COGNITION AND EMOTION, 2004, *18* (8), 1009–1026

Unequally masked: Indexing differences in the perceptual salience of ''unseen'' facial expressions

Jeffrey S. Maxwell and Richard J. Davidson
University of Wisconsin – Madison, USA

Two experiments investigated the degree to which different facial expressions, when backwardly masked by neutral faces, are blocked from different levels of perceptual access as indexed by explicit (self-reported) awareness, forced-identification performance, and stimulus sensitivity in a signal detection paradigm. Results indicate that a 16.67 ms target-mask stimulus onset asynchrony (SOA) is insufficient to reliably block either the simple detection or the forced-identification of backwardly masked expressions. This finding holds even for participants who deny possessing any explicit awareness of the masked images whatsoever. Furthermore, ''unseen'' happy faces are less effectively masked by neutral expressions than are unseen angry or neutral faces, an observation that might also extend to other affective expressions (e.g., fear). These results expand upon previous findings (Esteves & Öhman, 1993) and provide new information about the multitude of factors that may be operating in research that involves backwardly masked facial expressions.

The number of elements to which an individual can simultaneously attend and the quantity of items that he or she can juggle at any moment is markedly limited (Holender, 1986; Pashler, 1998). As such, when the mind is preoccupied or a great deal of stimulation bombards the senses, the majority of incoming information necessarily remains outside of the ken of explicit awareness. Several investigations have suggested that, in many cases, such information is able to directly influence cognition and affectively guided processing and behaviour even while it remains isolated from an individual's phenomenal experience (Holender, 1986; Mathews & MacLeod, 1986).

A sizeable body of clinical and behavioural evidence has now accumulated in support of this notion (e.g., Bechara, Damasio, & Damasio, 2000; De Houwer,

Correspondence should be addressed to Richard J. Davidson, Department of Psychology, University of Wisconsin, 1202 West Johnson Street, Madison, Wisconsin 53706, USA; e-mail: rjdavids@facstaff.wisc.edu; or email: maxwell@psyphw.psych.wisc.edu

We thank Paul Whalen, Alexander Shackman, Philip Merikle, and two anonymous reviewers for providing valuable feedback, Sarah Lindberg for assistance with image processing and data acquisition, and Adrian Pederson for computer programming.

http://www.tandf.co.uk/journals/pp/02699931.html DOI:10.1080/02699930441000003

Hendrickx, & Baeyens, 1997; Merikle & Reingold, 1990; Meyer & Waller, 1999; Murphy & Zajonc, 1993; Seitz & Watanabe, 2003; Winkielman, Zajonc, & Schwarz, 1997). Additional evidence that demonstrates the processing and influence of information presented outside of awareness comes from the psychophysiological (Bunce, Bernat, Wong, & Shevrin, 1999; Esteves, Dimberg, & Öhman, 1994a; Öhman & Soares, 1998), electrophysiological (Bernat, Shevrin, & Snodgrass, 2001; Wong, Bernat, Bunce, & Shevrin, 1997; Wong, Shevrin, & Williams, 1994), and neuroimaging (Berns, Cohen, & Mintun, 1997; Dehaene et al., 1998; Whalen et al., 1998; Morris, Öhman, & Dolan, 1998) domains.

One particularly useful technique for investigating processing without awareness involves the use of backward visual masking. Backward visual masking entails the presentation of a masking stimulus in such a way as to disrupt the processing of a previously presented target image. Given the increasing and widespread interest in clarifying the roles of awareness and automatic processes in the appraisal and generation of emotion (Robinson, 1998; LeDoux, 1996; Niedenthal, 1990; Mogg & Bradley, 1999) and autonomic activity (Saban & Hugdahl, 1999; Wong et al., 1994; Öhman & Soares, 1998), backward visual masking has proved extremely valuable for investigations in the affective domain.

Specifically, the rate of appearance for publications using masked facial expressions is rising, due in part to an increasing interest in examining automatic processing with electrophysiological and haemodynamic techniques. Face images are notably suited for paradigms of this sort because emotional expressions possess cross-cultural validity (Ekman, Sorenson, & Friesen, 1969), are biologically grounded (Dimberg, Thunberg, & Elmehed, 2000), and are ecologically appealing as cues for generating affective evaluations and responses (Dimberg et al., 2000; Öhman, 1986). The purpose of this article is to present new empirical data concerning backwardly masked facial expressions and to discuss their potential implications.

Several different opinions exist over what the minimum adequate requirements are for determining that a stimulus is processed outside of awareness (Bernat et al., 2001; Holender, 1986; Reingold & Merikle, 1988). In the present discussion, *explicit awareness* is comparable to the definition of awareness adopted by Wiens and Öhman (2002) and used to characterise whether an individual is able to report having seen and recognised a given stimulus. Specifically, in the current investigation, we classified a participant as being explicitly aware if he/she reported seeing one or more of the backwardly masked faces during the course of the experiment. Individuals falling on the threshold of explicit awareness, who reported seeing *something* associated with the backward visual masking (e.g., an image flickering) but were unable to specifically determine that it was a face, have been excluded from the first experiment in this investigation. An individual is classified here as being *explicitly unaware* only if he/she reported seeing *nothing* associated with the backward visual masking whatsoever. Note, however, that the complete absence of explicit awareness does not mean that such processing occurs exclusively outside of the domain of higher order cognition or that this information is inaccessible through other perceptual measures (e.g., forced-choice discrimination), which may or may not index other aspects of awareness. It merely indicates that an individual is unable or unwilling to report having any phenomenal experience of that stimulus' existence or identity (Lovibond & Shanks, 2002; Shanks & Lovibond, 2002).

