Aftereffects for Face Attributes With Different Natural Variability: Adapter Position Effects and Neural Models

Rachel Robbins, Elinor McKone, and Mark Edwards
Australian National University

Adaptation to distorted faces is commonly interpreted as a shift in the face-space norm for the adapted attribute. This article shows that the size of the aftereffect varies as a function of the distortion level of the adapter. The pattern differed for different facial attributes, increasing with distortion level for symmetric deviations of eye height and decreasing for asymmetric deviations. These results are interpreted in terms of different coding ranges for the 2 facial attributes, arising from differences in eye-height variability in natural face images (large for symmetric, small for asymmetric). Neural models developed in low-level vision also are applied to facial attributes, contrasting a 2-pool (norm-based) and a multichannel (exemplar-based) model. The adapter position effects generally support a norm-based model, as did a finding that perception of stimuli further from the norm than the adapter was shifted in the direction of the norm, rather than repulsed away from the adapter.

Keywords: face recognition, face space, adaptation, neural models, aftereffects

Adaptation aftereffects occur throughout the visual system, from early vision (e.g., line orientation and color processing) to high-level vision (including face recognition). For faces, “opposite” aftereffects have been demonstrated for many different distortion types and facial attributes. For example, adaptation to a contracted face makes an undistorted face appear expanded (Webster & MacLin, 1999; also see MacLin & Webster, 2001; Rhodes, Jeffery, Watson, Clifford, & Nakayama, 2003; Watson & Clifford, 2003), and adaptation to a male face makes a sex-neutral face appear female (Webster, Kaping, Mizokami, & Duhamel, 2004). Similar aftereffects have been shown for race (e.g., Caucasian–Japanese), expression (e.g., happy–sad; Hsu & Young, 2004; Webster et al., 2004), eye position (eyes in–out; Little, DeBruine, & Jones, 2005), and individual identity (Leopold, O’Toole, Vetter, & Blanz, 2001).

Such aftereffects are commonly explained in terms of a shift in the norm of a multidimensional “face space,” along a trajectory associated with the facial attribute(s) being distorted (e.g., Leopold et al., 2001; Rhodes et al., 2003). Face space is a heuristic used to explain the coding of face identity and distinctiveness in terms of deviations from the average (e.g., Valentine, 1991; Valentine & Bruce, 1986; see also Atick, Griffin, & Redlich, 1996; O’Toole,

Abdi, Deffenbacher, & Valentin, 1995). Each individual face is coded in terms of its position on a series of underlying dimensions, such that the average face is located at the center of face space and atypical individuals lie toward the periphery on one or more dimensions. The nature of the dimensions is not known (e.g., they could be simple attributes such as eye height and face width or more complex attributes such as the eigenfaces produced by principal-components analysis), but the presumption is that they are built up from ongoing experience with faces and will thus reflect the natural image statistics of faces, including the amount of variability of any given attribute in the face population.

To understand how adaptation has been explained in terms of face-space norm shifts, consider the identity adaptation procedure of Leopold et al. (2001; see also Leopold, Rhodes, Muller, & Jeffery, 2005; Little et al., 2005; Rhodes, Jeffery, & Leopold, 2004 [described in Rhodes et al., 2005]). Here, face–antiface pairs were created by taking four real individuals and morphing them on a trajectory through the center of face space (the average face determined from many individuals) and past it an equal distance on all face attributes in the opposite direction. Thus, for example, if Bill had eyebrows that were 20% more bushy than the average and lips 15% narrower, anti-Bill would have eyebrows 20% less bushy than the average and lips 15% wider. Results showed that adapting to anti-Bill made the average face look like Bill. This was explained as the norm having shifted toward the adapter along the trajectories describing the relevant facial attributes (i.e., toward less bushy eyebrows and wider lips).

Possible Neural Models of Face Distortion Aftereffects

In low-level and midlevel vision, adaptation and adaptation aftereffects have been used extensively to explore neural processing mechanisms (e.g., Blakemore & Campbell, 1969; Clifford, Wenderoth, & Spehar, 2000; Kourtzi & Kanwisher, 2001; Rogers & Graham, 1985; Smith & Edgar, 1994). Several researchers have suggested that aftereffects may similarly provide insight into neu-

Rachel Robbins, Elinor McKone, and Mark Edwards, School of Psychology, Australian National University, Canberra, Australia.

This work was supported by Australian Research Council Grants DP0208630 (to Elinor McKone) and DP0450636 (to Elinor McKone and Mark Edwards). A subset of the data from Experiment 1 have been summarized briefly as part of a book chapter (Rhodes et al., 2005), attributed as Robbins and McKone (2004; unpublished raw data).

We thank Colin Clifford and Gillian Rhodes for helpful discussions regarding two-pool neural models of adaptation. We also thank Isabel Gauthier for her helpful comments.

Correspondence concerning this article should be addressed to Rachel Robbins, who is now at the Department of Psychology, Neuroscience & Behaviour, McMaster University, Hamilton, Ontario L8S 4K1, Canada. E-mail: dr.r.robbins@gmail.com

ral mechanisms of face processing (e.g., Leopold et al., 2001; Rhodes et al., 2003). However, with the exception of some preliminary work we described in Rhodes et al. (2005), no actual models have been put forth. One aim of the current article is to begin the process of considering neural models of face aftereffects. We start by comparing two approaches that have been commonly used for perceptual aftereffects in low- and midlevel vision, namely, a two-pool model and a multichannel model.

Traditionally, two-pool models have been considered relevant for explaining aftereffects for which a perceptual dimension appears to behave in an opponent fashion, with a neutral point or norm falling between two ends of a dimension (e.g., Over, 1971; Regan & Hamstra, 1992; Rogers & Graham, 1985). This situation can occur in perception of depth (a near–far dimension with the fixation plane as the neutral point), motion (e.g., upward–downward with stationary as the neutral point), color (e.g., red–green with white as the neutral point), and various aspects of simple shape (e.g., vertical–horizontal elongation with square or circular as the neutral point). In contrast, multichannel models have been considered appropriate for cases in which a dimension does not have a neutral point. This occurs in perception of spatial frequency (where no particular spatial frequency has special status) and line orientation (where although more neurons code vertical and horizontal than other orientations, neither has special status for adaptation aftereffects; addressed below).

Two-pool models of aftereffects are based on three simple assumptions. First, at some level in the visual pathway, the coding of the aspect in question relies on two broadly tuned mechanisms that respond maximally to stimuli on opposite sides of the norm. For example, as illustrated in Figure 1, Regan and Hamstra (1992) explained coding of aspect ratio by presuming that neurons are organized functionally into two pools, one of which is preferentially excited by vertically elongated stimuli (Pool 1) and the other of which is preferentially excited by horizontally elongated stimuli (Pool 2).

Second, the neural tuning functions are overlapping so that any given stimulus produces activation in both pools. Unless the stimulus is at the norm, this activation will be stronger in one pool than the other. Perception is then determined by considering the balance of the output from the two pools. Thus, for example, a vertically elongated stimulus is perceived as “tall and thin” because Pool 1 is firing more than Pool 2; a square is perceived as square because Pool 1 and Pool 2 are firing equally; and a horizontally elongated stimulus is perceived as “short and fat” because Pool 2 is firing more than Pool 1.

The third basic assumption is that adaptation causes each pool to reduce its firing rate in proportion to the strength of its initial unadapted response to the adapting stimulus; that is, the pool that is most sensitive to that stimulus becomes the most adapted. This assumption is supported by physiological data (e.g., Maddess, McCourt, Blakeslee, & Cunningham, 1988; Movshon & Lennie, 1979). After adaptation to a horizontally elongated stimulus, therefore, the response of Pool 2 will be reduced more than the response of Pool 1. This will shift the point of equal response levels—and the physical stimulus now perceived as being square—in the direction of horizontal elongation.

