A Tiger and a President: Imperceptible Celebrity Facial Cues Influence Trust and Preference

ROBIN J. TANNER
AHREUM MAENG

Neuroscientific research suggests that the brain has evolved specific capabilities enabling automatic social judgments of others to be made based on facial properties alone. However, little research in marketing has considered the consequences of how facial imagery is automatically processed. We explore automatic perceptions of familiarity by using morphing software to digitally combine unfamiliar faces with those of Tiger Woods and George Bush. Despite a complete lack of conscious recognition, trustworthiness ratings of the composite faces are clearly influenced by the celebrities in question. This appears to be due to implicit recognition being sufficient for individuals to automatically access their own summary valence judgments of either Woods or Bush. Alternative explanations based on a perceptual-fluency account, or implicit recognition sufficient to perceive specific trait ratings, are ruled out. These findings suggest that the marketing practice of digitally manipulating the attractiveness of facial imagery risks overlooking the important influence of familiarity.

Human faces are arguably the most important visual stimuli we perceive. They are quite staggeringly complex and represent an exceptionally rich source of social information, which we have evolved to be extremely adept at decoding (Engell, Haxby, and Todorov 2007). For example, faces not only reveal a tremendous amount of information about our underlying emotions, but the mere adoption of certain facial expressions is sufficient to moderate the underlying physiology and subjective experience of emotion (Ekman 1992). Furthermore, we begin to mimic facial expressions soon after birth (Hatfield, Cacioppo, and Rapson 1993), infer movement tra-

jectories and intentions from peoples' gazes (Nummenmaa, Hyona, and Hietanen 2009), and are astonishingly adept at rapidly identifying angry faces in crowds (Hansen and Hansen 1988). In fact, evidence from cognitive, developmental, and evolutionary psychology, as well as from cognitive neuroscience, supports the face-specificity hypothesis, which suggests that the face represents a "special" (and possibly unique) class of stimuli whose processing is handled by discrete and specialized neural circuits (Kanwisher and Yovel 2006).

Indeed, converging evidence suggests that the human brain has evolved specific capabilities that enable automatic social judgments of others (such as how trustworthy they are) to be made based on facial properties alone (Critchley et al. 2000; Engell et al. 2007; Winston et al. 2002). That humans have evolved to be highly sophisticated (and automatic) interpreters of facial information has potentially profound implications for marketers. After all, from decorative models in cosmetics ads to celebrity endorsers of sodas, facial imagery is nearly ever-present across all forms of marketing communication. It is thus rather surprising how little research in marketing has directly explored the implications of how people automatically code (and respond to) facial stimuli.

While there has certainly been a growing awareness among marketing researchers that environmental cues can automatically influence consumer behavior, this work has

to date largely focused on examining the effects of lexical and brand primes on specific goals and traits (Chartrand et al. 2008; Fitzsimons, Chartrand, and Fitzsimons 2008). Although automatic responses to some more overtly social cues (such as the presence of interpersonal mimicry; Tanner et al. 2008) have been studied, relatively little research has specifically examined automatic evaluation of facial cues. The research that does exist has tended to focus on certain effects of specific expressions, such as in research on emotional contagion (Howard and Gengler 2001; Small and Verochi 2009) rather than chronic (emotionally neutral) features of the face. However, given the emerging neuropsychological evidence that our brains are essentially hardwired to automatically process facial information, we believe that research on automatic processing of facial imagery merits considerably broader attention from marketing scholars.

The current research was conceived to explore one aspect of automatic facial processing that has the potential to be highly relevant to marketers. Specifically, we are interested in whether perceptions of the trustworthiness of unfamiliar faces (i.e., of apparently unknown individuals) can be influenced by incorporating subtle facial cues of familiar people into the faces in question. This question was motivated, in part, by the finding from the facial morphing literature that subtly (and imperceptibly) incorporating elements of an individual's own face into an otherwise unfamiliar face increases reported trust in, and/or preference for, the unfamiliar face (Bailenson et al. 2009; DeBruine 2002). Importantly, since the self is simultaneously both similar and familiar, it is unclear from this stream of research whether a similarity- or familiarity-based process underlies the observed positive effects. A primary goal of our research was to better understand these competing explanations. We do this by relying on faces that are familiar but not similar to those of our participants. Specifically, we investigate the effect of subtly incorporating elements of celebrity images into otherwise unfamiliar faces.

The remainder of this article is structured as follows. First, we provide a brief overview of research on automatic facial processing, which suggests that perceptions of trust can be automatically determined based on facial cues alone. Second, we review the existing finding from the digital morphing literature that subtle incorporation of self-images into unfamiliar faces increases perceived trust in the individuals pictured, and we discuss both similarity- and familiarity-based explanations that could underlie this effect. Finally, we present four studies that explore the effects of morphing unknown faces with highly familiar celebrity images. These studies enable us to make the following claims. First, morphing unknown faces to contain imperceptible cues of familiar individuals can lead to similar positive effects on perceived trustworthiness as does self-morphing. Second, this effect appears to be driven by implicit recognition of the specific familiar individual concerned rather than by either a more general fluency effect or by some aspect of their facial structure that they may share with others. Third, this implicit recognition appears sufficient for a broad approach/avoid

valence pertaining to the familiar individual to be perceived but is insufficient for specific trait assessments of them to be accessed from memory.

AUTOMATIC FACIAL PROCESSING

Accurate social judgments have in all likelihood been a significant determinant of human survival and indeed of the survival of our primate ancestors (Winston et al. 2002). For example, from an evolutionary biology perspective, the ability to quickly discern whether a stranger might be hostile and/or untrustworthy could significantly enhance an individual's ultimate chance of reproductive success (Cacioppo 2004). In fact, research suggests that almost all the stimuli we encounter are evaluated on a positive/negative valence dimension (Bargh 1997; Osgood, Suci, and Tannenbaum 1957), with these evaluations evoking approach and avoid behavioral tendencies (Corwin 1921; Lewin 1935). Moreover, in social contexts, specifically, these valence evaluations serve as a central input into person-impression formation (Fiske, Cuddy, and Glick 2007; Kim and Rosenberg 1980).

Considerable evidence now suggests that many of these social judgments are actually made entirely automatically, that is, without any conscious deliberation or intent (Bargh and Chartrand 1999; Devine 1989; Hassin, Uleman, and Bargh 2005). This automatic "social intelligence" appears to be supported by specific cognitive and neural mechanisms (Winston et al. 2002), especially in the area of facial perception (Kanwisher and Yovel 2006). The amygdala region, in particular, appears to play a central role in making automatic socially relevant judgments on the basis of facial appearances (Adolphs, Tranel, and Damasio 1998; Winston et al. 2002). Indeed, Todorov and Engel (2008) argue that a primary function of the amygdala may be to provide continuous vigilance by evaluating objects and agents prior to interacting with them and that one way this is achieved is by automatically evaluating the valence of faces (see also Amaral 2002; Whalen 1998). Once rendered, these valence evaluations are then used to guide approach/avoidance behaviors (Chen and Bargh 1999).

Of particular relevance to the continuing debate as to the extent to which judgments can truly be determined entirely automatically, a notable feature of these apparently automatic face-based judgments is that they originate in a part of the brain that is neuroanatomically dissociable from the location where explicit judgments based on facial information are made (Critchley et al. 2000). Put differently, in layman's terms, automatic (and unconscious) processing of facial information is potentially unique in that it can actually be attributed to a specific location in the brain. That evolution has chosen to endow us with such specific capabilities for automatic facial processing only serves to reinforce our contention that marketers would be wise to pay greater attention to how the implicit processing of facial information might influence consumer behavior.

From a marketing perspective, a particularly important insight from this body of research is that individuals appear

to automatically categorize faces in terms of their perceived trustworthiness (Engell et al. 2007). Indeed, when assessing unfamiliar faces, trust ratings have been demonstrated to provide the best approximation of the core underlying valence evaluation (Oosterhof and Todorov 2008). This capability enables individuals to form judgments about a novel individual's trustworthiness after as little as a 50-millisecond exposure to their face (Todorov, Pakrashi, and Oosterhof 2009; Willis and Todorov 2006). Moreover, the social consequences of these rapid automatic judgments are not restricted to trust ratings in the laboratory. For example, ratings of competence based solely on a 1-second exposure to faces of unfamiliar political candidates were shown to predict 69% of the 2004 US Senate races and to be linearly related to the margins of victory (Todorov et al. 2005).

It is important to note, however, that while at this point a fairly significant body of research has demonstrated that automatic valenced judgments can be made on the basis of facial features alone, much of the research to date has (understandably) been phenomenological in nature, with less emphasis having been placed on understanding precisely how these judgments are arrived at. Put simply, while we know that automatic valence judgments are made, we know a lot less about the specific inferential processes that underlie them. While not the central focus of this article, we note that there are a number of different possible facial cues that automatic judgments might rely on. Such cues include facial expressions (Critchley et al. 2000), the actual physical structure of the face (Carre, McCormick, and Mondloch 2009), and resemblance to a specific known individual (Verosky and Todorov 2010).