If the question regards only whether, and how, information about which an individual is not explicitly aware is able to influence thought and action, then establishing lack of explicit awareness with subjective self-report may be sufficient. However, intersubject variability in perceptual thresholds for this information has the potential to introduce considerable noise into the data and may unintentionally influence the experimental results. For example, Murphy and Zajonc (1993) have suggested that when affect is elicited in the absence of conscious awareness, it is more readily attributed to unrelated neutral stimuli than when the source of that emotional activation is readily apparent. And more recently, Rotteveel and colleagues (2001) have reported obtaining stronger affective priming effects using suboptimal (less conscious) than optimal (conscious) emotional cues. If these observations are correct, then those individuals who possess a particularly keen faculty for registering backwardly masked stimuli may correspondingly yield small or no treatment effects and attenuate the overall pattern of experimental results. Similarly, if differences exist in the average participant sensitivity across separate study conditions or independent investigations, the ability to observe legitimate differences between experimental conditions or to replicate previously reported results might be compromised. Therefore, in addition to ensuring that backwardly masked stimuli fall outside of explicit awareness, setting and matching the relative perceptibility of distinct backwardly masked images across individual participants may be a profitable endeavour.

Equally and perhaps more importantly, if one wishes to directly contrast the processing of distinct objects that are presented outside of explicit awareness within individual participants, a more sensitive test of perceptual accessibility should be employed to ensure that these stimuli have been matched on their relative degree of perceptual degradation. Failing to match stimuli in this manner allows them to potentially vary on two dimensions: The primary experimental manipulation of interest and the tangential dimension of perceptual salience. In other words, the scale of explicit (self-reported) awareness may be insufficiently sensitive to categorically determine whether differences exist in the perceptibility of different backwardly masked visual stimuli. Solely relying upon this measure, therefore, opens up the distinct possibility of misidentifying sources of variance in one's data (Figure 1) and failing to detect potentially meaningful individual differences in perceptual sensitivity for distinct masked stimuli.

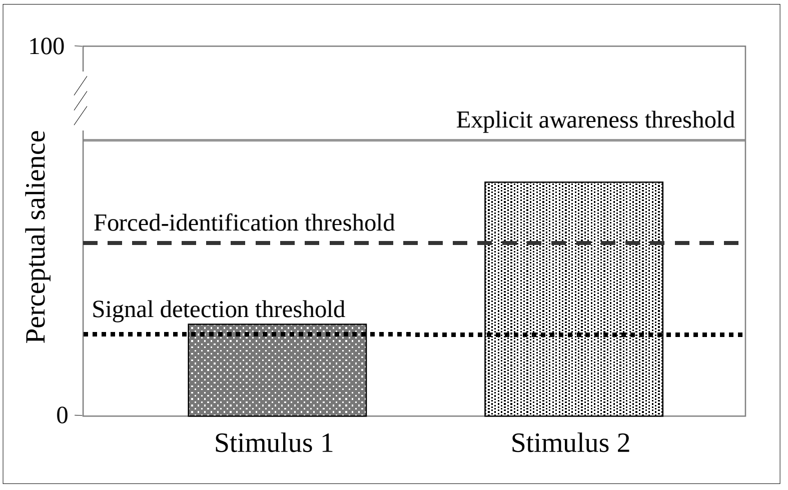

Figure 1. Hypothetical illustration of how two ''unseen'' stimuli (*x*-axis) that each fall below the threshold for explicit (self-reported) awareness can possess a different degree of perceptual salience from one another (*y*-axis; arbitrary units). Furthermore, note that alternative tests of perceptual sensitivity (e.g., forced-identification), which may or may not index other aspects of awareness, may also differ in their relative sensitivity for indicating how much perceptual information is available from an incoming series of stimuli. This point is important since the degree of perceptual salience that a given stimulus possesses may partially determine the extent to which that stimulus has access to particular neural structures and different aspects of cognition. When determining whether two stimuli have been masked or perceptually degraded to an equal degree, therefore, one's conclusions may specifically depend on which test of perceptual sensitivity is employed.

This point is important since, either at a group level or within individual participants, the majority of published investigations using masked facial expressions have not explicitly sought to verify that different masked images were perceptually degraded to an equal degree. However, many have contrasted the effects exerted by, or compared the neural processing of, discrete masked facial expressions and drawn conclusions from these analyses. Additionally, due in large part to the conceptual complexity surrounding these issues and the varying methods used by different laboratories, investigations that have utilised masked facial expressions have employed a variety of target-mask stimulus onset asynchonies (SOAs), a diverse range of stimuli, and a heterogeneous sampling of methods for assessing the effectiveness of their chosen masking parameters.

To date, only three investigations have sought to directly quantify and compare the relative perceptual salience of multiple facial expressions when presented under identical masking conditions. One obscure and infrequently cited study has suggested that, when masked by a plaid pattern at SOAs of 20–32 ms, Japanese expressions of happiness are more reliably recognised than are several other facial expressions (Ogawa & Suzuki, 1999). However, the methods and stimuli used in this investigation are very different from those used in other investigations and the findings in this report have not been replicated with face images used as masks.