Multichannel models differ primarily in the first assumption. It is presumed that a dimension is coded by many pools of neurons, each tuned to slightly different stimulus values. For example, each

Figure 1. Two-pool neural mechanism for coding an opponent dimension, with an explanation of how the adaptation aftereffect occurs. Two pools of cells exist, coding deviations on each side of the norm. The crossover point of these curves (indicating equal firing from both pools) gives the stimulus perceived as most normal. Solid lines show the response of pools of cells without adaptation. Dashed lines show the response after adaptation: The perceived norm shifts toward the adapter because Pool 2 is reduced by a larger amount than Pool 1. The notations A1–A4 mark different possible adapter positions: A1–A3 = small, moderate, and extreme deviations, respectively, that fall within the coding range; A4 = an extreme deviation that falls outside of the primary coding range.

pool of neurons coding line orientation responds maximally to one particular orientation, and then the full range of possible orientations is covered by having many narrowly tuned pools with different peak tunings (Hubel & Wiesel, 1974). The perception of a given stimulus is determined by the combined output of all pools (or “channels”) that respond to that stimulus. Adaptation again reduces the response of these pools in proportion to their initial firing rate, and this can shift the perceived orientation of a stimulus away from the adapting orientation.

Two-pool and multichannel models can be distinguished via two key predictions. The first arises from the fact that a two-pool system has a true norm or neutral point, whereas a standard multichannel model does not. Consider test stimuli that are further from the proposed neutral point than the adapter. In depth perception, for example, this could correspond to adapting to a moderately “far” stimulus (i.e., one appearing as moderately behind the fixation plane) and then testing with a very far stimulus (i.e., appearing further behind fixation than the adapter). A two-pool model predicts that this test stimulus will appear shifted toward the direction of the norm (i.e., shifted in the “near” direction so that it appears less far behind fixation than before), owing to the balance

shift of the two pools. This was indeed the result obtained for depth perception by Rogers and Graham (1985). In contrast, a multichannel model predicts that the very far test stimulus would appear to have been shifted away from the adapter position (i.e., further in the far direction). This type of pattern is observed for line orientation, where although adaptation to a slightly right-tilted line makes upright seem left tilted (potentially consistent with upright as a neutral point), testing with a more right-tilted line gives a percept of greater right-tilt again (Gibson, 1937).¹

The second prediction concerns the effect of the position of the adapter. In multichannel models, at least where each pool has a fairly narrowly tuned range of stimulus values to which it responds (as in coding of line orientation), only test stimuli relatively similar to the adapter will show an aftereffect, because adapters further away will not affect firing rates in the pools coding the test stimulus. In contrast, two-pool models assume very broad tuning (see Figure 1), and so adapters very far from the test stimulus can still shift perception. Indeed, in terms of measuring the amount of shift specifically at the norm position, two-pool models make a prediction that the size of the aftereffect will increase as the adapter is placed further from the norm, at least within the coding range of the two neural pools. This is because the change in the relative output of the two pools due to adaptation—and thus the amount of shift in the crossover point—will be greatest when the adapter is located in a region producing very different initial responses from each pool (e.g., very strong response from Pool 1, very weak from Pool 2). To understand this, assume the response curve shapes are as illustrated in Figure 1. With a small deviation adapter placed at A1, the shift in the norm after adaptation will be relatively small, because the vertical and horizontal pools are initially firing at only moderately different levels, meaning that the reduction in firing rate for the vertical pool will be only a small amount more than for the horizontal pool. For a moderate deviation adapter placed at A2, the shift in norm will be larger, because the initial difference in firing rate between the two pools is larger. For an extreme deviation adapter placed at A3, the initial difference is larger again, and indeed, the firing of the horizontal pool would barely be affected whereas the vertical pool would be strongly affected, leading to the largest possible shift in the norm.

In summary, a two-pool model can be distinguished from a multichannel model on the basis of the predictions that (a) perception of stimuli further from the norm than the adapter should always be shifted in the same direction as those falling between the adapter and the norm, rather than being repulsed away from the adapter, and (b) the size of the aftereffect for test stimuli at the original norm should increase as the adapter is made more extreme within the coding range of the neural pools. Neither of these predictions has ever been tested for faces. In previous studies, only one adapter position has been used, generally a rather extreme one, and test stimuli have always been closer to the norm than the adapter (e.g., Rhodes et al., 2003, obtained normality ratings for radial expansions and contractions ranging from -50% to 50%, and the adapters were either -50 or 50).

Another question that has not previously been explored is the effect of the exact coding range on the size of the aftereffects. As noted above, in a multichannel model the coding range of each pool of neurons is narrow, and so an extreme adapter would be expected to have little effect on perception of an undistorted face (in apparent conflict with the extant data). Within a two-pool

approach, however, there is a genuine question as to what would be predicted. In Figure 1, we have left unspecified what happens outside the range covered by the two response curves drawn (as did Regan & Hamstra, 1992). If we were to assume that the maximum aspect ratio coded by the "vertical" pool was height 100 times width, then what effect would arise from an adapter with a height of 300 times width? One possibility is that outside the primary coding range (e.g., with an adapter placed at position A4 in Figure 1), the pool response drops rapidly to zero; this would predict no adaptation aftereffect. Alternatively, the response curves might turn around and approach zero more slowly; this would predict reducing but nonzero aftereffects.

Varying Coding Range via Symmetric and Asymmetric Distortions of Eye Height

A reasonable assumption is that the coding range of different facial dimensions will be tuned to natural image statistics (cf. Simoncelli & Olshausen, 2001, in low-level vision), that is, to the variability of a given facial attribute in the face images to which a subject's perceptual system has been historically exposed. Two aspects of faces that differ substantially in natural variability are symmetric and asymmetric eye-height distortions. These are illustrated in Figure 2A. In the symmetric distortion, an original face (Bill in Figure 2B) has been altered by moving the two eyes up together or the two eyes down together. An adaptation aftereffect for this distortion manifests as, for example, the original undistorted face appearing to have its eyes too low after adapting to an eyes-up face (Carbon & Leder, 2006). (Readers can observe this phenomenon for themselves in the figure.) In the asymmetric distortion, the original face has been altered by moving either the right eye up and the left eye down or vice versa; again, this manipulation produces an adaptation aftereffect (also see Morikawa, 2005).

Physical measurements of faces show that these two distortions correspond to different amounts of variability in the population. For Caucasians, the distance between the mean position of both eyes and the hairline in young adult males has a standard deviation of 7.5 mm (around a mean of 67.1 mm), and the distance from eyes to chin has a standard deviation of 5.7 mm (around a mean of 124.7 mm; Farkas, Hreczko, & Katic, 1994). In contrast, the asymmetry between eyes varies to a much smaller degree: The

¹ Note that there are also more complicated ways to create a neutral point than a simple two-pool model. Motion perception arguably has a neutral point of zero motion (Weiss, Simoncelli, & Adelson, 2002) but is encoded by many cells that are tuned to a range of directions, and the tuning functions of these cells to direction are narrowband and nonoverlapping. The direction of motion is determined by activity pooled across all of these motion cells, not just those cells coding opposite directions (e.g., up and down; Mather, 1980). Outputs of cells tuned to similar directions of motion are positively pooled, and those tuned to markedly different directions are negatively pooled (i.e., they inhibit each other; Wilson, Ferrera, & Yo, 1992), allowing for the motion aftereffect. The difference between motion and orientation processing, a process that also uses many cells with narrowband and nonoverlapping tuning functions but does not have a neutral point, is how the outputs of the cells are pooled. With orientation processing, there is only pooling of cells that are tuned to similar orientations.