DIGITAL MORPHING RESEARCH

Digital morphing technology, which enables photographs of two individuals to be blended together into a highly realistic-looking virtual face, has provided researchers with a new technique with which to explore the implicit aspects of facial processing. For example, Bailenson et al. (2009) examined participant responses to composite images created by morphing the faces of politicians with the faces of participants themselves (henceforth "self-morphing") in a ratio of 60% politician, 40% participant. This ratio is significant, as the authors demonstrate that it produces composite faces within which individuals could not perceive their own image. Put differently, participants universally believed they were viewing unaltered pictures of the politicians in question. However, attitudes toward the composites reflected an entirely different story. When the candidate was relatively unfamiliar to them (e.g., the 2006 Florida gubernatorial candidates for non-Florida residents), participants consistently demonstrated stronger relative preference for the morphed images containing elements of their own faces.

Moving beyond unfamiliar politicians, Bailenson et al. demonstrated that this self-morphing effect similarly influenced assessments of the 2004 presidential candidates made by weak partisans and independents (though it had no effect on strong partisans). These positive assessments once again

accrued to the self-morphed images without any participants reporting recognition of their own face, suggesting that they occurred automatically, and that any recognition that did occur was entirely implicit. Nor is this positive effect of self-morphing limited to simple liking alone. In research that directly aligns with the neuroscientific work on automatic trust assessment, DeBruine (2002) used a very similar morphing manipulation to demonstrate that individuals playing a computer-based sequential trust game were more likely to trust a player represented by a self-morphed avatar.

While these findings demonstrate that individuals have a clear preference for self-morphed images, major questions remain about the specific process underlying this effect, which makes it hard to fully assess the extent of the potential marketing relevance. In particular, a broad theoretical question is whether the observed preference for self-morphed images was driven by narrow perceptions of similarity (i.e., by resemblance to the self specifically) or by a broader familiarity-based process (i.e., by resemblance to a familiar individual, of which the self is but one example), since the two are essentially confounded when morphing an individual's own face into a composite image. As Bailenson et al. (2009, 938) state, "facial similarity and familiarity are difficult, if not impossible, to disentangle." This theoretical distinction has genuine practical significance, as while technology might enable limited use of individuals' own faces (or elements of their own faces) in personally-targeted messages, such self-endorsing strategies are highly problematic to scale to mass media advertising and may in any case cross the privacy line for many consumers. However, subtle manipulation of facial familiarity, should it prove possible, is innately more scalable, assuming that the familiarity manipulation relies on a cue that is likely to be broadly efficacious.

From a purely applied perspective, the self-morphing findings raise the possibility that automatic facial processing may considerably influence consumer behaviors in ways that are currently entirely overlooked by marketers. For example, while digital manipulation of images is certainly nothing new to advertisers, the default focus is almost exclusively on attractiveness, with models being "airbrushed and digitally manipulated beyond the human plane" (Ives 2008, 17). However, both the Bailenson at al. (2009) and the DeBruine (2002) findings suggest that attractiveness (for which both sets of authors controlled) may be too narrow a lens through which to consider the full implications of how consumers automatically process faces. As such, we believe that there are both theoretical and practical reasons for attempting to deepen our understanding of the roles played by familiarity and similarity in the automatic evaluation of faces. In the next section, we use the existing self-morphing findings to frame a broader discussion outlining different theoretical accounts that could explain how consciously imperceptible facial cues might automatically influence how faces are perceived.

A Similarity-Based Account

Both Bailenson et al. (2009) and DeBruine (2002) suggest that automatically perceived similarity to the self underlies the

positive preference for self-morphed images they observed. Certainly social psychologists have long observed that individuals have a preference for people similar to themselves, being more helpful toward them (Park and Schaller 2005), judging them as more attractive and persuasive (Shanteau and Nagy 1979), and even favoring them in adoption decisions (DeBruine 2004a). Moreover, evolutionary psychology suggests a specific reason to explain this preference for similarity: since visual similarity is a cue of potential genetic kinship, then treating similar others preferentially is one strategy to improve the likelihood of survival of the broader family group. Consistent with this logic, humans and other primates do appear capable of recognizing unfamiliar kin (Parr and de Waal 1999; Wells 1987). This skill likely relies on a visual form of phenotype matching, which is the ability to detect the outward manifestations of a specific genotype (Hauber and Sherman 2001). This ability to automatically recognize kin certainly suggests that (automatically) perceived facial similarity could underlie the positive effects observed to accrue to self-morphed images.

Familiarity-Based Accounts

While the aforementioned similarity-based account has the potential to explain the self-morphed findings, in this article we explore the possibility of a broader familiarity-based account. Such an account would suggest that morphing an unknown face with the face of a familiar (but genetically unrelated) individual would have the potential to similarly influence how individuals perceive the resulting composite face. To be absolutely clear, we use the term "familiarity-based account" broadly to encompass any process that involves an individual being influenced by virtue of their prior exposure to the specific face in question. In our view, the influence of such familiarity could potentially occur with at least three distinct levels of specificity and/or recognition.

Perceptual Fluency-Based Account. One of the most robust findings in social psychology is that repeated exposure to almost any stimulus increases preference for it. This mere-exposure effect (Zajonc 1968) has proven to be highly robust and persists even when the repeated exposures occur outside of conscious awareness (Kunst-Wilson and Zajonc 1980). Most accounts of the mere-exposure effect suggest that perceptual fluency (the subjective feeling of ease when processing information) increases with repeated exposure to a stimulus and enhances preference either by being misattributed to liking (Bornstein and D'Agostino 1994; Schwarz 2004) or by generating positive affect that directly enhances stimuli evaluations (Winkielman and Cacioppo 2001). As far as potentially explaining preference for self-morphed images is concerned, the reality is that since the invention of the mirror most humans are repeatedly exposed to their own faces, and as such the positive perceptions observed to accrue to self-morphed images could potentially be explained by the increased perceptual fluency that comes from familiarity with one's own face.

Valence Accessibility Account. A pure perceptual fluency-based account for automatic facial processing suggests that an unknown face morphed with a familiar face would be beneficially evaluated due solely to its increased familiarity, not because any particular individual-specific meaning or valence associated with the familiar face is perceived (i.e., a fluency process does not rely on specific recognition). From an evolutionary perspective, however, one might imagine that such a process could be maladaptive. For example, if an unknown individual resembled a known threat, then despite their familiarity, the appropriate automatic reaction would presumably be to not trust and/or avoid this individual. More adaptive, then, would be for implicit recognition to enable the individual to automatically perceive some level of approach-avoid motivation. Indeed, this possibility is consistent with fundamental research on primitive approach-avoid motivations, which, depending on the exigencies of the current situation, are thought to govern the deployment of predispositions to approach or withdraw (Bargh 1997; Neumann and Strack 2000). Moreover, visual cues are one likely trigger of these motivational systems (Neumann and Strack 2000), which accords with the research in neuroscience suggesting that the amygdala preconsciously evaluates the valence of a face (Amaral 2002; Whalen 1998). This suggests a second familiarity-orientated explanation for participants' observed positive preference toward self-morphs —namely, that some level of implicit recognition of their own faces may have led them to experience similar approach-avoid motivations toward the self-morphs as they would with full conscious exposure to their own image. Since the majority of (nondepressed) individuals are positively self-disposed, automatically perceived approach motivations could certainly explain the positive ratings assigned to self-morphed images. We refer to this possible process as a valence accessibility account—that implicit recognition of a familiar individual in a morphed face is sufficient to enable an underlying (and pre-existing) valence judgment of the familiar individual to be automatically perceived.

Trait Accessibility Account. A third possibility is that implicit recognition of a familiar individual in a morphed image is sufficient for an individual not only to perceive a broad approach/avoid motivation specific to the known individual but also to automatically access more specific trait-based assessments about that person from memory. If implicit recognition does enable trait assessments to be automatically accessed from memory, then since most individuals view themselves as being above average on many dimensions (Alicke et al. 1995), full implicit recognition could also explain the positive trait ratings assigned to self-morphed images. We refer to this final possibility as a trait accessibility account—the possibility that implicit recognition of a familiar face actually enables specific trait beliefs of the specific individual to be automatically accessed from memory.

In summary, at a high level, two broad accounts exist for the positive preference observed to accrue to self-morphed images (Bailenson et al. 2009; DeBruine 2002). First, consistent with an innate ability to recognize our own kin, it is possible that an automatically perceived perception of sim-

ilarity (to the self) was the primary mediating process. Second, implicit perceptions of familiarity (as opposed to similarity to the self specifically) may have driven the observed effects. Such perceptions of familiarity could operate with varying levels of actual implicit recognition. First, a fluency-based account suggests that generally perceived familiarity (i.e., with no specific recognition of the familiar individual) could influence preference for morphed images via increased perceptual fluency. Second, a valence accessibility account suggests that implicit individual-specific recognition may occur (i.e., of the self in the self-morph studies reported) at a level sufficient to enable individuals to access a basic approach/avoid judgment (henceforth "valence judgment") of the individual concerned. Finally, a trait accessibility account would suggest that this implicit recognition might actually be sufficient to access specific trait information from memory.

