latency than at longer latency. Consistent with the affective priming literature, it appears that affective responses are influenced by temporal factors.

Appendix A. Stimuli

Experiment 1

Film/TV actors: Woody Allen, Gillian Anderson, Jennifer Aniston, Sandra Bullock, Michael Caine, Jim Carrey, Sean Connery, Kevin Costner, Tom Cruise, Jamie Lee Curtis, Robert deNiro, Leonardo Dicaprio, Michael Douglas, David Duchovny, Harrison Ford, Sarah Michelle Gellar, Richard Gere, Mel Gibson, Hugh Grant, Anthony Hopkins, Liz Hurley, Bruce Lee, Marilyn Monroe, Demi Moore, Roger Moore, Jack Nicholson, Brad Pitt, Oliver Reed, Burt Reynolds, Julia Roberts, Arnold Schwarzenegger, William Shatner, Sylvester Stallone, Sharon Stone, Liz Taylor, John Travolta, Bruce Willis, Catherine Zeta-Jones.

Pop stars: Damon Albarn, Victoria Beckham, Liam Gallagher, Geri Halliwell, Janet Jackson, Michael Jackson, Mick Jagger, Elton John, John Lennon, Jennifer Lopez, Madonna, Freddie Mercury, George Michael, Kylie Minogue, Elvis Presley, Cliff Richard, Robbie Williams.

Comedians: Rowan Atkinson, Craig Charles, Martin Clunes, Harry Enfield, Stephen Fry, Joanna Lumley, Nicholas Lyndhurst, Neil Morrissey.

Royal family: Sarah Ferguson, Prince Andrew, Prince Charles.

Politicians: Gerry Adams, Jeffrey Archer, Cherie Blair, Tony Blair, George W. Bush, Bill Clinton, William Hague, Adolf Hitler, Saddam Hussein, J.F. Kennedy, Ken Livingston, John Major, John Prescott, Ronald Reagan, Margaret Thatcher.

TV presenters: Jeremy Beadle, Cilla Black, Paul Daniels, Charlie Dimmock, Chris Evans, Judy Finnigan, Bruce Forsyth, Rolf Harris, Clive James, Des Lynam, Richard Madeley, Michael Parkinson, Anne Robinson, Chris Tarrant, Alan Titchmarsh, Carol Vordeman, Terry Wogan.

Sports: Boris Becker, David Beckham, Paul Gascoigne, O.J. Simpson, Mike Tyson.

Other: Richard Branson (entrepreneur), Stephen Hawking (scientist), Osama Bin Laden.

Faces excluded from analysis: Charles Bronson, Kathy Burke, John Cleese, Glenn Close, Jodie Foster, Dawn French, Tom Hanks, Woody Harrelson, Lenny Henry, Myra Hindley, David Jason, Hugh Laurie, Paul Mccartney, Eddie Murphy, Liam Neeson, Michael Palin, Michelle Pfeiffer, Jennifer Saunders, Robin Williams, Princess Diana.

Experiment 2

Film/TV actors: Russell Crowe, Tom Cruise.

Pop stars: Liam Gallagher, Geri Halliwell, Whitney Houston, Michael Jackson, Victoria Beckham, Cliff Richard, Britney Spears, Robbie Williams.

Royal family: Sarah Ferguson, Camilla Parker-Bowles, Prince Charles.

Politicians: Tony Blair, G.W. Bush, William Hague, John Major, Colin Powell, John Prescott, Margaret Thatcher, Anne Widdecombe.

TV presenters: Michael Barrymore, Jeremy Beadle, Bill Clinton, Paul Daniels, Chris Evans, Jeremy Paxman, Anne Robinson.

Other: Naomi Campbell (model).