Figure 2. A: Symmetric and asymmetric distortions of Bill (shown undistorted in B). The three distortion levels (± 5 , ± 20 , and ± 50 pixels per eye) correspond to the three adapter positions used in our experiments. Readers should be able to observe the opponent aftereffect by scanning one of the distorted faces for 30 s and then looking at the undistorted face (e.g., after adapting to eyes up, the normal face seems to have its eyes down). The aftereffect is largest for the symmetric ± 50 adapter. B: The 4 individuals used in the experiments. All 4 were used as test faces (pre- and postadaptation phases). Only Bill and Sam appeared as adapters. R = right; L = left.

height difference between left and right eyes in adult males has a standard deviation of only 1.1 mm (around a mean difference of 2.8 mm; Hreczko & Farkas, 1994). These differences in front-on views of faces are further exacerbated by the effects of head rotation. Information about up-down rotation of the head (e.g., nodding) is known to be neurally coded (e.g., Perrett et al., 1985), and this movement substantially changes the apparent position of eyes with respect to other features. In contrast, there is no three-dimensional head movement that can increase the physical asymmetry between the eyes, relative to the main axis of the face, in the manner shown in Figure 2A (i.e., tilting the head sideways creates an asymmetry relative to true vertical, but not relative to the face axis, and also tilts the eyes themselves).

These observations suggest that the range of eye-height values coded will be very different for symmetric and asymmetric eye-height attributes. For symmetric distortions, considering both the differences arising from individual face structure and the even

larger differences arising from three-dimensional up-down head rotation, we suggest that the required range probably covers from the eyes being very close to the tip of the nose (looking down) to the eyes being very close to the hairline (looking up). The required coding range for asymmetric distortions of eye height should be much smaller. In the natural population of faces it would be unusual to observe anything more than a relatively small deviation, and head rotation does not increase this level. Our presumption is that the coding range for asymmetric distortions might be not much more than enough to cover the most extreme natural individual seen in a front-on view.

The choice of an eye-height-based manipulation of coding range has several advantages. First, the distortion type is a relational change that alters the distances between features without altering feature shape itself. Given that the norm for eye position can be defined only with reference to the rest of the face, such relational distortion types may help to locate aftereffects as arising from face

processing as opposed to midlevel vision. In previous studies, in contrast, aftereffects such as those arising to radially expanded and contracted faces (e.g., Rhodes et al., 2003) could arise partly from adaptation of midlevel mechanisms such as convexity and concavity (Suzuki, 2001, 2003). (It is standard to rule out contributions from low-level vision by changing stimulus size between adapter and test; see Zhao & Chubb, 2001.)

Second, the level of distortion can be matched across the symmetric and asymmetric eye-height deviations for the total amount of metric change from the original face. For example, a 50-pixel-per-eye distortion in the symmetric condition (e.g., each eye moved up by 50 pixels) could be matched to a 50-pixel-per-eye distortion in the asymmetric condition (e.g., left eye up by 50 pixels, right eye down by 50 pixels). This means that adapter position can be varied in a coherent fashion across the two distortion types. Deviation level can be matched on the same scale in absolute terms (i.e., number of pixels). Of importance, it is then also matched in terms of the proportion of the distance to the edge of the face (e.g., distance to hairline or chin). It is not known how feature position is represented, and other possible deviation types that could potentially be matched for physical distortion (e.g., eyes up-down with eyes in-out) could not simultaneously be matched for distance relative to other landmarks in the face (e.g., head outline, nose position).

Finally, a previous study (McKone, Aitkin, & Edwards, 2005) has established that symmetric and asymmetric eye-height deviations produce effects on perception that track the natural variability associated with each face attribute, thus supporting the idea that the neural coding range for each attribute must differ. In McKone et al. (2005), no adaptation was used. Unadapted stimuli were presented at multiple levels of distortion ranging from eye(s) down almost to the tip of the nose to eye(s) up, approaching the hairline (as in Figure 2A), and subjects rated normality of the face. The relationship between physical deviation of eye height and perceived normality followed a Fechner law function (Fechner, 1860/1966), in which a certain amount of deviation from the original face could be tolerated without affecting perceived normality at all (the threshold value), beyond which a log relationship (Weber's law) related perception to physical deviation.² The relevant finding was that although the same general Fechner law pattern was obtained for both symmetric and asymmetric deviations, subjects were much less sensitive to symmetric deviations: Threshold deviation levels to perceive a change in normality from the original were 3.4% of the hairline-chin distance for symmetric deviations and only 0.7% for asymmetric, and the slope of the log function above threshold was also lowest for symmetric. These results are consistent with the idea that a face-space trajectory associated with symmetric eye-height differences codes a large range of eye-height values, whereas a trajectory associated with asymmetric eye-height differences codes a smaller range of eye-height values.

Present Study

Across two experiments, we used symmetric and asymmetric eye-height distortions to explore aftereffect patterns relevant to two-pool versus multichannel models. We tested the two-pool prediction that perception of face distortions further from the norm than the adapter would be shifted toward the adapter direction (e.g., after adaptation to 20 pixels eyes up, a 30-pixel-eyes-up test

stimulus would be seen as having its eyes lower, or more normal, than prior to adaptation) rather than repulsed away from the adapter (30 pixels eyes up seen as higher, or less normal).

We also explored the effects of adapter position. The most extreme adapters had deviation levels of ± 50 pixels per eye, corresponding to eyes approaching tip of nose or hairline, as illustrated in Figure 2A. We also tested adapter deviations of ± 20 pixels per eye and adapter deviations of ± 5 pixels per eye (see Figure 2A). If the coding ranges we have suggested for symmetric and asymmetric eye-height deviations are correct, then for symmetric distortions, all three adapter positions will fall within the coding range (e.g., in a two-pool model, their location relative to the neural response curves might be similar to A1, A2, and A3 in Figure 1). In contrast, for asymmetric distortions, the ± 5 adapter could well be within the coding range (e.g., a location somewhere in the region A1–A3 in Figure 1), whereas ± 20 and ± 50 adapters would clearly be outside it (location similar to A4 in Figure 1). A two-pool model would then predict that for symmetric deviations, the size of the adaptation aftereffect should increase as the adapter is moved from ± 5 through ± 20 to ± 50 pixels. However, for asymmetric deviations, this pattern should not occur, and the aftereffect could decrease or drop immediately to zero as the adapter is moved beyond ± 5 pixels (depending on the assumptions made about the shape of the tuning function outside the primary coding range). Another prediction arises from the fact that the ± 20 and ± 50 adapter positions should be outside the coding range of asymmetric eye-height distortions but inside it for symmetric distortions. For these values, the size of the aftereffect should then be larger for symmetric than for asymmetric distortions.

In each experiment, a three-phase procedure was used: a preadaptation phase; an adaptation phase in which a distorted face was presented for 2 min; and a postadaptation phase identical to the preadaptation phase, with the addition of top-up adaptation stimuli presented for 5 s between each of the test stimuli. To avoid a contribution of low-level retinotopic aftereffects, we presented the adapting faces at a smaller size than the test faces. Similarly, subjects were instructed to make eye movements around the adapter during the adaptation phase, and top-up adapter faces were presented at several locations to require eye movements during the postadaptation test.

In Experiment 1, the adapter position was set to ± 20 pixels, and subjects rated many levels of deviation for how much the stimulus "looked like the original person." This allowed testing of perceptual shifts on both sides of the adapter position; indeed, we were able to test the two-pool model prediction that the Fechner curves relating physical to perceived deviation in the unadapted state (McKone et al., 2005) should show a simple shift left or right along the deviation level axis, moving toward the direction of the adapter. Because this procedure was very time consuming to test,

² A standard example of Fechner's law from early vision is the relationship between physical and perceived light intensity. Here, a very small amount of light added to complete darkness is not perceived at all (the below-threshold region). Once above threshold, perceived intensity increases logarithmically with physical intensity (e.g., lighting a single candle has more perceptual effect in a semidarkened room than in bright sunlight).

we tested a small number of subjects and analyzed individual subject data.