Of note, a similarity-based account would suggest that the positive preference observed to accrue to self-morphed images (Bailenson et al. 2009; DeBruine 2002) would not extend to morphed images containing familiar faces. However, all of the familiarity-based processes discussed above would predict that similar preferences might also be observed for faces morphed to contain familiar faces. As such, while in this article we do not attempt to specifically rule out a unique similarity-based process, a primary goal was to establish whether a broader familiarity-based process might operate. Second, while it is not clear that the three familiarity-based accounts discussed above should necessarily be entirely mutually exclusive, a secondary goal was to begin to understand the relative merits of each, especially in the context of the facial morphing paradigm.

Research Overview

Four studies explore the effects of morphing unfamiliar faces with the faces of highly familiar celebrities (Tiger Woods and George W. Bush). Consistent with a familiarity-driven process (and inconsistent with a pure similarity-based explanation), these studies demonstrate that adding imperceptible celebrity facial cues to pictures of unfamiliar people increases participants' perceived trustworthiness in (studies 1 and 2A) and reported likelihood to buy from (study 3) the individuals pictured. Furthermore, this effect appears consistent with a valence accessibility account; that is, it involves a level of implicit recognition that appears to be sufficient to enable a broad approach/avoid evaluation of the underlying celebrity to be perceived (study 2B). However, it is inconsistent with accounts based on trait accessibility (study 2A) or perceptual fluency (studies 2B and 3).

STUDY 1: TRUSTING A CELEBRITY-MORPHED IMAGE

Study 1 was designed to provide an initial test of whether subtly incorporating elements of a familiar celebrity image into an unfamiliar face can lead the resulting composite image to be perceived as more trustworthy than the original face. To do this we compare the trustworthiness of an un-

familiar face to the same unfamiliar face digitally blended with the face of Tiger Woods. Tiger was chosen because we expected him to be highly familiar to our participants. Of note, all data was collected before the Tiger Woods scandal materialized. We chose to focus on trust for several reasons. First, much of the neuroscientific work on automatic trait assessment has identified perceptions of trust as being significantly determined by facial properties (e.g., Engell et al. 2007) and to proxy an underlying valence judgment (Oosterhof and Todorov 2008). Second, trust has long been identified as a critical marketing variable, influencing firm metrics as diverse as salesman performance (Doney and Cannon 1997), the efficacy of celebrity endorsers (Ohanian 1990), and ultimately perceptions of the brand itself (Ambler 1997; Chaudhuri and Holbrook 2001).

Design and Method

Study 1 randomly assigned 109 undergraduate participants to either a Tiger-morph condition or a control-morph condition in a single-factor between-subjects design. Participants were instructed to spend a few moments looking at a picture of a face and then to record their impressions of the trustworthiness of the individual pictured on an 11-point scale anchored from -5 (well below average) to $+5$ (well above average). Participants were asked to not overthink the exercise and to rely on their gut instincts and immediate reactions. In the Tiger-morph condition, a composite face was created by digitally blending a stock model face with the face of Tiger Woods. In the control-morph condition, the same stock model face was digitally blended with the face of an individual chosen to be of approximately similar age and attractiveness to Tiger Woods. Using a control image that has itself been created using a digital blending process is required to ensure that any observed effects could not be attributable to the effect of morphing alone, since composite faces (which are more symmetrical) have been shown to be perceived as being more attractive than the source faces from which they are constructed (Langlois and Roggman 1990). Furthermore, "the what is beautiful is good" stereotype (Dion, Berscheid, and Walster 1972) would suggest a possible link between such morph-induced attractiveness and subsequent trait ratings. As such, to ensure both facial morphs were indeed perceived as being similar in attractiveness, we conducted a pretest ($N = 75$) utilizing an 11-point scale anchored from -5 (very unattractive) to $+5$ (very attractive) that revealed no differences in the perceived attractiveness of the two faces ($M_{\text{Tiger\_morph}} = 3.64$, $M_{\text{control\_morph}} = 3.68$; $t(73) = .17$, $p = .86$). The original and composite faces are shown in figure 1.

Both composite images (henceforth the "Tiger-morph" and "control-morph" images) were created using Abrosoft Fanta Morph Pro 4 software and were composed of 65% of the unfamiliar stock model face and 35% of either Tiger Woods's face or the control face. The morphing process involves marking corresponding points at standard landmarks on both faces (such as the corners of the eyes and mouth), which allows the software to generate a composite face from a weighted

FIGURE 1

CONSTRUCTION OF FACES USED IN STUDY 1
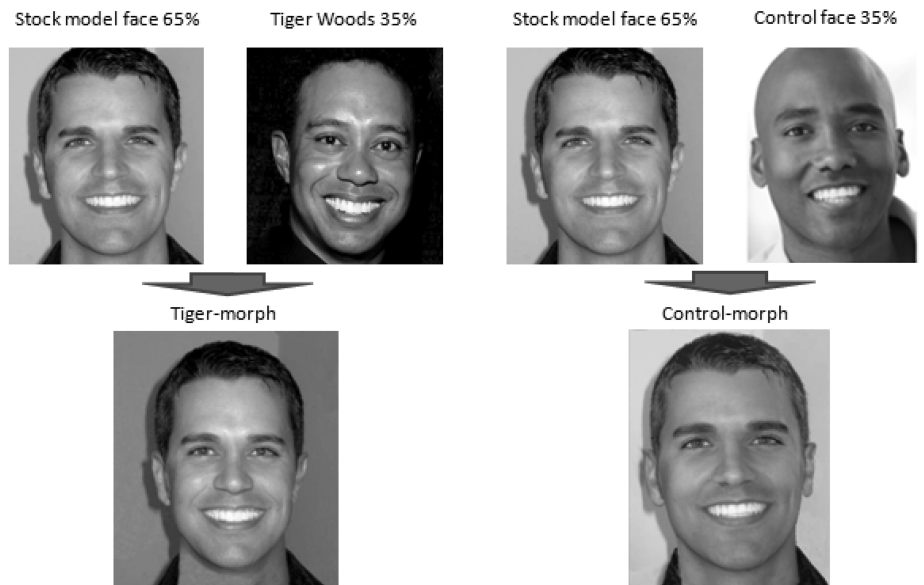
average of landmark-specific parameters of the source faces. While prior morphing research has demonstrated that individuals are unable to consciously recognize themselves in composite images containing 40% of their own face (Bailenson et al. 2009; DeBruine 2004b), we conservatively chose to use the 65 : 35 ratio in order to minimize the chance of conscious recognition.

Once participants had provided their trustworthiness rating, they completed a funnel debriefing (Bargh and Chartrand 2000) that first asked if there was anything odd or unusual about the face pictured, then asked if the face reminded them of anyone they knew, before finally providing an open-ended opportunity to report if the face looked like anyone famous.

Results and Discussion

First, we considered the funnel debriefing data to assess possible differences in how the faces were perceived between conditions. While 16 participants reported that the face they viewed was odd, there was no difference in the percentage reporting this between the control-morph condition (13%) and the Tiger-morph condition (17%; $\chi^2(1) = .34$, $p = .56$). Of greater importance, none of the explanations for the faces looking odd betrayed any underlying suspicion of the morphed nature of the faces pictured. In fact, the explanations offered were universally entirely unrelated to the experimental manipulation and included such reasons as "smiling too much" and "eyes a little closed." As far as general (conscious) perceptions of familiarity were concerned, we observed no differences across the two con-

ditions, with 22% of participants in both conditions indicating that the face they viewed reminded them of someone they knew ($\chi^2(1) = .003$, $p = .96$). However, none of the individuals (in either condition) who suggested the face was familiar indicated that it resembled Tiger Woods; instead, they tended to suggest that the face resembled someone they had a personal relationship with. As such it appears that, at least as far as conscious recognition is concerned, there was no difference in either general perceptions of familiarity or specific recognition of Tiger Woods between the two composite faces.

However, despite participants in the two conditions reporting the two images as being equally familiar, their reported trust ratings revealed a different pattern. In fact participants in the Tiger-morph condition rated the face as being more trustworthy ($M = 1.75$) than did participants in the control-morph condition ($M = 0.95$; $t(107) = -2.2$, $p < .05$). Thus it appeared that subtly incorporating elements of a highly familiar face into an otherwise unfamiliar image increased perceived trust in the pictured individual. These initial data thus conceptually mirror the prior research, which has demonstrated positive changes in perceptions of the trustworthiness of self-morphed images (DeBruine 2002, 2004b). However, unlike the self, which is both similar and familiar, Tiger Woods is highly unlikely to be on average similar in appearance to our participant population. As a result, these data are more consistent with a broader familiarity-based process as underlying the positive perceptions of the trustworthiness of the Tiger-morphed image.