Results reported by Esteves and Öhman (1993), on the other hand, have been widely cited. In that set of investigations, neutral faces were used to mask neutral, angry, and happy expressions with target-mask SOAs ranging from 20 ms to 300 ms. They observed lower thresholds for identifying happy faces at SOAs of 50 ms and above, but concluded that SOAs of 20–30 ms were effective for blocking the conscious identification, and establishing the subliminality, of all facial expressions as evidenced by chance performance on a forced-choice test of discriminability. These results were subsequently replicated and it was concluded that ''with an SOA of 30 ms, subjects had recognition performance at the chance level'' (p. 377) and that the previously used masking parameters were indeed effective and did ''not allow subjects to consciously perceive the masked stimuli'' (Esteves, Parra, Dimberg, & Öhman, 1994b).

However, the small numbers of trials used in these investigations were acknowledged as being insufficient for conducting conventional signal detection analyses and potentially limited any ability to observe effects at the shorter, more difficult, SOAs. Since neither investigation reported power analyses or formally tested the null hypothesis, the question is still open as to whether the identification of masked facial expressions is possible at shorter SOAs; and if so, whether distinct facial expressions are masked to an equal degree.

We suggest that, in the absence of data to suggest otherwise, there is no a priori reason to expect that identical backward masking parameters will be equally effective at perceptually degrading different facial expressions. In fact, based on a number of published observations, we expect that the perception of different facial expressions will not be blocked to an equal degree when identical masking stimuli and target-mask temporal relations are employed. For example, a considerable amount of evidence has now accumulated to suggest that the processing and recognition of different facial expressions involves a distributed and partially dissociable set of neural circuits in the brain (Adolphs, Damasio, Tranel & Damasio, 1996; Blair, Morris, Frith, Perrett, & Dolan, 1999; Harmer, Thilo, Rothwell, & Goodwin, 2001; Kesler-West et al., 2001). The specific brain regions that underlie the perception of a given facial expression, in turn, may be more or less sensitive to changes in an incoming stimulus' sensory information; and as such, play a role in determining how difficult it is to effectively mask, or perceptually degrade, that particular facial expression.

Specifically, we propose that happy faces are perceptually more salient and should therefore be more difficult to backwardly mask than angry faces even at extremely short SOAs. This hypothesis derives from several reported findings in the literature. First, and most transparently, this effect has previously been observed at moderate SOAs (greater than or equal to 50 ms; Esteves & Öhman,

1993) and not convincingly refuted at shorter SOAs. Second, as reported by Whalen (1998), adults subjectively rank having had more experience with happy or smiling facial displays than with angry, sad, fearful, disgusted, or surprised facial expressions. This repeated exposure to smiling faces over the course of one's lifetime may serve to selectively increase the recognition familiarity associated with happy facial expressions. Third, Pollak and colleagues (2000) have reported that even as early as 3–6 years of age nonmaltreated children are better at discriminating happy faces than facial displays of anger, fear, disgust, or sadness.

If the relative frequency of exposure to different facial displays affects each expression's corresponding degree of discriminability and perceptual salience, these effects should manifest at very brief exposure durations in studies of backward visual masking. We designed the current pair of experiments to investigate the degree to which distinct facial expressions, when masked at brief SOAs (17 ms) by neutral faces, are blocked from different levels of perceptual access as indexed by explicit (self-reported) awareness, forced-identification performance, and stimulus sensitivity in a signal detection paradigm. These data were subsequently analysed to determine whether discrete facial expressions presented outside of explicit awareness possess a comparable degree of perceptual salience when measured with two arguably more sensitive scales.

EXPERIMENT 1

Method

Participants and materials. Students and employees of the University of Wisconsin–Madison (N = 46, 16 male) viewed multiple pairs of target-mask images on a ViewSonic 17GS monitor (operated at 60 Hz; RGB medium-short persistence phosphors) at a distance of 50 cm, head stabilised with chin rest. Mask images consisted of 60 unique Caucasian male faces displaying neutral expressions. Target pictures were five male models selected from the *Pictures of Facial Affect* (Ekman & Friesen, 1976) displaying each of three facial expressions (neutral, angry, and happy). Greyscale images were matched on luminance to within 0.05 standard deviations of the group mean, subtended a visual angle of 9.95° horizontally and 11.29° vertically, and were surrounded by an isoluminant border (violet or grey) that accounted for an additional .91° and 1.03°.

Procedure. Trials contained a target expression (neutral, angry, or happy) that was presented for a single frame and immediately masked for an additional six, corresponding roughly to a target duration of 4 ms, mask duration of 87 ms, and interstimulus interval (ISI) of 13 ms (total stimulus duration = ~104 ms, target-mask SOA = 16.67 ms).[1] Intertrial intervals (ITIs) varied between 1750 ms and 2000 ms. Stimulus and trial timing, single frame presentation of targets, and target-mask SOA were verified using Scopemeter FLUKE 105B Series II Oscilloscope in addition to examining the computer's internally generated timestamps. Blocks of trials consisted of four presentations of each of the 15 targets in random order. The 60 masks were randomly assigned to trials and each was used once per block. Trial order and target-mask pairings were newly generated for each participant.