In Experiment 2, we used an adjustment task, in which subjects adjusted the eye height of a starting point stimulus until it "looked like the original person." This technique measures adaptation effects on the position of the perceived norm only, rather than assessing perception for a full range of stimuli including those more deviated than the adapter, but was quicker to test, allowing comparison of all three adapter positions with data averaged over multiple subjects.

A final point is that aftereffects were examined for same-identity and different-identity pairs. Of the four individuals used as test faces (Bill, Sam, John, Fred), only two (Bill and Sam) were ever used as the adapters, allowing us to compare aftereffects for adapted-identity individuals (Bill and Sam) and unadapted-identity individuals (John and Fred). Previous studies have found that adaptation generalizes across identity change (e.g., Rhodes et al., 2003) and that transfer relative to a same-identity condition tends to be strongest when the faces are similar in overall appearance (e.g., have similar hair; Webster & MacLin, 1999; Yamashita, Hardy, De Valois, & Webster, 2005).

Experiment 1: Full Rating Curves

In Experiment 1, 27 levels of the symmetric eye-height distortion and 27 levels of the asymmetric eye-height distortion were used as test stimuli. The amount of distortion in each face was defined by the amount of change per eye; thus, a 50-pixel-per-eye symmetric distortion (e.g., 50-pixel right-eye shift up and 50-pixel left-eye shift up) was matched for total metric deviation to the 50-pixel-per-eye asymmetric distortion (e.g., 50-pixel right-eye shift up and 50-pixel left-eye shift down). All four test individuals (Bill, Sam, John, and Fred) were presented at each of these deviation levels, and subjects rated how much the stimulus looked like the original undistorted person (e.g., how much it "looked like Bill"). Deviation direction was coded as positive for eyes-up (symmetric) and right-up/left-down (asymmetric) and as negative for eyes-down (symmetric) and right-down/left-up (asymmetric). The adapter position was an intermediate value of ± 20 pixels per eye. Only Bill and Sam appeared as the adapters.

The predictions derived from a two-pool model of face aftereffects for this experiment were as follows. First, adaptation should produce a simple left-right shift in the curves relating rated likeness to physical deviation; that is, all faces on the adapter side of the original norm (regardless of which side of the adapter they lie on) should be perceived as "more like Bill" than previously, and all faces on the nonadapter side of the original norm should be perceived as "less like Bill" than previously. Second, if the coding ranges we proposed for each face attribute are correct, the amount of shift in the face rated as "most like Bill" (i.e., the size of the aftereffect) should be larger for symmetric than for asymmetric distortions, given the use of the ± 20 -pixel adapter.

The primary theoretical focus was on aftereffects for upright faces, using results for a condition in which both the adapter and the test stimuli (pre- and postadaptation) were shown in this orientation (U-U). We also tested inverted faces and included an inverted-adapter/inverted-test condition (I-I), an inverted-adapter/upright-test condition (I-U) and an upright-adapter/inverted-test condition (U-I). Previous studies have found strong aftereffects in

the I-I conditions (Leopold et al., 2001; Watson & Clifford, 2003; Webster & MacLin, 1999) and rather smaller aftereffects in the two orientation-transfer conditions (U-I and I-U; Watson & Clifford, 2003; Webster & MacLin, 1999), indicating that the aftereffects for upright and inverted faces arise from largely different neural populations (also see Rhodes et al., 2004). We expected to replicate these findings. In total, the experiment required 10 hr per subject, and analysis was based primarily on single-subject data.

Method

Subjects

There were 5 subjects in total (age range = 20–42; 2 men, 3 women), 4 who were naive as to the purposes of the experiment (S1–S4) and took part in all conditions, and the first author (Rachel Robbins), who completed only the U-U conditions. All subjects were Caucasian and thus ethnically matched to the face stimuli. All reported normal or corrected-to-normal vision. Naive subjects were paid \$10 per session.

Design

The experiment began with a training phase in which subjects learned to recognize the four test individuals (Bill, Sam, John, and Fred; see Figure 2B) in their undistorted form. Each session, lasting 30–45 min, then contained the pretest, adaptation, and posttest phases for one distortion type (e.g., symmetric), adapter direction (e.g., negative), and orientation condition (e.g., U-U). Within each session, preadaptation ratings of how much a deviated face looked like the original individual ($1 = \text{looks exactly like [Bill/Sam/John/Fred]}$, $9 = \text{looks nothing like [Bill/Sam/John/Fred]}$) were determined for 27 deviation levels ($0, \pm 1, \pm 2, \pm 3, \pm 6, \pm 8, \pm 10, \pm 12, \pm 15, \pm 20, \pm 25, \pm 30, \pm 40, \text{ and } \pm 50$ pixels per eye), for all four individuals (Bill, Sam, John, and Fred), presented intermixed and in random order. Subjects were adapted to either $+20$ or -20 versions of the adapter faces (Bill and Sam) showing the same distortion type as the test stimuli. The ± 20 adapter stimuli are illustrated in Figure 2A. Postadaptation ratings were then taken for all four individuals again presented intermixed and in random order. ("Adapted" individuals were always Bill and Sam, and "unadapted" individuals were always John and Fred.)

For each subject, the U-U and I-I orientation conditions were tested first, with orientation, distortion type, and adapter direction counterbalanced as far as possible with only 5 subjects (sessions were run so that upright and inverted orientations alternated). After this, the I-U and U-I conditions were tested, again counterbalanced in order as far as possible (and with orientation alternated). The delay between sessions was at least 20 hr in all except 7 of the 68 cases, where it was 1–1.5 hr. Given recent evidence that face distortion aftereffects can endure for at least 80 min (Carbon & Leder, 2006), we checked our data for any evidence of carryover of adaptation between sessions testing a common distortion type. Preadaptation curves were not shifted in the direction that would be predicted by carryover effects of the preceding adapter.

Stimuli

All stimulus manipulations were carried out with Adobe Photoshop 5.5. Stimuli were derived from four undistorted 0-deviation

faces taken from McKone et al. (2005). These faces had been made by cutting the internal regions from photographs of four real people (gray-scale Caucasian faces obtained from the Stirling PICS database³) and pasting them onto a common background head derived from a fifth person. The natural distances between eyes, nose, and mouth were retained, and the overall set of internal features was placed at an average height within the new head. A common background head with stiff upright hair was used to ensure that eyes could be moved relative to a clearly seen hairline. The resulting faces were all natural in appearance.

From each of the four undistorted individuals, 27 deviation levels were made for each distortion type (symmetric and asymmetric). The total number of pixels from hairline to chin in the images was 263; thus, for example, the most extreme deviation of 50 pixels per eye corresponded to 19.01% of the hairline–chin distance, 20 pixels per eye to 7.60% of this distance, and 1 pixel per eye to 0.38% of this distance. At the viewing distance of 40 cm, test stimuli (i.e., those shown in the preadaptation and postadaptation phases) subtended 10° from the chin to hairline by 7.9° across the widest point of the cheeks (and thus a 10 pixel per eye deviation, for example, corresponded to a 0.4° position difference from the original, if fixating the original eye position). Adapter stimuli were shown smaller than the test stimuli, at 7.9° by 5.7°.

Procedure

Stimuli were presented on an iMac computer with a 36-cm screen set to a resolution of 1,024 × 768 pixels, using PsyScope software (Cohen, MacWhinney, Flatt, & Provost, 1993). A chin-rest was used.