While our chosen focus was trust, a natural question to

FIGURE 2

FACES USED IN STUDY 2

Stock model face   Bush-morph   Control-morph
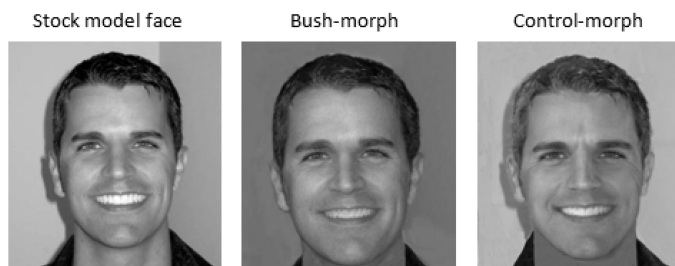
ask is whether the observed effect is unique to trust or might generalize to other trait assessments. To explore this question we ran a replicate experiment with 35 MBA students who rated trustworthiness, liking, and athletic ability of the same two faces used in the main study. The only difference in procedure was that 1–10 scales were used (as opposed to −5 to +5). Not only did these data reveal a (marginally) significant replication of the prior trust result ($M_{\text{Tiger}} = 5.3$, $M_{\text{control}} = 4.0$; $t(31) = -1.9$, $p < .07$) but a similar result was also observed for reported liking ($M_{\text{Tiger}} = 5.3$, $M_{\text{control}} = 3.8$; $t(31) = -2.1$, $p < .05$). Furthermore, the trust and liking ratings were themselves very highly correlated ($r = .96$, $p < .01$), with a significant number of participants actually reporting the exact same score for both trust and liking. However, we observed no difference in reported athletic ability between the two face conditions ($M_{\text{Tiger}} = 7.1$, $M_{\text{control}} = 7.4$; $t(31) = 0.51$, $p > .6$), and we also observed no correlation between the athletic ability ratings and the ratings of either trust or liking (both $p > .4$).

These follow-up data enable us to draw a number of additional conclusions. First, the fact that participants did not report the Tiger-morph to be more athletic provides initial evidence against a full trait accessibility account, as presumably if implicit recognition was sufficient to enable full trait accessibility from memory, then this would have led to higher athletic ratings of the Tiger-morph. This conclusion is admittedly a little speculative, however, as we did not measure our participant's beliefs about the athletic ability of Tiger Woods, and it is certainly possible (perhaps due to golf not being perceived as being a physically intense sport) that our participants simply did not perceive Woods to be especially athletic. Second, the reported liking data clearly suggest that the observed positive effect is not specific to trust and likely generalizes to a variety of general personality assessments. We believe that this result is consistent with a valence accessibility account, as it appears likely that automatically perceived positive valence would manifest via individuals making positive assessment of broad personality-type variables rather than simply trustworthiness alone (e.g., liking or honesty).

STUDY 2A: EXPLORING A TRAIT ACCESSIBILITY EXPLANATION

Study 2A was designed to build on study 1 in a number of ways. First, in order to generalize the results of study 1, we switch to a different highly familiar figure: George W. Bush. Second, in order to address possible concerns that the results in study 1 might have been driven primarily by the morph-control image (i.e., by some idiosyncratic negative feature of the control face), we include two control conditions in study 2: the original unaltered model face as well as a control-morph using an individual of similar age and attractiveness to George Bush. Finally, while no participants explicitly recognized Tiger Woods in study 1, implicit recognition may well have taken place. As previously discussed, such implicit recognition might have enabled either a broad valence judgment of Woods to be accessed or even for specific trait beliefs held about him to be perceived in the morphed image (though the athletic rating data provided preliminary data suggesting that this is unlikely). To better rule out the latter possibility, in study 2A we also collected participants' actual ratings of the trustworthiness of George W. Bush.

Design and Method

Study 2 randomly assigned 179 undergraduate participants to either a Bush-morph condition, a control-morph condition, or an unaltered-face condition in a single-factor between-subjects design. In the two morph conditions, composite faces were created by blending the stock model face (same face as study 1) with either 35% of the face of George W. Bush or 35% of a control face chosen to be of similar age and attractiveness to Bush. Participants in the unaltered face condition viewed the original unaltered stock model face. As in study 1, a pretest ($N = 66$) suggested that the two morphed faces were equally attractive ($M_{\text{Bush\_morph}} = 3.62$, $M_{\text{control\_morph}} = 3.73$; $t(64) = .53$, $p = .60$). The unaltered face and both composite faces used are shown in figure 2.

All other aspects of the experiment were identical to study 1 except that in a supposedly unrelated study participants

subsequently rated the trustworthiness of six celebrities (including George W. Bush) on the same 11-point scale.

Results and Discussion

First, we considered the funnel debriefing data and once again observed no reliable differences in conscious perceptions of familiarity across the three conditions (percentage reporting face resembled someone they knew: Bush-morph 28%; control-morph 19%; neutral image 15%; $\chi^2(2) = 2.50$, $p = .29$). As far as specific recognition of George Bush was concerned, none of the participants in any condition suggested that the face they viewed reminded them of George Bush.

Turning to our primary dependent measure of interest, there was once again an overall effect of face condition on reported trust ($F(2, 176) = 4.56$, $p < .02$). A planned comparison revealed that this effect was primarily driven by participants in the Bush-morph condition rating the face as being significantly more trustworthy ($M = 2.16$) than either participants in the unaltered face condition ($M = 1.27$) or the control-morph condition ($M = 1.30$; $t(176) = -3.02$, $p < .01$). There was no difference in perceived trust between the control-morph and unaltered-face conditions ($t(176) = .08$, $p > .90$). Next we considered whether participants' trust ratings in the Bush-morph condition were related to their actual beliefs about the trustworthiness of George Bush. This did not appear to be the case; in fact in the Bush-morph condition, ratings of trustworthiness of the composite face were essentially uncorrelated with actual ratings of the trustworthiness of George Bush ($r = .07$, $p = .62$). Thus, in addition to not being consciously recognized, any implicit recognition that may have been taking place was apparently not sufficient for specific beliefs about Bush's trustworthiness to be perceived in the morphed image.

While these data rule out a trait accessibility account, that is, the possibility that implicit recognition enables specific trait assessments to be accessed from memory, they do not speak to the alternative possibility that an intermediary level of implicit recognition might exist, one sufficient to perceive an overall valence judgment of the individual in question (i.e., whether to approach or avoid them) but insufficient to access preexisting beliefs about specific traits. Recall that, when assessing faces believed to be unknown, trust ratings have been demonstrated to provide a good approximation of automatically made valence evaluations of faces (Oosterhof and Todorov 2008). By this account, the positive trust ratings assigned to the Bush-morph might have reflected automatically perceived perceptions of positive valence (i.e., that an approach motivation was induced). However, it was impossible for us to assess an overall valence explanation for the study 2A data, as we had no independent measure of whether Bush might trigger an approach or avoid motivation in our participants. Certainly, despite his rather mixed reputation (both across traits and between individuals), we would also expect Bush to be inextricably associated with the primary role for which he is known, that of president of the United States (McCracken 1989). As such, his close association with the

presidency suggests that implicit recognition of him in a composite image might be expected to automatically cue approach-consistent concepts such as freedom and safety, thus leading to a positive valence appraisal of the composite image. However, rather than simply speculate about the possibility of a positive approach motivation, we designed study 2B to explore if we could measure this directly.

STUDY 2B: EXPLORING A VALENCE ACCESSIBILITY EXPLANATION

Study 2B had two main goals. First, we wanted to investigate whether the results of study 2A could be attributed to the Bush-morph image inducing an approach motivation in our participants. Second, we wanted to rule out a pure fluency-based explanation for the results of studies 1 and 2A (i.e., that the presence of the Woods or Bush facial cues led to increased fluency in processing the morphed images, with this fluency then being misattributed to the trust and liking ratings). The study utilized a modified version of a methodology used by Mehta and Zhu (2009) to evidence that the color red induces an avoid motivation while the color blue invokes an approach motivation. Specifically, these authors demonstrated that a blue/(red) background facilitated participants solving anagrams based on approach-/(avoid-) related words. In this study, instead of manipulating the background color, we located the appropriate anagrams directly beneath either the Bush-morph or the control-morph image from study 2A.

Design and Method

Four hundred and forty-five participants from a general online panel participated in this study for payment. The study was designed as a 2 (face presented: Bush-morph vs. control-morph) × 2 (solve time: control vs. approach related words) mixed design, with the second factor being the within-subject factor. Participants were randomly assigned to one of the two face-presented conditions, and the study was administered via computer. Participants were told that the researchers were interested in how people solve anagrams in the presence of another person and that in this first study they would simply be asked to solve a series of anagrams presented underneath pictures of faces. It was also explained that the first one or two letters of the target word would be underlined in the letter string and an example (ATSKABREF for which the solution is BREAKFAST) was presented. Finally, participants were instructed to write "don't know" if they could not solve a given anagram within 2 minutes.

Next, participants proceeded to attempt to solve nine anagrams, each of which were presented on a separate screen. Six of the anagrams represented control words (i.e., count, ranch), while three were associated with an approach motivation (i.e., advance, adventure). The presentation order of the anagrams was randomized. Depending on the condition, each anagram appeared on the screen directly below either the Bush-morph or the control-morph image from study 2A.