Participants initially viewed 15–20 target-mask pairings without any instructions regarding where to focus attention, and were then walked through the task instructions while the target-mask images continued to be presented on the screen. This was done for two reasons. First, to familiarise participants with the stream of images that they would be watching during the actual task and to facilitate their understanding of the instructions. Second, to allow participants to directly examine the images (target-mask pairings) in a natural way, while moving their eyes and switching their attention between different aspects of the stimuli according to their own volition.

Participants next viewed 420 target-mask pairings (140 of each target expression) in a border colour discrimination cover task. Subsequently, an unexpected questionnaire was administered asking them to comment on everything that they could about the faces, and to specifically elaborate upon the *nature* of the image presentation. Participants were then asked if at any point they had seen multiple faces presented during the course of a single trial or noticed a perceptual change (e.g., flickering) within any given image. Finally, participants were debriefed that on every single trial, including on the trials that they viewed during the naturalistic viewing and instruction portions of the experiment, exactly *two* faces had been presented and that they had never viewed an unpaired image. They were again asked if they had noticed a double image or change within a given image.

Subsequently participants were given three buttons corresponding to the facial expressions and instructed to identify the target expression on each trial. To prevent extreme response bias, we informed participants that in every block each target expression would be presented exactly 20 times and that they should equally distribute their identification responses across the three buttons. Three blocks totalling 180 target-mask pairs (60 per expression) were presented. Participants were required to press a button on every trial, even if they felt as if they could not reliably discriminate the images.

[1] As cautioned by Bridgeman (1998), most researchers assume that the duration of a single frame is equal to the reciprocal of the video display terminal's raster rate in hertz. This is incorrect. The duration of a single frame is actually equal to the phosphor persistence (~4 ms for standard computer monitors). Reported calculations, therefore, typically overestimate the total length of any given image by approximately 12–13 ms on a 60 Hz monitor.

Perceptual sensitivity indexes

A participant's sensitivity for detecting or identifying a signal in a distribution of noise can be determined from a pair of hit and false alarm rates. Hit rate (H) is the probability of correctly making a response (R_j) given the corresponding stimulus (S_j). False alarm rate (F) is the probability of making a particular response (R_j) when the corresponding stimulus is absent ($S_{\neq j}$). To index participant sensitivity we calculated A′, the nonparametric analogue of d' from signal detection theory (Equation 1). For a comprehensive discussion of signal detection theory's assumptions and the A′ sensitivity index, we refer the reader to Snodgrass and Corwin (1988); Haase, Theios, and Jenison (1999); and Macmillan and Creelman (1990, 1991).

$$A' = .5 + (H - F)(1 + H - F)/[4H(1 - F)] \quad (1)$$

Results

No participants reported seeing double faces during the naturalistic viewing of the target-mask pairs. We defined unaware participants as those individuals who, even after being informed that they had viewed multiple pairs of faces during the stimulus familiarisation, task instruction, and border colour discrimination portions of the experiment, adamantly maintained that they had not seen any multiple images at any point during the entire experiment. Criteria for inclusion in this group were extremely stringent as we excluded even those participants who, when pressed, reported that they might have seen a double image at some point during the session. Four subjects fell on threshold between explicit awareness and complete unawareness (i.e., reported an image flickering but not a masked face) and were excluded from these analyses.

Confidence intervals (CIs; 95%) were separately constructed for those individuals who exhibited an explicit awareness of the masked target images (n = 20) and those participants who did not exhibit any awareness whatsoever (n = 22; Table 1) using target emotional expression as a within-subjects factor. As indicated by the A′ sensitivity index, participants were able to reliably identify all of the masked target expressions above chance and to identify each of the distinct target expressions to a different degree $F(2, 80) = 28.66$, $p < .01$ (Table 1). Specifically, planned contrasts revealed that masked happy facial expressions were better identified than angry or neutral expressions by both the explicitly aware $F(1, 19) = 51.75$; 23.42 and explicitly unaware $F(1, 21) = 5.59$; 17.04 participants, all $p < .05$, respectively.

The explicitly aware and unaware subgroups did not reliably differ in their overall ability to identify masked facial expressions $F(1, 40) < 1.00$. However, a significant interaction between masked target expression and participant awareness revealed that explicitly aware participants were relatively better at

TABLE 1
Experiment 1: Forced-identification

	A′	
Target	*Mean*	*CI*
Chance performance	0.500	
Explicitly aware (*n* = 20)		
Neutral	0.601	0.555–0.647
Angry	0.519	0.478–0.561
Happy	0.702	0.643–0.761
Unaware (*n* = 22)		
Neutral	0.552	0.508–0.596
Angry	0.566	0.532–0.601
Happy	0.642	0.588–0.696

Note: Chance performance denotes null sensitivity. Mean = group average awareness. CI = minimum and maximum values defining the 95% confidence interval for the group mean.

identifying happy and neutral faces than their explicitly unaware counterparts, but relatively worse at identifying masked angry expressions $F(1,40) = 6.71(A')$, $p < .05$. Offering insight into the source of this interaction, an examination of the specific types of false alarms that were committed indicates that masked neutral faces were more often classified as being angry (31%) than happy (19%) by the explicitly aware participants $t(19) = 3.92$, $p < .01$. This observation compliments and expands on results previously reported by Esteves and Öhman (1993) in which participants more often interpreted unmasked neutral faces as being angry than happy. However, in contrast to the explicitly aware participants and diverging from the findings of Esteves and Öhman (1993), explicitly unaware individuals in our investigation did not differentially interpret masked neutral expressions as being angry (25%) versus happy (26%) $t(21) < 1.00$.