Training phase. An identity training phase began with each 0-deviation face (Bill, Sam, John, and Fred) shown once upright and once inverted with its name presented beneath. This was followed by 32 trials of training with feedback (4 trials of each 0-deviation face in each orientation presented in random order). Here, stimuli were presented without names; subjects pressed the key matching the face's initial (B, S, J, or F), and a feedback beep sounded for incorrect responses.

At the beginning of his or her first session, each subject completed four cycles through this training phase (i.e., 40 exposures to each individual), taking approximately 15–20 min. Identification accuracy was measured during the training-with-feedback task, and all subjects passed a criterion of 80% correct (performance for upright faces was above 89%). At the beginning of each additional session, a single cycle through the training phase was given to ensure subjects had reinstated an active representation of each face in its undistorted form (accuracy in these was always above 94% overall).

Preadaptation test phase. In the preadaptation phase of each session, each of the four individuals (Bill, Sam, John, and Fred) was presented once at each of the 27 deviation levels, for a total of 108 trials. Face stimuli were presented at the center of the screen with no fixation point (this procedure avoided providing easy cues as to absolute location of the eyes on the screen, rather than within the face). Each trial comprised the stimulus face (1,000 ms), a blank screen (200 ms), and the prompt to enter an identity rating. The name of the person was always given in the prompt. For example, the prompt following a picture of Bill was "How much

does this look like Bill? 1 = *exactly*, 9 = *not at all*." Trials were presented in random order for each subject.

To encourage appropriate use of the full scale, we showed six examples from the relevant distortion type (symmetric or asymmetric) before the rating task began. Examples covered both mild and extreme deviation levels in both directions, using a nontested individual ("Ralph").

Adaptation phase. In the adaptation phase of each session, the relevant adapter deviation (e.g., 20 pixels eyes up, symmetric distortion) was shown for 2 min. The adapter individual was Bill for 1 min then Sam for 1 min, or vice versa, with the individual presented first chosen randomly. Adapters were presented in the center of the screen. Subjects were instructed not to fixate the adapter but instead to scan the whole face for the full exposure time.

Postadaptation test phase. The postadaptation phase in each session was a repeat of the preadaptation phase, with the addition of top-up adaptation. Before each rating trial began, one of the adapter individuals (Bill or Sam) was shown for 5,000 ms, with alternation of adapter identity on successive trials. The top-up adapter was shown in one of four locations around the test face, to force subjects to make eye movements to view the adapter. The locations were centered 5.7° vertical and 5.7° horizontal away from the test face, to the top-left, top-right, bottom-right, or bottom-left corner, with locations chosen in clockwise order on successive trials.

Results

Our primary interest was in the results for upright faces. These are considered first, with results for the other orientation conditions summarized at the end. Preliminary information on curve fitting is relevant to all orientations.

Curve Fits

In general, as face deviation increased away from 0 in either direction, subjects initially perceived no change in rated identity from the undistorted face. Outside this range, ratings were related to physical deviation in a logarithmic fashion; that is, for a unit increase in deviation, changes in perception were large immediately past the threshold but were smaller toward extreme deviation levels. Data were therefore fitted using a Fechner law (threshold-plus-log) function. As explained in Figure 3, this allowed extraction of five parameters of interest. (In reading all graphs, note again the direction of the identity rating scale: 1 = *most like*; 9 = *least like*.) The left threshold value (T_l) indicates the number of pixels per eye by which the stimulus can be deviated in the negative direction before ratings of identity start to suddenly increase. The right threshold value (T_r) is the corresponding value for positive deviations. The center of the full subthreshold region (c) provides the best estimate of the stimulus perceived as "most like" the individual. The final two parameters are the left slope at threshold (A_l) and the right slope at threshold (A_r); in each case, higher slopes indicate greater sensitivity to changes, this time in

³ This database can be found at <http://pics.psych.stir.ac.uk/>.

Figure 3. The Fechner law function used to fit the data.

the above-threshold region. All fits were done using MacCurveFit (Kevin Raner Software, 1993).⁴

Fechner law curves provided good fits to every subject's data, in preadaptation and postadaptation conditions, upright and inverted (fits were even better when data were averaged across a larger number of subjects in McKone et al., 2005). Figure 4A shows an example fit, for a single subject in a preadaptation symmetric distortion condition. For data in the asymmetric condition, the Fechner curves sometimes failed to capture the full degree of "pointedness" close to the zero deviation, as illustrated for one subject in Figure 4B. For this reason we also tried fitting Pearson IV curves (four-parameter modified Gaussians that allow more pointedness—kurtosis—than the standard three-parameter Gaussian). Although the overall quality of the fits was still good for Pearson IV, average R^2 did not increase as compared with the Fechner fits, and the ability to capture the pointedness did not improve (see Figure 4C). We chose to retain Fechner law fits, because they do not force the curve to be symmetrical (thus providing better fits in cases like Figure 4A) and because there is a well-established theoretical rationale behind Fechner's law. Conclusions regarding adaptation were the same from both fitting methods.

Confirming Greater Perceived Variability for Symmetric Than for Asymmetric Distortions (Unadapted Upright Faces)

We wished to confirm that in the preadaptation conditions, subjects were less sensitive to symmetric deviations of eye height than to asymmetric deviations (replicating McKone et al., 2005). To obtain the most reliable results, we averaged over all relevant unadapted scores for upright faces. This included collapsing over direction of future adapter (+20, -20) and identity of the face

(Bill, Sam, John, and Fred were all equivalent here because no adaptation had yet taken place), as well as including preadaptation trials from both the U-U condition (upright-adapter/upright-test) and I-U condition (inverted-adapter/upright-test). Figure 5 shows resulting ratings and fitted Fechner curves for symmetric (left panel) and asymmetric (right panel) distortions. The range of faces perceived as normal was larger for the symmetric distortion than for the asymmetric distortion (i.e., the subthreshold region was wider, indicating subjects tolerated more physical deviation in the face while still rating it as much like Bill as the original face), consistent with a greater coding range for symmetric eye-height deviations than asymmetric eye-height deviations.

Adaptation Aftereffects (Upright Faces)

The primary theoretical questions were then about adaptation effects. We examined both the form of postadaptation change in the rating curves and the size of the aftereffect for symmetric versus asymmetric distortions. Analysis of adaptation effects used only preadaptation data collected in the same session as the postadaptation data.

Form of adaptation. Figure 6 shows pre- and postadaptation rating curves, averaged across subjects, for each adapter direction separately (positive and negative). Patterns of results did not differ as a function of identity status (adapted individuals = Bill and Sam; unadapted individuals = John and Fred), and so results are presented collapsed across all four test individuals.

In general, results are consistent with the predictions of a two-pool model, in that they show that adaptation produced an overall shift in the curve. Note also that the direction of the overall shift along the x -axis was always the predicted direction, namely toward the adapter. To make this prediction clear, consider that after adapting to an eyes-up face (+20), the unaltered face (0) should appear to have its eyes moved down; thus, the face perceived as most like the original will be a stimulus with its eyes physically higher than the original position. This corresponds to the center of the subthreshold region shifting in the positive direction, for a positive adapter.

In the case of the symmetric distortion (top two panels of Figure 6), results exactly matched the predictions of the two-pool model. On the same side of 0 as the adapter, perception was affected in the same way regardless of whether the stimulus was more extreme than the adapter (e.g., +25 for a +20 adapter; Regions A in Figure 6) or less extreme than the adapter (e.g., +15; Regions B in Figure 6); in both cases, ratings shifted downward, indicating the test face appeared more like the target individual. In addition, on the opposite side of 0 (e.g., negative test stimuli for a positive adapter; Regions C in Figure 6), ratings shifted upward, indicating the test faces appeared less like the target individual.