Results and Discussion

First, for each participant, we calculated their average solve time for control and approach words. Only correctly solved anagrams were included in the average calculations, and data from participants who got one or fewer words correct (in either category) were discarded. As a result, data from 45 participants were excluded from further analyses, leaving useable data for 397 participants.

These computed average solve times were then used in a 2 (face presented: Bush-morph vs. control-morph) × 2 (solve time: control words vs. approach words) repeated measures ANOVA, which revealed a significant interaction ($F(1, 395) = 4.50, p < .04$). A planned comparison revealed that participants who viewed the Bush-morph solved the approach-related anagrams significantly faster (M = 14.2 seconds) than did those who viewed the control-morph (M = 16.9 seconds, $p < .04$) but that there was no difference in the speed with which the control anagrams were solved ($M_{\text{Bush-morph}}$ = 15.2 seconds, $M_{\text{control-morph}}$ = 15.1 seconds; p = .95). These data thus build on the findings from study 2A by revealing that exposure to the Bush-morph image appeared to induce an approach motivation in participants. Put differently, participants appeared to implicitly recognize Bush in the morphed image to an extent sufficient to access a positive valence judgment. As such, these data support a valence accessibility account, that is, that the positive trust ratings observed to accrue to the Bush-morph in study 2A were attributable to an automatically induced approach motivation. Moreover, these data are very hard to reconcile with a pure fluency-based account, as not only is it rather problematic to explain why increased fluency in processing an image would influence the solve times of anagrams positioned adjacent to that image, but we can think of no reason why such an effect would only pertain to approach-related words.

STUDY 3: LIKELIHOOD TO BUY FROM A TIGER WOODS–MORPHED IMAGE

The first three studies demonstrated that blending a familiar (but not similar) face into an otherwise unfamiliar face can produce a composite face that is rated as being more trustworthy and likeable than the original. This effect on perceived trustworthiness obtained despite a complete lack of apparent recognition of the presence of the familiar image in the composite image. The original goal of study 3 was to begin to explore whether this "hijacking" of familiarity has the potential to affect behaviors of specific interest to marketers. As such, in study 3 we again use a Tiger-morphed image to examine whether the morphing procedure could influence perceived likelihood to buy from a pictured salesperson. Rather serendipitously, data were collected in two time periods: both before and at the peak of the Tiger Woods scandal. This enabled us to further explore the valence accessibility account we have argued underlies the results to date and to rule out explanations based on facial structure or expression alone.

Design and Method

Study 3 randomly assigned 319 undergraduate participants to either a Tiger-morph condition or a control-morph condition. In the Tiger-morph condition, a composite face was created by digitally blending a stock model face with the face of Tiger Woods. In the control-morph condition, the same stock model face was digitally blended with the face of an individual chosen to be of approximately similar age and attractiveness to Tiger Woods. As before, Tiger and the similarly attractive/aged face represented 35% of the respective morphed images. The stock model face in this study was a neutral image of a salesperson pictured from the chest upward. The original salesman image and both composite faces are shown in figure 3. Participant instructions were identical to those of studies 1 and 2A except that instead of rating trustworthiness they were asked to rate how likely they would be to buy from the pictured salesman on an 11-point scale anchored from −5 (very unlikely) to +5 (very likely).

The data were collected in two tranches. The first tranche (167 participants) was collected in the last week of September 2009, before the Tiger Woods scandal emerged. The second tranche (152 participants) was collected in the second week of December 2009, when the Tiger Woods scandal was at its peak and was almost completely dominating the media. Thus study 3 was a 2 face presented (Tiger-morph vs. control-morph) × 2 timing (pre-scandal vs. post-scandal) between-subjects design.

Results

Results indicated no main effect of face presented or timing but a significant interaction between the two ($F(1, 318) = 12.01, p < .001$). As shown in figure 4, while pre-scandal participants indicated that they were more likely to buy from the Tiger-morph than from the control-morph ($M_{\text{Tiger-morph}}$ = 1.46, $M_{\text{control-morph}}$ = 1.04), this pattern reversed itself post-scandal ($M_{\text{Tiger-morph}}$ = −.11, $M_{\text{control-morph}}$ = 1.02). Of note, given the particular nature of the scandal, there were no effects of gender.

Discussion

Despite a complete lack of conscious recognition of Tiger Woods, assessments of the likelihood to buy from the Tiger-morphed salesman appeared to be significantly influenced by the precipitous decline in Woods's real world reputation. This result is consistent with an associative learning/evaluative conditioning view of the scandal (DeHouwer, Thomas, and Baeyens 2001). Specifically, the (highly negative) scandal appears to have acted as an affective stimulus, which, since it was paired with Tiger Woods, was sufficient to change valence judgments of him. Furthermore, implicit recognition of Tiger was apparently sufficient to access these valence judgments. Of note, the fact that we observed a valence judgment being influenced by behaviors undertaken by the individual in question strongly

FIGURE 3

FACES USED IN STUDY 3

Stock model face   Tiger-morph   Control-morph
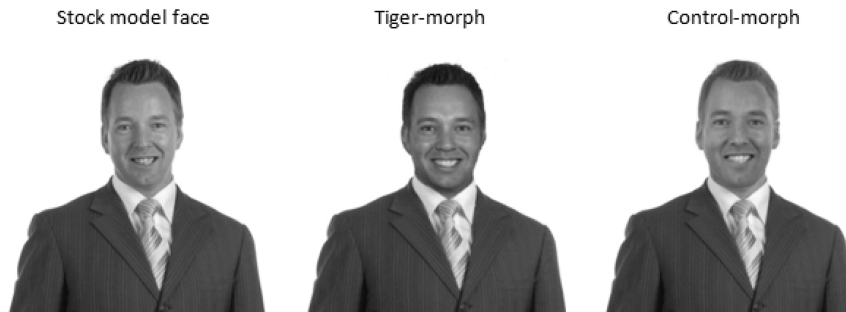
suggests that this judgment is related to the actual identity (i.e., Tiger Woods) of the known individual in the morphed image.
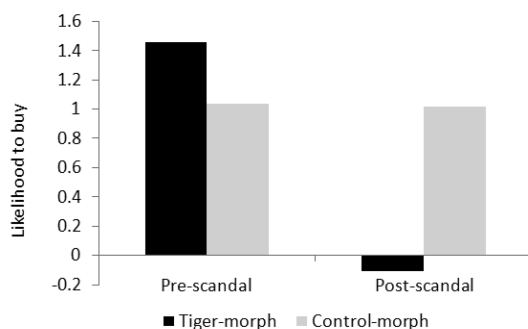
This distinction is important, as while study 2B revealed that the Bush-morph appeared to induce an approach motivation in participants, the data were entirely silent as far as the actual cause of that motivation is concerned. Put differently, when discussing implicit recognition of the individual in the morphed images (i.e., Woods or Bush), to this point we have essentially made an assumption that a valence accessibility account involves implicit recognition of the celebrities' facial properties sufficient for the perceiver to access a valence judgment of that specific individual in particular. However, an alternative possibility is that there is something about the facial structure of these two individuals (or even, possibly, of their specific expressions in the photos of them we used) that in and of itself drove our results to date (i.e., independent of their unique identities). For example, the possibility exists that both Bush's and Woods's facial features fit a genetic pattern generally perceived to appear more trustworthy or approachable. However, since in study 3 we utilized the exact same Tiger-morph image at two points in time, the

observed difference in participant ratings across these two points in time cannot be solely due to some idiosyncratic feature of Woods's facial structure itself or of his specific facial expression in the photo we used to create the morph.

Also of note, these results mirror study 2B in that they are very hard to reconcile with a pure fluency-based process. Indeed, we note that a fluency-based account might actually predict the opposite result as, despite Woods being arguably the world's most recognizable sports celebrity, the obsessive coverage of the scandal presumably could only have led participants in the post-scandal condition to be exposed to additional images of Tiger Woods. This repeated exposure should, all other things aside, have increased their fluency in processing the Tiger-morph, leading to enhanced preference. Certainly, even though one might argue that the relatively short duration of the scandal might be insufficient to make much positive difference to an individual's fluency in processing Woods's image, it is problematic for a fluency account to explain a decline in how Tiger was rated. Thus these data not only suggest that participants were implicitly recognizing Tiger Woods to a degree sufficient to access an overall valence judgment of him but that this implicitly perceived valence was sufficient to overcome the possible opposing effect of fluency.

GENERAL DISCUSSION

The current research provides the first demonstration of the important finding that our impressions of unfamiliar faces can be influenced by digitally morphing them with those of highly familiar figures. Study 1 showed that digitally blending an unfamiliar face with 35% of the face of Tiger Woods resulted in a composite face that was perceived as more trustworthy (and in the replicate more likeable) than the original face. This effect occurred despite 100% of our participants failing to identify any resemblance between the Tiger-morphed face and Tiger Woods himself. Study 2A both replicated this result with a different celebrity (George W. Bush) and additionally demonstrated that participants' ratings of the trustworthiness of the Bush-morphed image were uncorrelated with their consciously held beliefs about

FIGURE 4

WILLINGNESS TO BUY FROM SALESMAN AS A FUNCTION OF FACE PRESENTED AND SCANDAL TIMING

Bush's actual trustworthiness. These data were thus inconsistent with a trait accessibility account, as it did not appear that participants were implicitly recognizing Bush to an extent sufficient to automatically perceive specific trait assessments of him.