Faces excluded from analysis: Woody Allen, Yasser Arafat, David Beckham, Cilla Black, Cherie Blair, Helena Bonham-Carter, Gordon Brown, Jim Carrey, Cher, Puff Daddy, Danny DeVito, Leonardo DiCaprio, David Duchovny, Clint Eastwood, Eminem, Bruce Forsyth, Sarah-Michelle Gellar, Hugh Grant, Rutger Hauer, Liz Hurley, Janet Jackson, Samuel L. Jackson, Mick Jagger, Elton John, Ross Kemp, Annie Lennox, Ken Livingston, Jennifer Lopez, Richard Madeley, Rik Mayall, Ally McBeal, Ewan McGregor, George Michael, Liza Minelli, Mike Myers, Leonard Nimoy, Gwynneth Paltrow, Luciano Pavarotti, Alan Rickman, Jonathan Ross, Arnold Schwarzenegger, O.J. Simpson, Iain Duncan Smith, Sylvester Stallone, Patrick Stewart, Barbra Streisand, Mike Tyson, Oprah Winfrey, Terry Wogan, Princess Anne, The Queen.

Experiment 3

All faces included: Michael Barrymore, Tony Blair, Naomi Campbell, Chris Evans, Liam Gallagher, Geri Halliwell, John Major, Cliff Richard, Anne Robinson, Margaret Thatcher.

Appendix B. Examples of stimuli and the mask

References

Banse, R. (1999). Automatic evaluation of self and significant others: Affective priming in close relationships. *Journal of Social and Personal Relationships*, 3, 803–821.

Banse, R. (2001). Affective priming with liked and disliked persons: Prime visibility determines congruency and incongruency effect. *Cognition and Emotion*, 15, 201–220.

Bargh, J. A., Chaiken, S., Raymond, P., & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology*, 32, 104–128.

Bargh, J. A., Litt, J., Pratto, F., & Spielman, L. A. (1989). On the preconscious evaluation of social stimuli. In A. F. Bennett & K. M. McConkey (Eds.), *Cognition in individual and social contexts*. North-Holland: Elsevier.

Burton, A. M., Bruce, V., & Johnston, R. A. (1990). Understanding face recognition with an interactive activation model. *British Journal of Psychology*, 81, 361–380.

Carr, T. H., & Dagenbach, D. (1990). Semantic priming and repetition priming from masked words: Evidence for a centre-surround attentional mechanism in perceptual recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16, 341–350.

Dagenbach, D., Carr, T. H., & Wilhelmsen, A. L. (1989). Task-induced strategies and near-threshold priming: Conscious influences on unconscious perception. *Journal of Memory and Language*, 28, 412–443.

De Houwer, J., & Hermans, D. (1994). Differences in the affective processing of words and pictures. *Cognition and Emotion*, 8, 1–20.

De Houwer, J., Hermans, D., & Eelen, P. (1998). Affective and identity priming with episodically associated stimuli. *Cognition and Emotion*, 12, 145–169.

Dimberg, U., & Ohman, A. (1996). Behold the wrath: Psychophysiological responses to facial stimuli. *Motivation and Emotion*, 20, 149–182.

Dimberg, U., Thunberg, M., & Elmehed, K. (2000). Unconscious facial reactions to emotional facial expressions. *Psychological Science*, 11, 86–89.

Esteves, F., & Ohman, A. (1993). Masking the face: Recognition of emotional facial expressions as a function of the parameters of backward masking. *Scandinavian Journal of Psychology*, 34, 1–18.

Farah, M. J., O_Reilly, R. C., & Vecera, S. P. (1993). Dissociated overt and covert recognition as an emergent property of a lesioned neural-network. *Psychological Review*, 100, 571–588.

Fazio, R. H. (2001). On the automatic activation of associated evaluations: An overview. *Cognition and Emotion*, 15, 115–141.

Glaser, J., & Banaji, M. R. (1999). When fair is foul and foul is fair: Reverse priming in automatic evaluation. *Journal of Personality and Social Psychology*, 77, 669–687.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102, 4–27.

Hermans, D., de Houwer, J., & Eelen, P. (1994). The affective priming effect: Automatic activation of evaluative information in memory. *Cognition and Emotion*, 8, 515–533.

Hermans, D., De Houwer, J., & Eelen, P. (2001). A time course analysis of the affective priming effect. *Cognition and Emotion*, 15, 143–165.