Results for the asymmetric distortion (bottom two panels of Figure 6) are less clear cut but are more consistent with a two-pool model than a multichannel model. There is a general pattern of the curve staying the same basic shape and shifting toward the adapter, although most of the change appears to happen on the adapter side

⁴ This program allows the threshold points to be determined from the data, unlike the majority of programs, which require these to be prespecified.

Figure 4. Examples of fits with different functions. Data fitted were single-subject responses in a single condition. A and B: Fechner law fits for two data sets (S1 = Subject 1; RR = Rachel Robbins). C: The same data as Panel B to illustrate that a Pearson IV curve does not improve the fit. D: The same data again with a seventh-order polynomial fit, illustrating that the mere inclusion of seven parameters (as in the Fechner curve) does not produce an appropriate fit if the curve shape is wrong.

of the original norm; however, this might be attributable simply to the fact that the overall aftereffect is very small. There was no evidence for the key prediction of a multichannel model, namely, that stimuli more extreme than the adapter (i.e., Regions A) should have ratings shifted upward after adaptation. Averaging across all eight deviation values beyond the adapter (+25, +30, +40, and +50 for the positive adapter; -25, -30, -40, and -50 for the negative adapter), there was essentially no change in rating scores between pre- and postadaptation (a tiny increase of 0.04 rating units on a scale of 1–9). This was also true considering only the values closest to the adapter (+25 and +30 for positive; -25 and -30 for negative) to avoid possible ceiling effects at the most extreme positions (a tiny decrease of 0.01 rating units).

Size of the aftereffect for symmetric versus asymmetric distortions. Figure 6 also indicates that shift in the curve after adaptation was greater for symmetric distortions of eye height (top two panels) than it was for asymmetric distortions (bottom two panels), consistent with the predictions of a two-pool model (given the ± 20 pixel adapter). To quantify this effect and allow a concise presentation of results from individual subjects, we calculated a simple

measure of shift, as the difference between the center of the subthreshold region (i.e., the test stimulus perceived as most like the target individual) before and after adaptation. Shift scores were defined such that a change in the direction predicted for adaptation always corresponded to a positive shift value (i.e., for positive adapters, $\text{shift} = c_{\text{post}} - c_{\text{pre}}$; for negative adapters, $\text{shift} = c_{\text{pre}} - c_{\text{post}}$).

Shift scores did not differ noticeably as a function of the direction of the adapter (for symmetric, mean shift positive adapter = 6.39 pixels, negative adapter = 5.35; for asymmetric, positive adapter = 1.03 pixels, negative adapter = 1.25), and so results were collapsed over this variable to maximize stability of the data. Figure 7 then shows shift scores for the five individual subjects, for both adapted (Bill and Sam) and unadapted (John and Fred) test faces. It is apparent that the greater adaptation effect for symmetric than asymmetric distortions was consistent across subjects. Indeed, the difference between these conditions (collapsed over identity and direction of adapter) was significant even with only 5 participants, $t(4) = 3.19$, $p < .05$.

Generalization over identity. Although identity status (adapted individuals = Bill and Sam; unadapted individuals =

Figure 5. Experiment 1: Preadaptation identity ratings for upright faces averaged over subjects (as well as identity of face and direction of future adapter) and the corresponding Fechner law fit, showing a wider subthreshold region for symmetric than for asymmetric distortions. The inset box indicates the width of the subthreshold region ($|T_l| + |T_r|$, in pixels per eye) for each individual subject (S1–S4 = Subjects 1–4; RR = Rachel Robbins). T_l = left threshold value; T_r = right threshold value.

John and Fred) did not interact noticeably with the main factors of interest, another question was whether there was an overall effect of identity status, that is, larger aftereffects overall for the adapted individuals than for the unadapted individuals. This did not occur. Indeed, after adaptation to Bill and Sam, the mean shift score averaged across subjects for John and Fred (symmetric = 5.37, asymmetric = 1.32) was, if anything, slightly larger than for Bill and Sam themselves (symmetric = 4.37, asymmetric = 0.75). Thus, transfer across identity was strong, for both symmetric and asymmetric distortions.⁵ This result is consistent with previous findings; two studies (Webster & MacLin, 1999; Yamashita et al., 2005) have previously found close to complete transfer across identity change where the face stimuli were similar in overall appearance, as were our stimuli (all had the same hair). Variation between subjects in the amount of transfer (see Figure 7) is also consistent with previous findings (Yamashita et al., 2005).

Inverted Faces and Orientation-Transfer Conditions

We now briefly describe results for inverted faces (I-I) and the orientation transfer conditions (I-U and U-I). Fechner's law fits indicated a quantitative difference between upright and inverted faces, with sensitivity to distortion weaker in inverted faces. The subthreshold region was wider inverted than upright for both symmetric distortions (upright = 11.62 pixels, inverted = 14.83) and asymmetric distortions (upright = 8.49 pixels, inverted = 9.89). This is consistent with extensive evidence suggesting that changes to relational information (spacing between features) are perceived less strongly in inverted faces than in upright faces (e.g., Bartlett & Searcy, 1993; McKone et al., 2005).

With respect to adaptation aftereffects, the form of the aftereffect (overall shift in Fechner curve) was the same as for the U-U conditions. Figure 8 plots shift scores for I-I and the two orientation mismatch conditions, collapsed over identity and adapter direction. For the I-I condition, aftereffects were similar in strength

to those for U-U (Subjects 1–4 in Figure 7) and, like U-U, larger for symmetric than asymmetric distortions. For the U-I and I-U conditions, aftereffects were much smaller and, indeed, were barely different from zero. The pattern of strong aftereffects for I-I but weak transfer across orientations is consistent with previous findings using different distortion types (Watson & Clifford, 2003; Webster & MacLin, 1999).

Discussion

The results support the predictions of a two-pool model over multichannel models. Subjects were asked to rate a range of deviations from normal (to a maximum of ± 50 pixels) according to how much each looked like the original undistorted face, allowing the curve relating physical to perceived deviation to be determined. After adaptation to an intermediate (± 20 pixels) adapter, adaptation took the form of a general shift in the curve about the original norm (most clearly so in the case of symmetric distortions) and not perceptual repulsion away from the adapter position (for either distortion type). The form of adaptation was thus consistent with a model in which two broadly tuned pools have maximal response to opposite extremes of eye position (e.g., far up and far down). Note that previous studies that have measured perception at multiple test positions have also found an overall shift of the curve between pre- and postadaptation for other

⁵ The generalization of adaptation over identity should not be taken to suggest that subjects did not reference identity in this task. McKone et al. (2005) used the same identity rating instructions as here and contrasted these to normality ratings. In the unadapted state (adaptation was not tested), both tasks produced a Fechner law relationship between physical and perceived deviation. However, sensitivity to deviations was noticeably stronger in the identity task (i.e., smaller subthreshold region and steeper slope at threshold). The present fit parameters agree well with those of McKone et al.'s identity task.

Figure 6. Experiment 1 upright-adaptor/upright-test condition: Mean identity ratings averaged across subjects, and Fechner's law fits, for preadaptation (open circles, darker fit line) and postadaptation (filled diamonds, lighter fit line). Region A = deviation levels more extreme than the adaptor; Region B = deviation levels between the adaptor and zero (the original norm); Region C = deviation levels on the opposite side of zero from the adaptor. A two-pool model predicts that following adaptation, ratings should decrease in Regions A and B and increase in Region C; a multichannel model predicts that ratings should increase in Region A and decrease for at least some of Region B (with exact prediction depending on the bandwidth of the channels).

distortion types (e.g., identity distortion, Leopold et al., 2001; global expansion/contraction, Rhodes et al., 2003), suggesting opponent coding for a range of different facial aspects. The present study strengthens the conclusion in favor of two-pool models by having, for the first time, tested perception of stimuli in the region beyond the adaptor position.