Study 2B built on this result by demonstrating that the Bush-morph image appeared to induce an approach motivation in participants, a result consistent with a valence accessibility account, that is, the idea that implicit recognition is sufficient only to access an overall valence judgment of the known individual concerned (Todorov and Engel 2008). Study 3 expanded on this finding by revealing that a marked decline in Tiger Woods's overall reputation appeared to negatively influence how a Tiger-morphed composite face was perceived, a result that suggests that overall valence judgments of morphed faces are (at least partially) determined by actual perceptions of the specific individual concerned rather than by some non-individual-specific aspect of their facial expression or facial structure. A second contribution of studies 2B and 3 was that they served to rule out a pure fluency-based explanation, as it is very hard to reconcile a fluency-based account with the pattern of data we observed in these studies.

It is important to note that all these results were obtained without a single participant in any condition (celebrity-morph or otherwise) ever consciously suggesting resemblance to the celebrity in the composite image. As such, it is worth highlighting that all these results were observed using a manipulation of facial imagery that is subtle enough to universally elude conscious awareness. Also worthy of note, with the exception of the post-scandal likelihood-to-buy rating for the Tiger-morph in study 3, all the ratings our participants provided (for all the faces presented) were positive in an absolute sense (i.e., greater than the midpoint of the scale). This result, which may seem a little probabilistically surprising, is actually broadly consistent with the work of Cacioppo and colleagues on the positivity offset (Cacioppo, Gardner, and Bernston 1999). Cacioppo argues that neutral stimuli actually tend to induce a weak approach motivation, which he suggests serves to motivate an organism to explore and learn about the environment (since novel stimuli may otherwise not be approached). As such, the finding that truly unknown faces (e.g., our control faces) are positively rated is consistent with the idea that exposure to them induced a weak approach motivation in our participants (and thus is consistent with a valence assessment process more generally).

Theoretical Contributions

The current results contribute to the nascent facial morphing literature and to the wider facial processing literature in several important ways. First, recall that both Bailenson et al. (2009) and DeBruine (2002, 2004b) demonstrated that, despite a complete lack of conscious recognition, self-morphed images benefited from enhanced preference compared to the original unaltered source face. Both sets of authors primarily attributed this result to automatic perceptions of similarity. In our studies, we conceptually replicate these findings when we incorporate celebrity faces rather than participants' own images into composite faces. Since the celebrities we chose (Tiger Woods and George W. Bush) are unlikely to be similar to our participants, our results are clearly inconsistent with a perceived similarity account (or more precisely with a similarity-to-the-self account). While our data do not explicitly rule out similarity as being a unique factor in these prior studies, they do suggest that familiarity-based accounts may have more potential to parsimoniously explain preference for both self-morphed and celebrity-morphed images.

Second, turning to familiarity-based accounts specifically, our work suggests that some measure of implicit recognition likely contributes to how morphed faces are perceived (be they self-morphed or celebrity-morphed). While study 2A demonstrated that participants rating a Bush-morph image were unable to access specific beliefs about George Bush's trustworthiness, study 2B revealed that a broader approach/avoid type evaluation was possible. Together, studies 2A and 2B combine to suggest that implicit recognition, at least in our facial morphing paradigm, occurs at an intermediary level of specificity. It appears sufficient to access an overall valence judgment of the individual in question (as in study 2B) but insufficient to access more specific trait-based ratings that may be held and that would be reportable with full conscious recognition (as in study 2A). Study 3 further supported this valence accessibility account by demonstrating that an actual behavior undertaken by a specific individual (i.e., Tiger Woods) led to a subsequent change in how an unknown face (morphed to include Woods's face) was perceived. As such, these findings resonate with the neuroscience literature that suggests that a primary function of the amygdala is to automatically provide a high-level valence-based evaluation of agents prior to interacting with them (Amaral 2002; Todorov and Engel 2008; Whalen 1998).

Finally, we believe that our data directly contribute to the rather thin discussion in the literature concerning precisely how these valence-based evaluations are constructed. As we emphasized in our introduction, there is surprisingly little research on the nature of the informational building blocks from which the implicit evaluations are presumed to be built. For example, does a specific face induce an approach motivation (and thus garner high trustworthy ratings) due to a subtle (and hard to control for) aspect of its particular expression, due to its chronic structural features (which may fit a genetic pattern indicative of trust), or because it resembles a specific known individual for whom an overall positive evaluation exists? While our research is clearly insufficient to fully characterize what we suspect is a complex and interdependent set of inferential processes, we do note that in study 3 the Tiger-morphed face was identical in both timing conditions. Thus facial expression and idiosyncratic structural features of the face were both perfectly controlled for. As such, study 3 suggests that the inferential process involved appeared to involve sufficient implicit recognition of a specific individual (Tiger Woods) for an individual-

specific valence judgment of him to be perceived. Thus, while in no way ruling out the potential for facial expressions and structural facial features to also contribute to automatic judgments, our research does identify that the specific role resemblance to a known individual can play in determining automatic valence judgments. Hence, while our experimental paradigm may have been the digital morphing of faces, a broader implication is that our automatic attentional resources appear to be extremely sensitive to cues of facial familiarity. We believe that future research should increasingly focus on better understanding the full range of inferential processes underlying the implicit processing of facial information. Indeed, this challenge appears to represent a particularly suitable arena in which neuroscientists and decision scientists could fruitfully collaborate.

Implications for Marketers

At a high level, we believe that the most important implication of our work is that it suggests a general need for marketers to pay more attention to the full potential consequences of how facial information is automatically processed. With facial imagery being such an ever-present cue in the world of marketing communications, the simple reality is that there is a tremendous variety of circumstances and contexts where automatic processing of facial stimuli may influence consumer behaviors of interest. In our view, while the influence of attractiveness and certain emotional expressions have received some specific attention, the potentially broader implications of facial familiarity have been largely unexplored by marketers.

More specifically, these data suggest a novel way for marketers to use celebrities and highly familiar figures. Instead of solely placing them front and center as endorsers, an alternative way to utilize them is to subtly morph their faces with those of otherwise unfamiliar stock model images. Indeed, mirroring the success of "forgotten spokespersons" (Weisbuch and Mackie 2009), such a strategy may even have the potential to outperform explicit use of celebrities in some circumstances. Prior research has demonstrated that the persuasive impact of spokespersons can be enhanced when prior exposure cannot be recalled, as this facilitates misattribution of the perceptual fluency experienced while processing the endorser (Weisbuch, Mackie, and Garcia-Marques 2003). Put differently, when the endorser is correctly identified, any effect of fluency is more likely to be correctly attributed to their fame and familiarity. Similarly, we believe it is possible that implicitly perceived valence judgments of individuals are likely candidates for misattribution given the absence of conscious awareness of their source.

As well as this possible valence misattribution, another potential advantage of using celebrity-morphed images is that the lack of conscious recognition of the celebrity, which we routinely observed, might also mean persuasion knowledge (Friestad and Wright 1994) is less likely to be activated, which, in and of itself, might lead to greater potential for successful persuasion. As such, we believe that there are multiple reasons to believe that digital morphing may offer marketers a way to capitalize on celebrities' familiarity without paying all the costs (be they to the persuasion process or perhaps financial) of the attendant fame. We note, however, that the legal ramifications of using celebrities in morphed images are likely to be rather complex. While a celebrity owns the rights to the use of his or her image, it is not entirely clear how these rights might extend to composite faces containing a minority proportion of their face, especially when they are not consciously recognizable in the composite (which would also make policing problematic).

Perhaps, however, some of the benefits of using celebrities in morphed images might be obtainable without actually using the actual celebrity in question. Indeed, an intriguing implication of our work involves the potential for using celebrity look-alikes in morphed images. If look-alikes are used directly in an advertising context, then unless their resemblance to the celebrity of interest is unusually close, it is likely that consumers will be consciously aware that the look-alike is not in fact the real deal. This conscious perception of illegitimacy may interfere with the possible benefits that might otherwise be expected to (automatically) accrue from the resemblance. However, if the look-alike is used in a morphed image, then an interesting possibility is that this may offer a way to capitalize on the benefits of the resemblance without the attendant risks of conscious correction.