It should be noted that while there were marked individual differences in ability to identify the masked images, a majority of participants were able to reliably identify all of the masked facial expressions above chance (Figure 2; 81% were able to identify happy; 81%, neutral; and 64%, angry). Furthermore, the majority of individuals exhibited a differential degree of sensitivity for identifying each of these discrete facial expressions in a manner consistent with the overall group (Figure 2; 74% of participants identified happy expressions better than angry and 74% identified happy better than neutral).

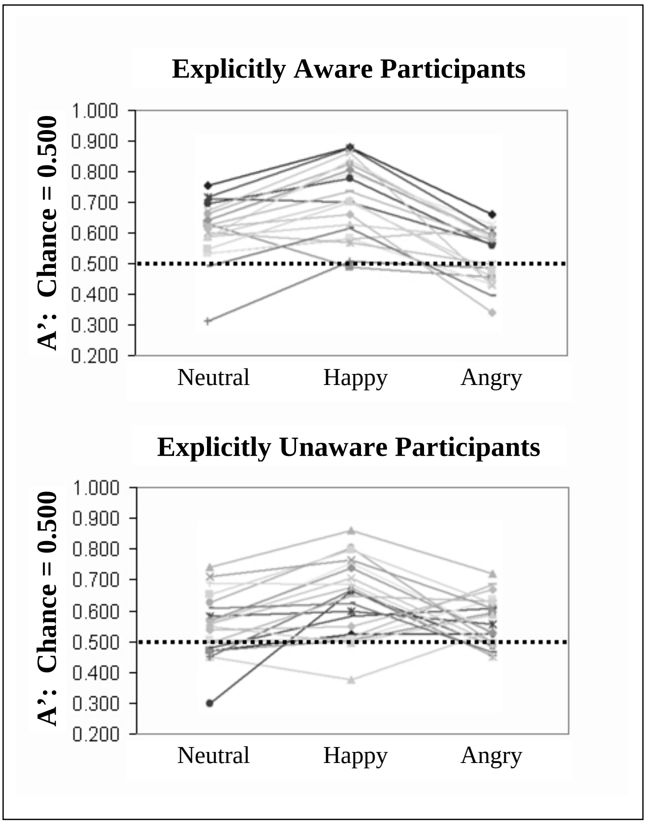

Figure 2. A′ (Chance performance = 0.500; dotted line). Individual subject sensitivity plots for identifying each of the masked facial expressions employed in Experiment 1. The plots are split into two groups which show those individuals who exhibited an explicit awareness of the masked images (top) and those participants who did not exhibit any explicit awareness whatsoever (bottom).

EXPERIMENT 2

Method

A second experiment was run in which the primary task was to detect, as opposed to identify, the presence of backwardly masked faces. We employed this alternative test of perceptual sensitivity to ensure that the results obtained in the first experiment were not artifactual. For example, forced-identification performance in Experiment 1 may not have been driven solely by perceptual factors, but also by the relative confusability of the masked facial expressions.

Additionally, participants may have been basing their forced-identification judgements on gut feeling, stemming from the registration of automatically evoked autonomic reactions (Katkin, Wiens, & Öhman, 2001), rather than on differences in the perceptual salience of the masked images. Participants in the second experiment (*N* = 19; 10 males), therefore, were not required to make any emotional evaluations of facial expression. In fact, emotion and facial expression were never mentioned during this experiment. In Experiment 1, the overall ability of explicitly aware and unaware individuals to identify backwardly masked facial expressions did not reliably differ. Furthermore, participants in both groups were significantly better at identifying masked happy than masked neutral or angry expressions. Therefore, it was not an aim of the second experiment to directly compare individuals based on self-reported explicit awareness.

In addition to the three facial expressions used in the first experiment—neutral, anger, and happy—the facial expression of fear was selected and matched with the other target stimuli to yield four unique affective expressions for Experiment 2. Target-mask pairings (signal + noise) were randomly intermixed with solitary mask (noise) trials in a 1:2 ratio (20 target-mask, 40 noise trials per block). A single frame was added to the solitary mask trials in order to match stimulus duration across trial types, meaning that all stimuli consisted of exactly seven frames and were perceptually identical except for the perceptual content of the first frame. Failure to match solitary and compound stimuli in this manner would leave open the possibility of detecting masked images based on temporal instead of strictly perceptual cues.

Within any single block, only one emotional expression was presented on the target-mask trials. However, in order to maximise participants' chances of noticing multiple faces during naturalistic viewing, only double images (target-mask pairings) were presented during the familiarisation and instruction portions of this experiment. Block presentation order was counterbalanced across participants, and each individual's entire block sequence was repeated three times yielding 12 blocks (3 per expression) per participant.

Instructions stated that within every block, exactly one third of the trials would contain the sequential presentation of two images, and two thirds of the trials would be single image presentations. Facial expression was not mentioned. Participants were simply instructed to press a button each time they noticed two images, and to refrain from responding otherwise. The initial cover task for assessing explicit awareness of the target expressions was the same as that used in the first experiment, but with two modifications. It was run here with fear incorporated into the randomised trial order (4 target expressions displayed 3 times by each of 5 models), and as a single block (60 trials; 15 of each expression) in order to minimise any stimulus repetition effects that might have been present in the first experiment.