The study also showed that with the intermediate adaptor tested, there were differences in the amount of adaptation for the two distortion types. The size of the aftereffect was larger for symmetric distortions than for asymmetric ones. This is also consistent with the prediction of a two-pool model, presuming that the ± 20 adaptor is inside the neural coding range for symmetric deviations of eye height but outside the coding range for asymmetric deviations.

Of interest, these conclusions apply to inverted faces (I-I condition) as well as upright faces (U-U). Despite using relational distortion types in the present study, which have been argued to be particularly sensitive to inversion (e.g., Leder & Bruce, 1998; but see Yovel & Kanwisher, 2004), we found the amount and patterns of adaptation to be essentially identical for U-U and I-I. Note, however, that the weak transfer across orientations (I-U and U-I) indicates that the U-U and

I-I effects derive from largely independent sets of neurons, even though these neurons are behaving in similar ways.

Experiment 2: Varying Adaptor Position

In Experiment 2 three adaptor positions were tested (± 5 , ± 20 , and ± 50 pixels per eye), with the size of the aftereffect measured as the amount of shift in the face perceived as most normal, determined using an adjust-to-most-like technique. We argued in the introduction that the three adaptor positions should lie in different positions relative to the coding range for the different distortion types. For symmetric distortions of eye height, all three adaptors should fall within the coding range. Considering both the differences arising from individual face structure and the even larger differences arising from three-dimensional up-down head rotation, we suggested the required range extends from the eyes being very close to the tip of the nose (looking down) to the eyes being very close to the hairline (looking up). In our stimuli, these eye positions correspond to a range of approximately -60 to $+60$ pixels per eye. For asymmetric distortions, in contrast, only the ± 5 adaptor would fall within or close to the coding range, with the

Figure 7. Experiment 1 upright-adaptor/upright-test condition: Shift scores (number of pixels shift in the center of the subthreshold region between pre- and postadaptation) for individual subjects (S1–S4 = Subjects 1–4; RR = Rachel Robbins). Data are collapsed over direction of adaptor (positive or negative). Each plot also shows the mean (± 1 SEM, denoted by the error bars) of these individual subjects' scores (denoted "Mean of Ss" on the figure), as well as a score derived by averaging data across subjects first and then fitting (denoted "Fit to mean data").

± 20 and ± 50 adaptors falling outside it. There is no three-dimensional head rotation that produces asymmetric changes in eye height, and even a ± 5 pixel-per-eye deviation would be unusual (although not impossible) to observe in the natural population of faces. Thus, our guess is that the coding range for asymmetric distortions might be at most -10 to $+10$ pixels per eye.

In addition to these a priori arguments about neural coding range, also note that there is direct evidence that the three adaptor positions lie in different positions relative to subjects' perception of the two types of face distortion. Symmetric and asymmetric deviations of eye height produce the same qualitative form of effect on perception, but the quantitative parameters of the Fechner law curves are quite different for the two cases. Figure 5 (from Experiment 1) allows readers to determine the location of the three adaptor positions relative to the Fechner curves for the two distortion types in the unadapted state. An obvious difference is that the ± 5 pixel-per-eye adaptor is within the subthreshold region (i.e., the area for which faces are seen as equally normal or most like the original) for the symmetric distortion but is just outside it for the asymmetric distortion. In addition, the ± 20 adaptor is somewhat outside the subthreshold region and in the steep part of the log region for the symmetric distortion but is well outside the threshold region and approaching the flatter part of the log region for the asymmetric distortion.

The task in Experiment 2 was an adjust-to-most-like procedure, in which subjects were initially given a clearly nonnormal distortion and allowed to adjust the deviation level until the stimulus appeared "most like the target individual." Shift in the face stimulus perceived as most normal after adaptation, compared with before adaptation, was the measure. Note that in this procedure, only one test position was assessed (the face seen as most normal), rather than 27 deviation levels as in Experiment 1.⁶

Given this test position of "zero" (face perceived as undistorted) and adaptor positions of ± 5 , ± 20 , and ± 50 , predictions of the two-pool versus multichannel models can be derived. The two-pool model predicts that the amount of shift in the norm (the size

of the aftereffect at zero) will increase as the adaptor becomes more extreme in deviation level (i.e., is placed further from the original norm) but remains within the coding range of the two broadly tuned pools. Thus, presuming our arguments about the coding ranges are correct, the model would predict that for symmetric deviations, the aftereffect should be largest for ± 50 adaptors and smallest for ± 5 adaptors. For asymmetric distortions, in contrast, the aftereffect should be largest for ± 5 adaptors and either smaller or zero for ± 20 and ± 50 adaptors (depending on the assumptions made about the shape of the tuning functions outside the primary coding range; see introduction). A corollary of these predictions is that for ± 20 and ± 50 adaptors, the aftereffect should be larger for symmetric than for asymmetric distortions (replicating Experiment 1).

A multichannel model incorporating narrow tuning of pools makes rather different predictions. These arise from the fact that adaptation is predicted only for test stimuli close to the adaptor. Given that our test position was of the perceived "zero" face, the aftereffect should then be strongest for ± 5 adaptors, and weakest for ± 50 adaptors, for both symmetric and asymmetric distortions.

Overall, the predictions of the two models are quite similar for asymmetric distortions but opposite for symmetric distortions. For asymmetric, both models predict that increasing the deviation level of the adaptor should decrease the size of the aftereffect; for symmetric, however, the two-pool model predicts an increasing aftereffect whereas the multichannel model predicts a decreasing aftereffect.

In Experiment 2, only upright faces were tested. To assess all three adaptor positions for both distortion types required 2 hr per subject. Data were averaged over multiple subjects, and it was not feasible to analyze individual subject data because the estimate of

⁶ It was thus not possible to examine shifts to both sides of the adaptor, to check for overall shift versus repulsion away from the adaptor.

Figure 8. Experiment 1: Shift scores for conditions involving inverted faces. The aftereffect was strong for inverted faces following inverted adapters (I-I) but transferred poorly across changes in orientation (U-I and I-U). Data are collapsed over identity and direction of adapter. Mean = average of individual subjects ± 1 SEM (denoted by the error bars). Note that the scale is the same as for Figure 7 showing upright adapters (U-U).

the shift in an individual subject's center point (the face perceived as most like Bill/Sam/John/Fred) was less reliable with the adjust-to-normal procedure than with the rating-27-levels procedure used in Experiment 1.

Method

Subjects

Twenty-four naive Caucasian subjects participated (age range = 18–39 years; 9 men, 15 women). Subjects were paid \$25 for their participation. All reported normal or corrected-to-normal vision.

Design

All stimuli (training, preadaptation, adapters, postadaptation) were upright. Each subject was tested in three 40-min sessions. Each session began with identity training on the original undistorted faces. Adaptation effects were then assessed for one adapter position (e.g., -5 pixels per eye) for one distortion type (e.g., symmetric) and, after a 15-min filled delay, a different adapter position (e.g., -50 pixels per eye) for the other distortion type (i.e., asymmetric). There was at least 24 hr between sessions with tests of the same distortion type; as in Experiment 1, we checked that no carryover of adaptation between sessions had occurred. Within a

session, during the filled delay, subjects saw 240 trials of normal faces and other objects (at least 120 faces) and were required to press the space bar every time an item appeared for a second or successive time. The filled delay also included one cycle of identity training.

Across all sessions, a given subject received adapters of only one direction (all positive or all negative) but was tested on all three adapter levels (e.g., -5, -20, -50) for both distortion types. Order of the resulting six conditions was counterbalanced across subjects, within the constraints that each session contained one condition of each distortion type, and with different adapter levels for each. Direction of adapter (positive or negative) was manipulated between subjects, with the intention of collapsing over this variable, given that it had had no influence on adaptation effects in Experiment 1.