Limitations and Directions for Future Research

As we noted earlier, when evaluating unknown faces, ratings of trust have previously been demonstrated to be a good approximation for an overall valence evaluation of a face (Oosterhof and Todorov 2008). Put differently, when asking individuals to rate the trustworthiness of an unfamiliar face, the answer they provide appears to be significantly determined by an automatically perceived approach/avoid evaluation. In essence, then, these trust ratings might be innately different in nature (they certainly are in terms of their source) from perceptions of trustworthiness of known people, and they really are better conceptualized as a proxy for an overall evaluation of valence. As such, it is important not to overstate the extent to which such evaluations might influence consumer behavior in the way that trust (as traditionally conceptualized and measured) has been shown to do. However, we believe that study 3 clearly demonstrates the potential for these overall evaluations to affect more than simple trust ratings and to directly influence consumer behavior. Certainly, however, one fruitful avenue for future research is to begin to better delineate the circumstances in which these automatic evaluations of valence can influence downstream consumption behaviors. To give one example, imagine a typical print advertisement that makes some form of product claim and that includes one or more stock model images. The question is, would the positive valence that could be obtained by morphing the model face(s) with a known figure influence the extent to which the product claim is attended to and/or believed?

A second (and related) limitation is that we are some distance from being able to provide a complete theoretical account of how morphed faces are perceived and processed. In particular, in the current research, our primary focus has been on exploring the extent to which individuals can implicitly recognize celebrity images in composite faces. We have placed less emphasis on fluency explanations, although we note that the results of studies 2B and 3 are entirely inconsistent with a purely fluency-based process. That said, since the manipulation in study 3 was naturally occurring, the resulting consequences for internal validity require us to be conservative in how we interpret the results. We note, however, that in very recent research, Verosky and Todorov (2010) demonstrated a conceptually similar result in a laboratory environment. In particular, these authors found that the positive or negative behaviors they previously associated with novel faces appeared to transfer (in broad valence terms) to composite faces including the aforementioned novel faces. As such, it seems unlikely that fluency is the primary mechanism influencing how composite faces are automatically perceived. Our view, however, is that it is unlikely to play no role whatsoever and that future research could usefully examine it in isolation. One idea such research might pursue is to examine the consequences of digitally combining multiple unfamiliar faces into one composite face. We believe that this might result in a composite face that is perceived to be more familiar, thus enabling familiarity to be studied while perfectly controlling for valence effects (since all the source faces would be unknown.)

A final limitation relates to the valence judgments we have argued underlie the effects we have observed. We have argued that our data support the view that when evaluating an apparently unknown face, if that face subtly resembles a known individual (which we achieve via morphing), then a valence judgment of that known individual automatically colors ratings of the apparently (and consciously perceived to be) unknown face. Indeed, one of the major contributions of this work is revealing that such valence accessibility is possible. What this article does not examine, however, is the process via which such summary valence judgments are formed. For example, while study 2B empirically demonstrated that exposure to the Bush-morph appeared to invoke an approach motivation in our participants, it had nothing to say about why this might be. Indeed, the reader might legitimately wonder why such an apparently polarizing figure would invoke an approach motivation. While we believe understanding the origins of summary valence evaluations is a (highly complex) issue for future research to explore, one thing we would stress is that the approach and avoid distinction is a fundamental driver of motivation across all life forms (Elliot and Covington 2001) and is perhaps best conceptualized as a primal big picture distinction between stimuli likely to be beneficial or harmful (Elliot 2006). As we speculated earlier, our belief is that Bush's automatic association with the presidency is a likely candidate for explaining the approach motivation we observed, but it remains exactly that, speculation. We suggest that future research on valence judgments should focus both on the primary roles people are associated with (particularly relevant for celebrities whose fame relates to acting) as well as the actual behaviors they carry out in their real lives (which were clearly relevant in study 3). Indeed, it is worth noting that one benefit of the digital morphing technique we used is that it actually provides a method via which valence assessments of known individuals can be implicitly measured.

Conclusion

From salesmen to print ads, the human face is almost ever-present across the broad spectrum of marketing communication efforts. Furthermore, considerable research now suggests that humans, as befits innately social animals, have evolved with significant mental capabilities designed solely to facilitate rapid implicit processing of facial information. However, despite this nontrivial confluence between marketing practice and neural circuitry, little attention has been paid to understanding the possible ramifications such automatic facial evaluations may have in marketing contexts. While marketers narrowly focus on digitally manipulating attractiveness, our results suggest that manipulations of familiarity may have similar potential to influence consumers. As such, we hope that these initial findings may help to precipitate a broader stream of research in marketing on the automatic processing of facial imagery. We suspect that this research will reveal that, rather than being one environmental prime among many, faces are actually the recipients of a dramatically disproportionate amount of our automatic attentional resources. Indeed, we believe that faces may actually be the most primal of primes.

REFERENCES

Adolphs, Ralph, Daniel Trane, and Antonio R. Damasio (1998), "The Human Amygdala in Social Judgment," *Nature*, 393 (June), 470–74.

Alicke, Mark D., M. L. Klotz, David L. Breitenbecher, Tricia J. Yurak, and Debbie S. Vredenburg (1995), "Personal Contact, Individuation, and the Better-Than-Average Effect," *Journal of Personality and Social Psychology*, 68 (5), 804–25.

Amaral, David G. (2002), "The Primate Amygdala and the Neurobiology of Social Behavior: Implications for Understanding Social Anxiety," *Biological Psychiatry* 51 (1), 11–17.

Ambler, Tim (1997), "How Much of Brand Equity Is Explained by Trust?" *Management Decision*, 35 (4), 283–92.

Bailenson, Jeremy N., Shanto Iyengar, Nick Yee, and Nathan A. Collins (2009), "Facial Similarity between Voters and Candidates Causes Influence," *Public Opinion Quarterly*, 72 (5), 935–61.

Bargh, John A. (1997), "The Automaticity of Everyday Life," in *Advances in Social Cognition*, Vol. 10, ed. Robert S. Wyer Jr., Mahwah, NJ: Erlbaum, 231–46.

Bargh, John A., and Tanya L. Chartrand (1999), "The Unbearable Automaticity of Being," *American Psychologist* 54 (July), 462–79.

——— (2000), "Studying the Mind in the Middle: A Practical Guide to Priming and Automaticity Research," in *Handbook of Research Methods in Social Psychology*, ed. Harry T. Reis

and Charles M. Judd, New York: Cambridge University Press, 253–85.

Bornstein, Robert F., and Paul R. D'Agostino (1994), "The Attribution and Discounting of Perceptual Fluency: Preliminary Tests of a Perceptual Fluency/Attributional Model of the Mere Exposure Effect," *Social Cognition*, 12 (Summer), 103–28.

Cacioppo, John T. (2004), "Asymmetries in Affect Laden Information Processing," in *Perspectivism in Social Psychology: The Yin and Yang of Scientific Process*, ed. John T. Jost, Mahzarin R. Banaji, and Deborah A. Prentice, Washington, DC: American Psychological Association, 85–95.

Cacioppo, John T., Wendi L. Gardner, and Gary G. Berntson (1999), "The Affect System Has Parallel and Integrative Processing Components: Form Follows Function," *Journal of Personality and Social Psychology*, 76 (5), 839–55.

Carré, Justin M., Cheryl M. McCormick, and Catherine J. Mondloch (2009), "Facial Structure Is a Reliable Cue of Aggressive Behavior," *Psychological Science*, 20 (10), 1194–98.

Chartrand, Tanya L., Joel Huber, Baba Shiv, and Robin J. Tanner (2008), "Nonconscious Goals and Consumer Choice," *Journal of Consumer Research*, 35 (August), 189–201.

Chaudhuri, Arjun, and Morris B. Holbrook (2001), "The Chain of Effects from Brand Trust and Brand Affect to Brand Performance: The Role of Brand Loyalty," *Journal of Marketing*, 65 (April), 81–93.

Chen, Mark, and John A. Bargh (1999), "Consequences of Automatic Evaluation: Immediate Behavioral Predispositions to Approach or Avoid the Stimulus," *Personality and Social Psychology Bulletin*, 25 (2), 215–24.

Corwin, G. H. (1921), "The Involuntary Response to Pleasantness," *American Journal of Psychology*, 32 (4), 563–70.

Critchley, Hugo, Eileen Daly, Mary Phillips, Michael Brammer, Edward Bullmore, Steven Williams, Therese Van Amelsvoort, Dene Robertson, Anthony David, and Declan Murphy (2000), "Explicit and Implicit Neural Mechanisms for Processing of Social Information from Facial Expressions: A Functional Magnetic Resonance Imaging Study," *Human Brain Mapping*, 9 (2), 93–105.

DeBruine, Lisa M. (2002), "Facial Resemblance Enhances Trust," *Proceedings of the Royal Society, Series B*, 269 (1498), 1307–12.

——— (2004a), "Resemblance to Self Increases the Appeal of Child Faces to Both Men and Women," *Evolution and Human Behavior*, 25 (2), 142–54.

——— (2004b), "Trustworthy but Not Lust-Worthy: Context Specific Effects of Facial Resemblance," *Proceedings of the Royal Society, Series B*, 272 (1566), 919–22.

De Houwer, Jan, Sarah Thomas, and Frank Baeyens (2001), "Association Learning of Likes and Dislikes: A Review of 25 Years of Research on Human Evaluative Conditioning," *Psychological Bulletin*, 127 (6), 853–69.

Devine, Patricia G. (1989), "Stereotypes and Prejudice: Their Automatic and Controlled Components," *Journal of Personality and Social Psychology*, 56 (January), 5–18.