Results

No participants reported seeing double faces during the naturalistic viewing of the target-mask pairs. When asked, none of the participants reported having any explicit awareness of the backward masking; and when pressed, nine of 19 met the full criterion for complete explicit unawareness. However, as defined by the stimulus detection threshold, all participants exhibited a marked degree of perceptual sensitivity for each of the masked pictures with a main effect of target emotional expression $F(3, 54) = 4.33$, $p < .01$. Without exception, every single participant exhibited at least some degree of perceptual sensitivity for detecting every masked facial expression. Descriptive statistics are provided in Table 2.

Planned contrasts revealed that masked happy expressions were better detected than angry $t(18) = 2.70$ and neutral expressions $t(18) = 3.01$, $p < .01$ (one-tailed). Two-tailed tests revealed no reliable difference between happy and fearful $t(18) = 1.01$ or neutral and angry $t(18) < 1.00$ expressions; but did suggest that fearful expressions were better detected than angry $t(18) = 2.10$, $p = .05$ and neutral expressions $t(18) = 2.02$, $p < .06$. Furthermore, the majority of individuals exhibited a differential degree of sensitivity for detecting each of these discrete facial expressions in a manner consistent with the overall group: 68% detected happy better than angry, 79% detected happy better than neutral, 63% detected fear better than angry, and 79% detected fear better than neutral.

For all expressions, signal detection performance was near ceiling. To illustrate, note that combining a 90% hit rate with a 10% false alarm rate yields an A′ of 0.944, a value lower than that obtained for any masked expression in this study (Table 2). Therefore, detection of facial expressions with these masking parameters (17 ms target-mask SOA) using a standard video display terminal may be a data-limited approach for indexing the perceptual salience of backwardly masked facial expressions (Holender, 1986) and constrain the possibility of teasing apart other potentially meaningful differences in perceptual sensitivity (e.g., fear vs. happy).

TABLE 2
Experiment 2: Signal detection (N = 19)

	A′	
Target	*Mean*	*CI*
Chance performance	0.500	
Neutral	0.951	0.923–0.980
Angry	0.950	0.915–0.985
Happy	0.962	0.931–0.994
Fear	0.959	0.927–0.991

DISCUSSION

In the light of the fact that distinct facial expressions are processed by a partially dissociable underlying functional neuroanatomy, we suggested that there was no a priori reason to expect that identical backward masking parameters would be equally effective at perceptually degrading different facial expressions. Furthermore, based on results previously reported by Esteves and Öhman (1993), the finding that nonmaltreated children are better at discriminating happy expressions than other facial displays (Pollak et al., 2000), and the observation that normal adults subjectively rank having had more experience with smiling faces than other emotional expressions (Whalen, 1998), we proposed that happy faces are perceptually more salient and familiar than are other affective facial expressions; and should therefore be more difficult to backwardly mask at extremely short SOAs.

We designed two experiments to investigate the degree to which different facial expressions, when masked by neutral faces, are blocked from different levels of perceptual access as indexed by explicit (self-reported) awareness, forced-identification performance, and stimulus sensitivity in a signal detection paradigm. As defined by two different measures of perceptual sensitivity, these data suggest that an SOA of 16.67 ms between an emotional target expression and a neutral mask is not sufficient for reliably blocking either the simple detection or the forced-identification of masked facial expressions.

Furthermore, not all target-mask pairs are created equal, an observation that is not readily apparent when only explicit awareness is used to verify masking effectiveness. In line with our predictions, these results demonstrate that happy faces are less effectively masked by neutral expressions than are angry or neutral. These effects were replicated with two different measures of perceptual sensitivity and observed to occur even in the complete absence of explicit awareness. The possibility was also raised that such effects may extend to other affective expressions (e.g., fear).

However, other factors may also have contributed to this pattern of results. For example, all of the happy expressions used in this investigation contain big smiles and white teeth, whereas all of the neutral images have closed mouths. Therefore, it may be that happy faces are particularly difficult to backwardly mask with neutral expressions simply because of the stark changes in brightness and contrast that occur in the mouth region. However, we believe that this is unlikely because participants were never shown templates of the emotional expressions that could potentially cue them to look for particular features within the face images, such as teeth, to represent particular expressions. Furthermore, three of the five angry and all of the fearful expressions used in this investigation contain teeth. However, follow-up comparisons revealed identical hit rates for the angry faces with teeth (0.301) when compared to the angry faces without teeth (0.302) in Experiment 1; and perhaps more convincingly, that masked angry faces were not detected any more efficiently than masked neutral faces in Experiment 2 (Table 2). These results are inconsistent with the idea that changes in contrast and brightness in the mouth region can be solely responsible for determining the perceptual salience of masked facial expressions.

An important limitation of this report concerns how generalisable these findings are to other investigations that utilise different stimuli and equipment. Furthermore, this criticism can be justifiably levelled at any investigation that does not use a standardised set of equipment and methods; and in backward visual masking research, no standardised set-up currently exists. This point directly underscores the key importance of using within-experiment controls to assess the effectiveness of backward visual masking. However, despite these methodological limitations, these results have not been demonstrably refuted and should therefore concern any investigator who uses these techniques or refers to this literature.