Johnsen, B. H., & Hugdahl, K. (1991). Hemispheric asymmetry in conditioning to facial emotional expressions. *Psychophysiology*, 28, 154–162.

Johnsen, B. H., & Hugdahl, K. (1993). Right hemisphere representation of autonomic conditioning to facial emotional expressions. *Psychophysiology*, 30, 274–278.

Kahan, T. A. (2000). Negative priming from masked words: Retrospective prime clarification of centre-surround inhibition? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 26, 1392–1410.

Klauer, K. C., Rossnagel, C., & Musch, J. (1997). List-context effects in evaluative priming. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 23, 246–255.

Mogg, K., & Bradley, B. P. (1999). Orienting of attention to threatening facial expressions presented under conditions of restricted awareness. *Cognition and Emotion*, 13, 713–740.

Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology*, 64, 723–739.

Musch, J., & Klauer, K. C. (2001). Locational uncertainty moderates affective congruency effects in the evaluative decision task. *Cognition and Emotion*, 15, 167–188.

Niedenthal, P. M. (1990). Implicit perception of affective information. *Journal of Experimental Social Psychology*, 26, 505–527.

Ohman, A., Esteves, F., & Soares, J. J. F. (1995). Preparedness and pre-attentive associative learning—Electrodermal conditioning to masked stimuli. *Journal of Psychophysiology*, 9, 99–108.

Pizzagalli, D., Koenig, T., Regard, M., & Lehmann, D. (1998). Faces and emotions: Brain electric field sources during covert emotional processing. *Neuropsychologia*, 36, 323–332.

Pizzagalli, D., Lehmann, D., Hendrick, A. M., Regard, M., Pascual-Marqui, R. D., & Davidson, R. J. (2002). Affective judgements of faces modulate early activity (_160 ms) within the fusiform gyri. *Neuroimage*, 16, 663–677.

Robinson, M. D. (1998). Running from William James' bear: A review of preattentive mechanisms and their contributions to emotional experience. *Cognition and Emotion*, 12, 667–696.

Saban, S., & Hugdahl, K. (1999). Nonaware classical conditioning to pictorial facial stimuli in a between-groups paradigm. *Integrative Physiological and Behavioural Science*, 34, 19–29.

Snodgrass, M., Shevrin, H., & Kopka, M. (1993). The mediation of intentional judgments by unconscious perceptions: The influences of task strategy, task preference, word meaning, and motivation. *Consciousness and Cognition*, 2, 169–193.

Stapel, D. A., Koomen, W., & Zeelenberg, M. (1998). The impact of accuracy motivation on interpretation, comparison, and correction processes: Accuracy · knowledge accessibility effects. *Journal of Personality and Social Psychology*, 74, 878–893.

Stone, A. M., & Valentine, T. (2004). Better the devil you know? Non-conscious processing of identity and affect of famous faces. *Psychonomic Bulletin & Review*, 11, 469–474.

Stone, A. M., & Valentine, T. (in press). Orientation of attention to non-consciously recognised famous faces. *Cognition and Emotion*.

Stone, A. M., Valentine, T., & Davis, R. (2001). Face recognition and emotional valence: Processing without awareness by neurologically intact participants does not simulate covert recognition in prosopagnosia. *Cognitive, Affective and Behavioural Neuroscience*, 1, 183–191.

Strack, F., Schwarz, N., Bless, H., Kubler, A., & Wanke, M. (1993). Awareness of the influence as a determinant of assimilation versus contrast. *European Journal of Social Psychology*, 23, 52–62.

Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., & Jenike, M. A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *Journal of Neuroscience*, 18, 411–418.

Wong, P. S., Shevrin, H., & Williams, W. J. (1994). Conscious and nonconscious processes: An ERP index of an anticipatory response in a conditioning paradigm using visually masked stimuli. *Psychophysiology*, 31, 87–101.

Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35, 151–175.

Zajonc, R. B. (2001). Mere exposure: A gateway to the subliminal. *Current Directions in Psychological Science*, 10, 224–228.