Dion, Karen, Ellen Berscheid, and Elaine Walster (1972), "What Is Beautiful Is Good," *Journal of Personality and Social Psychology*, 24 (3), 285–90.

Doney, Patricia M., and Joseph P. Cannon (1997), "An Examination of the Nature of Trust in Buyer-Seller Relationships," *Journal of Marketing*, 61 (April), 35–51.

Ekman, Paul (1992), "Facial Expressions of Emotion: New Findings, New Questions," *Psychological Science*, 3 (1), 34–38.

Elliot, Andrew J. (2006), "The Hierarchical Model of Approach-Avoidance Motivation," *Motivation and Emotion*, 30 (2), 111–16.

Elliot, Andrew J., and Martin V. Covington (2001), "Approach and Avoidance Motivation," *Educational Psychology Review*, 13 (2), 73–92.

Engell, Andew D., James V. Haxby, and Alexander Todorov (2007), "Implicit Trustworthiness Decisions: Automatic Coding of Face Properties in Human Amygdala," *Journal of Cognitive Neuroscience*, 19 (9), 1508–19.

Fiske, Susan T., Amy J. Cuddy, and Pete Glicke (2007), "Universal Dimensions of Social Cognition: Warmth and Competence," *Trends in Cognitive Sciences*, 11 (2), 77–83.

Fitzsimons, Grainne M., Tanya L. Chartrand, and Gavan J. Fitzsimons (2008), "Automatic Effects of Brand Exposure on Motivated Behavior: How Apple Makes You Think Different," *Journal of Consumer Research*, 35 (June), 21–35.

Friestad, Marian, and Peter Wright (1994), "The Persuasion Knowledge Model: How People Cope with Persuasion Attempts," *Journal of Consumer Research*, 21 (June), 1–31.

Hansen, Christine H., and Ronald D. Hansen (1988), "Finding the Face in the Crowd: An Anger Superiority Effect," *Journal of Personality and Social Psychology*, 54 (6), 917–24.

Hassin, Ran R., James S. Uleman, and John A. Bargh, eds. (2005), *The New Unconscious*, New York: Oxford University Press.

Hatfield, Elaine, John T. Cacioppo, and Richard L. Rapson (1993), "Emotional Contagion," in *Review of Personality and Social Psychology*, Vol. 14, *Emotion and Social Behavior*, Newbury Park, CA: Sage, 151–77.

Hauber, Mark E., and Paul W. Sherman (2001), "Self-Referent Phenotype Matching: Theoretical Considerations and Empirical Evidence," *Trends in Neuroscience*, 24 (10), 609–16.

Howard, Daniel J., and Charles Gengler (2001), "Emotional Contagion Effects on Product Attitudes," *Journal of Consumer Research*, 28 (September), 189–201.

Ives, Nat (2008), "Despite Talk of Ethics Codes, Airbrushing Is Here to Stay," *Advertising Age* 79 (April 21), 24.

Kanwisher, Nancy, and Galit Yovel (2006), "The Fusiform Face Area: A Cortical Region Specialized for the Perception of Faces," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 361 (November), 2109–28.

Kim, Moonja, and Seymour Rosenberg (1980), "Comparison of Two Structural Models of Implicit Personality Theory," *Journal of Personality and Social Psychology*, 38 (3), 375–89.

Kunst-Wilson, William R., and Robert B. Zajonc (1980), "Affective Discrimination of Stimuli That Cannot Be Recognized," *Science*, 207 (February), 557–58.

Langlois, Judith H., and Lori A. Roggman (1990), "Attractive Faces Are Only Average," *Psychological Science*, 1 (2), 115–21.

Lewin, Kurt (1935), *A Dynamic Theory of Personality*, New York: McGraw-Hill.

McCracken, Grant (1989), "Who Is the Celebrity Endorser? Cultural Foundations of the Endorsement Process," *Journal of Consumer Research*, 16 (3), 310–21.

Mehta, Ravi, and Rui (Juliet) Zhu (2009), "Blue or Red? Exploring the Effect of Color on Cognitive Task Performances," *Science* 323 (5918), 1226–29.

Neumann, Roland, and Fritz Strack (2000), "Approach and Avoidance: The Influence of Proprioceptive and Extereoceptive Cues on Encoding of Affective Information," *Journal of Personality and Social Psychology*, 79 (1), 39–48.

Nummenmaa, Lauri, Jukka Hyona, and Jari K. Hietanen (2009), "I'll Walk This Way: Eyes Reveal the Direction of Loco-

motion and Make Passersby Look and Go the Other Way," *Psychological Science*, 20 (12), 1454–58.

Ohanian, Roobina (1990), "Construction and Validation of a Scale to Measure Celebrity Endorsers' Perceived Expertise, Trustworthiness, and Attractiveness," *Journal of Advertising*, 19 (3), 39–52.

Oosterhof, Nikolaas N., and Alexander Todorov (2008), "The Functional Basis of Face Evaluation," *Proceedings of the National Academy of Sciences, USA*, 105 (32), 11087–92.

Osgood, Charles E., George J. Suci, and Percy H. Tannenbaum (1957), *The Measurement of Meaning*, Urbana: University of Illinois Press, 1957.

Park, Justin H., and Mark Schaller (2005), "Does Attitude Similarity Serve as a Heuristic Cue for Kinship? Evidence of an Implicit Cognitive Association," *Evolution and Human Behavior*, 26 (2005), 158–70.

Parr, Lisa A., and Frans B. M. de Waal (1999), "Visual Kin Recognition in Chimpanzees," *Nature*, 399 (June), 647–48.

Schwarz, Norbert (2004), "Metacognitive Experiences in Consumer Judgment and Decision Making," *Journal of Consumer Psychology*, 14 (4), 332–48.

Shanteau, James, and Geraldine F. Nagy (1979), "Probability of Acceptance in Dating Choice," *Journal of Personality and Social Psychology*, 37 (4), 522–33.

Small, Deborah A., and Nicole M. Verrochi (2009), "The Face of Need: Facial Emotion Expression on Charity Advertisements," *Journal of Marketing Research*, 46 (December), 777–87.

Tanner, Robin J., Rosellina Ferraro, Tanya L. Chartrand, James R. Bettman, and Rick Van Baaren (2008), "Of Chameleons and Consumption: The Impact of Mimicry on Choice and Preferences," *Journal of Consumer Research*, 34 (April), 754–66.

Todorov, Alexander, and Andrew D. Engell (2008), "The Role of the Amygdala in Implicit Evaluation of Emotionally Neutral Faces," *Social, Cognitive, and Affective Neuroscience*, 3 (4), 303–12.

Todorov, Alexander, Anesu N. Mandisodza, Amir Goren, and Crystal C. Hall (2005), "Inferences of Competence from Faces Predict Election Outcomes," *Science*, 308 (5728), 1623–26.

Todorov, Alexander, Manish Pakrashi, and Nikolaas N. Oosterhof (2009), "Evaluating Faces on Trustworthiness after Minimal Time Exposure," *Social Cognition*, 27 (6), 813–33.

Verosky, Sara C., and Alexander Todorov (2010), "Generalization of Affective Learning about Faces to Perceptually Similar Faces," *Psychological Science*, 21 (6), 779–85.

Weisbuch, Max, and Diane Mackie (2009), "False Fame, Perceptual Clarity, or Persuasion? Flexible Fluency Attribution in Spokesperson Familiarity Effects," *Journal of Consumer Psychology*, 19 (2009), 62–72.

Weisbuch, Max, Diane Mackie, and Teresa Garcia-Marques (2003), "Prior Source Exposure and Persuasion: Further Evidence for Misattribution Processes," *Personality and Social Psychology Bulletin*, 29 (6), 691–700.

Wells, Peter A. (1987), "Kin Recognition in Humans," in *Kin Recognition in Animals*, ed. David J. C. Fletcher and Charles D. Michener, Chichester: Wiley, 395–415.

Whalen, Paul J. (1998), "Fear, Vigilance, and Ambiguity: Initial Neuroimaging Studies of the Human Amygdala," *Current Directions in Psychological Science* 7 (December), 177–88.

Willis, Janine, and Alexander Todorov (2006), "First Impressions: Making Up Your Mind after 100 Milliseconds Exposure to a Face," *Psychological Science*, 17 (7), 592–98.

Winkielman, Piotr, and John T. Cacioppo (2001), "Mind at Ease Puts a Smile on the Face: Psychophysiological Evidence That Processing Facilitation Increases Positive Affect," *Journal of Personality and Social Psychology*, 81 (6), 989–1000.

Winston, Joel S., Bryan A. Strange, John O'Doherty, and Raymond J. Dolan (2002), "Automatic and Intentional Brain Responses during Evaluation of Trustworthiness of Faces," *Nature Neuroscience*, 5 (3), 277–83.

Zajonc, Robert (1968), "Attitudinal Effects of Mere Exposure," *Journal of Personality and Social Psychology Monograph Supplement*, 9 (2, pt. 2), 1–28.