While these results replicate previously reported findings (Esteves & Öhman, 1993), they also fundamentally extend the work and provide new information about the multitude of factors that may be operating in research that involves backwardly masked facial expressions. In particular, we highlight three issues that should be taken into consideration by investigations that involve backwardly masked facial expressions. First, caution should be exercised when using the terms ''unconscious'', ''nonconscious'', or ''subliminal'' for addressing data acquired in the absence of explicit awareness. Lack of explicit awareness does not mean that incoming information is unavailable through other perceptual measures, which may or may not index other aspects of awareness; it speaks only to whether a given individual is able to report having a phenomenal experience of that information. Many participants here were able to reliably identify those same facial expressions which, when backwardly masked, did not even register within their subjective experience.

Second, substantial between-subject variability in perceptual thresholds may exist and go unnoticed unless a sufficiently sensitive index is used to assess the perceptibility of backwardly masked visual stimuli. As is illustrated in the bottom half of Figure 2, we observed substantial individual differences in perceptual sensitivity even among those participants who did not report seeing any of the backwardly masked faces. Although between-subject differences in the ability to detect or identify masked facial expressions will not necessarily influence within-subject contrasts in any systematic manner, large intersubject variability is likely to introduce considerable noise into the data and may unintentionally skew the overall group mean. Therefore, matching the perceptual degradation of masked facial expressions across individuals may be a desirable approach, though in other contexts such individual variability may be informative.

Third, these data indicate that striking dissociations may exist in the relative perceptibility of different backwardly masked facial expressions. These perceptual differences, in turn, may influence the degree to which each expression has access to particular neural structures and different aspects of cognition. It is possible, therefore, that systematic differences in the perceptual degradation of backwardly masked facial expressions may inadvertently influence experimental results by adventitiously contributing to positive experimental results, attenuating fundamental effects of interest, or contributing to compound observations that derive from a variety of factors.

Alternatively, observed differences in the relative perceptibility of masked emotional expressions may indicate, rather than cause, a dissociation in the underlying processes that are responsible for perceiving and responding to particular affective expressions. Relatedly, differences in perceptual salience may be integral to defining the very essence of particular expressions, both signifying and producing a dissociation in the underlying neural and cognitive processing. For example, if a particular affective process (e.g., amygdalar activation) were to selectively intensify the perceptual salience of a given masked stimulus (e.g., fear face), that stimulus might consequently engage a variety of neuroanatomical operations to which it would not otherwise have been granted access.

These possibilities are consistent with our observation that large-scale individual differences exist in both explicitly aware and unaware participants' perceptual sensitivities for different backwardly masked facial expressions. If the perceptual threshold for detecting or identifying a given masked facial expression is constitutionally bound to the underlying neuroanatomical processing of that particular expression, directly measuring participants' perceptual sensitivities for ''unseen'' facial expressions may prove to be a valuable tool for assessing overall face processing ability and indexing individual differences in affective perception and emotional reactivity. The summary point is that investigators should either make a concerted effort to verify that backwardly masked visual images have been perceptually degraded to a comparable degree, or quantify and give thoughtful consideration to the ways in which any observed perceptual dissociations might have arisen from or impacted the differential processing of these stimuli.

There currently exists no undisputed standard for assessing how effectively a stimulus has been occluded from awareness. This is due in part to the conceptual complexity that surrounds these issues as well as to the varying methods and equipment that are used in different laboratories. However, we have demonstrated that tractable solutions exist for indexing perceptual differences between backwardly masked expressions and therefore suggest that this issue can and should be appropriately addressed in future investigations of this nature. The chosen test of perceptual sensitivity should be based upon the particular research questions being asked and, at least until norms and a common set of procedures have been established, these tests should be conducted on a within experiment (or within laboratory) basis.

Manuscript received 18 April 2002
Revised manuscript received 20 July 2003

REFERENCES

Adolphs, R., Damasio, H., Tranel, D. & Damasio, A. (1996). Cortical systems for the recognition of emotion in facial expressions. *The Journal of Neuroscience*, *16*, 7678–7687.

Bechara, A., Damasio, H., & Damasio, A. R. (2000). Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex*, *10*, 295–307.

Bernat, E., Shevrin, H., & Snodgrass, M. (2001). Subliminal visual oddball stimuli evoke a P300 component. *Clinical Neurophysiology*, *112*, 159–171.

Berns, G. S., Cohen, J. D., & Mintun, M. A. (1997). Brain regions responsive to novelty in the absence of awareness. *Science*, *276*, 1272–1275.

Blair, R. J. R., Morris, J. S., Frith, C. D., Perrett, D. I., & Dolan, R. J. (1999). Dissociable neural responses to facial expressions of sadness and anger. *Brain*, *122*, 883–893.

Bridgeman, B. (1998). Durations of stimuli displayed on video display terminals: (n -1)/f + persistence. *Psychological Science*, *9*, 232–233.

Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, *77*, 305–327.

Bunce, S. C., Bernat, E., Wong, P. S., & Shevrin, H. (1999). Further evidence for unconscious learning: Preliminary support for the conditioning of facial EMG to subliminal stimuli. *Journal of Psychiatric Research*, *33*, 341–347.

Dehaene, S., Naccache, L., Le Clec'h, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., Van De Moortele, P. F., & Le Bihan, D. (1998). Imaging unconscious semantic priming. *Nature*, *395*, 597–600.

De Houwer, J., Hendrickx, H., & Baeyens, F. (1997). Evaluative learning with ''subliminally'' presented stimuli. *Consciousness and Cognition*, *6*, 87–107.

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

Ekman, P., & Friesen, W. V. (1976). *Pictures of facial affect.* Palo Alto, CA: Consulting Psychologists Press.

Ekman, P., Sorenson, E. R., & Friesen, W. V. (1969). Pan-cultural elements in facial displays of emotion. *Science*, *164*, 86–88.

Esteves, F., Dimberg, U., & Öhman, A. (1994a). Automatically elicited fear: Conditioned skin-conductance responses to masked facial expressions. *Cognition and Emotion*, *8*, 393–413.

Esteves, F., & Öhman, 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.

Esteves, F., Parra, C., Dimberg, U., & Öhman, A. (1994b). Nonconscious associative learning: Pavlovian conditioning of skin-conductance responses to masked fear-relevant facial stimuli. *Psychophysiology*, *31*, 375–385.

Haase, S. J., Theios, J., & Jenison, R. (1999). A signal detection theory analysis of an unconscious perception effect. *Perception & Psychophysics*, *61*, 986–992.

Harmer, C. J., Thilo, K. V., Rothwell, J. C. & Goodwin, G. M. (2001). Transcranial magnetic stimulation of medial-frontal cortex impairs the processing of angry facial expressions. *Nature Neuroscience*, *4*, 17–18.

Holender, D. (1986). Semantic activation without conscious identification in dichotic-listening, parafoveal vision, and visual masking: A survey and appraisal. *The Behavioral and Brain Sciences*, *9*, 1–23.

Katkin, E. S., Wiens, S., & Öhman, A. (2001). Nonconscious fear conditioning, visceral perception, and the development of gut feelings. *Psychological Science*, *12*, 366–370.

Kesler/West, M. L., Andersen, A. H., Smith, C. D., Avison, M. J., Davis, C. E., Kryscio, R. J., & Blonder, L. X. (2001). Neural substrates of facial emotion processing using fMRI. *Cognitive Brain Research*, *11*, 213–226.

LeDoux, J. E. (1996). *The emotional brain: The mysterious underpinnings of emotional life.* New York: Simon & Schuster.

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*, 3–26.

Mathews, A., & Macleod, C. (1986). Discrimination of threat cues without awareness in anxiety-states. *Journal of Abnormal Psychology*, *95*, 131–138.

Merikle, P. M., & Reingold, E. M. (1990). Recognition and lexical decision without detection: Unconscious perception. *Journal of Experimental Psychology. Human Perception and Performance*, *16*, 574–583.

Meyer, C., & Waller, G. (1999). The impact of emotion upon eating behavior: The role of subliminal visual processing of threat cues. *International Journal of Eating Disorders*, *25*, 319–326.

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.

Morris, J. S., Öhman, A., & Dolan, R. J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature*, *393*, 467–470.

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.

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

Ogawa, T., & Suzuki, N. (1999). Response differentiation to facial expression of emotion as increasing exposure duration. *Perceptual and Motor Skills*, *89*, 557–563.

Öhman, A. (1986). Face the beast and fear the face: Animal and social fears as prototypes for evolutionary analyses of emotion. *Psychophysiology*, *23*, 123–145.

Öhman, A., & Soares, J. J. F. (1998). Emotional conditioning to masked stimuli: Expectancies for aversive outcomes following nonrecognized fear-relevant stimuli. *Journal of Experimental Psychology. General*, *127*, 69–82.

Pashler, H. E. (1998). *The psychology of attention.* Cambridge, MA: MIT Press.

Pollak, S. D., Cicchetti, D., Hornung, K., & Reed, A. (2000). Recognizing emotion in faces: Developmental effects of child abuse and neglect. *Developmental Psychology*, *36*, 679–688.

Reingold, E. M., & Merikle, P. M. (1988). Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics*, *44*, 563–575.

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.

Rotteveel, M., de Groot, P., Geutskens, A. & Phaf, R. H. (2001). Stronger suboptimal than optimal affective priming? *Emotion*, *1*, 348–364.

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

Seitz, A. R., & Watanabe, T. (2003). Is subliminal learning really passive? *Nature*, *422*, 36.

Shanks, D. R., & Lovibond, P. F. (2002). Autonomic and eyeblink conditioning are closely related to contingency awareness: Reply to Wiens and Öhman (2002) and Manns et al. (2002). *Journal of Experimental Psychology. Animal Behavior Processes*, *28*, 38–42.

Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology. General*, *117*, 34–50.

Whalen, P. J. (1998). Fear, vigilance, and ambiguity: Initial neuroimaging studies of the human amygdala. *Current Directions in Psychological Science*, *7*, 177–188.

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.

Winkielman, P., Zajonc, R. B., & Schwarz, N. (1997). Subliminal affective priming resists attributional interventions. *Cognition & Emotion*, *11*, 433–465.

Wiens, S. & Öhman, A. (2002). Unawareness is more than a chance event: Comment on Lovibond and Shanks (2002). *Journal of Experimental Psychology: Animal Behavior Processes*, *28*, 27–31.

Wong, P. S., Bernat, E., Bunce, S., & Shevrin, H. (1997). Brain indices of nonconscious associative learning. *Consciousness and Cognition*, *6*, 519–544.

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.