The final design was thus a $2 \times 3 \times 2$ within-subject design, crossing distortion type (symmetric, asymmetric) with adapter level (5, 20, 50 pixels per eye) and identity status (adapted individuals = Bill and Sam; unadapted individuals = John and Fred, as for Experiment 1). The subject's task was to adjust each face until it appeared "most like the original [Bill/Sam/John/Fred]"; the dependent measure was the physical deviation level of the stimulus chosen.

Stimuli

For the pre- and postadaptation phases, stimuli were as for Experiment 1, except that additional deviation levels were created. To allow adjustments in 1 pixel-per-eye increments, we prepared stimuli for every 1 pixel deviation between -40 and +40, for each identity and distortion type. These were prepared at the larger "test face" size used in Experiment 1. For the adaptation phase, additional stimuli were the ± 5 and ± 50 versions of Bill and Sam, prepared at the smaller "adapter face" size.

Procedure

The procedure was as for Experiment 1, except where noted.

Identity training phase. At the beginning of each session, four cycles of identity training were provided. Unlike in Experiment 1, training included only upright faces, shortening the training phase to approximately 10 min. Each cycle first presented 8 trials of faces with names (2 each of Bill, John, Sam, and Fred), followed by 16 trials of training with feedback without names (4 of each individual). One cycle through this identity training preceded testing of the second distortion type in the session.

At the beginning of Session 1, subjects were required to achieve 80% correct in the last of the four training cycles. Three subjects who initially failed to reach this criterion were given an extra two cycles before continuing. All subjects remained above 80% correct in the training phase of Sessions 2 and 3.

Preadaptation phase. The preadaptation phase for a given condition (e.g., symmetric, -5 adapter) presented two adjust-to-most-like trials for each of the four individuals, intermixed and in random order. One of these trials began with an obvious positive distortion (e.g., eyes too high), and one with an obvious negative distortion (e.g., eyes too low). The starting position was chosen randomly for each trial within the range of +20 to +40 pixels per eye, or -20 to -40 pixels per eye, respectively. To adjust the stimuli, subjects used number keys at the top of the keyboard.

Keys 1 and 4 allowed coarse adjustment, moving the eyes by ± 5 pixels with each press (one key for each direction). Keys 2 and 3 allowed finer adjustments, moving the eyes by ± 1 pixel. A beep sounded if subjects tried to move the eyes outside the maximum range of ± 40 pixels. Subjects were instructed to make adjustments fairly quickly, taking approximately 10 key presses and at most 1 s per press. When the subject considered the stimulus to "look like the undistorted person," he or she pressed the space bar to accept this stimulus, and its deviation level was recorded.

Adaptation phase. This was exactly as for Experiment 1.

Postadaptation phase. This was as for the preadaptation phase with the addition of 5-s top-up adapters, shown in the same format as in Experiment 1.

Results

Trials on which subjects made no adjustments before pressing the space bar were treated as errors and discarded. This never occurred for more than one trial for any given subject.

Confirming Greater Perceived Variability for Symmetric Than for Asymmetric Distortions

As for the previous experiment, difference in perceived variability between the symmetric and asymmetric distortions was assessed in the unadapted state. For each subject, the preadaptation deviation level (measured in pixels per eye) of the stimulus chosen as "most like Bill/Sam/John/Fred" was computed for each distortion type. To obtain the most reliable data, we averaged over all four individuals and all three future adapter levels (e.g., -5, -20, and -50).

The variability of the symmetric and asymmetric conditions across subjects was then compared using Levene's test of equality of variances. This showed that the variance of 5.38 for the symmetric distortion (around a mean of -1.49 pixels) was significantly higher than the variance of 0.19 for the asymmetric distortion (around a mean of 0.27 pixels), $F(1, 46) = 30.26$, $p < .001$. That is, subjects were less certain about where the eyes should be placed in the symmetric condition, corresponding to the finding that there was a broader subthreshold region for ratings of the symmetric distortion than the asymmetric distortion in Experiment 1. Again, the implication is that subjects' perception tracked the fact that the symmetric distortion corresponds to more variability in the real world than does the asymmetric distortion.

Adaptation Aftereffects

As in Experiment 1, shift scores were defined such that a positive shift always corresponded to a change in the direction predicted for adaptation. After adaptation to positive adapters, the stimulus chosen as "most like" should be more positive than before adaptation (thus, $\text{shift} = c_{\text{post}} - c_{\text{pre}}$), and after adaptation to negative adapters, the chosen stimulus should become more negative (thus, $\text{shift} = c_{\text{pre}} - c_{\text{post}}$).

An initial four-way mixed analysis of variance (ANOVA) was conducted on the shift scores. This included distortion type (symmetric vs. asymmetric), adapted-unadapted identity status (Bill and Sam vs. John and Fred), adapter level (5, 20, 50), and direction of adapter (positive or negative). This showed no main effect of

direction or any interactions involving this variable: for Direction $\times$ Distortion Type, $F(1, 22) = 3.27, p > .08$ (all other F s < 1). Thus, we collapsed over direction in later analyses.

Generalization over identity. The four-way ANOVA revealed no main effect of identity status, $F(1, 22) = 2.88, p > .10$. However there was a significant interaction between identity status and distortion type, $F(1, 22) = 5.97, p < .05$. For the asymmetric distortion, results agreed with those of Experiment 1. Adaptation generalized perfectly over identity: If anything, unadapted individuals showed a slightly larger effect (John and Fred: $M = 0.78, SEM = 0.12$) than adapted individuals (Bill and Sam: $M = 0.61, SEM = 0.15$). For the symmetric condition, results differed slightly from those of the first experiment: Adaptation transferred well, but not perfectly, over identity changes. The effect for unadapted individuals (John and Fred: $M = 2.29, SEM = 0.48$) was significantly smaller than for adapted individuals (Bill and Sam: $M = 3.24, SEM = 0.68$), $t(23) = 2.22, p < .05$. These results are averaged across level of adapter, because there was no three-way interaction between identity status, distortion type, and adapter level ($F < 1$). A check of the three adapter levels separately confirmed that the same trend was present for symmetric distortions at all adapter levels (for ± 5 , John and Fred = 1.25, Bill and Sam = 1.75; for ± 20 , John and Fred = 2.97, Bill and Sam = 3.66; for ± 50 , John and Fred = 2.65, Bill and Sam = 4.32).

Effects of distortion type and adapter level. Given the lack of any three- or four-way interactions in the initial ANOVA, the questions of primary interest—how the adaptation aftereffect was influenced by distortion type and adapter position—were examined collapsing over identity and direction. Results are shown in Figure 9. A two-way ANOVA revealed a main effect of distortion type, $F(1, 23) = 12.64, p < .01$, and a significant interaction between distortion type and adapter position, Wilks's lambda = .70,⁷ $F(1, 23) = 4.67, p < .05$. As can be seen in Figure 9, the interaction takes the form that whereas shift scores increased with increasing adapter level for the symmetric distortion (the trend was $5 < 20 < 50$), the opposite was true for the asymmetric distortion (the trend was $5 > 20 > 50$).

To follow up the interaction statistically, we analyzed results for the two distortion types separately. Results of the analysis were neatest for the asymmetric distortion, so this is presented first. In

this condition, trend analysis revealed a significant linear trend, $F(1, 23) = 5.82, p < .05$, with no quadratic component ($F < 1$). Contrasts confirmed that the ± 5 pixel adapter produced a significantly larger aftereffect than the ± 50 pixel adapter, $F(1, 23) = 6.65, p < .05$. The ± 20 pixel condition fell in the middle and was not significantly different from either the ± 5 condition, $F(1, 23) = 1.21, p > .05$, or the ± 50 condition, $F(1, 23) = 2.18, p > .05$.

For the symmetric distortion, there was a trend in the opposite direction (see Figure 9). The sphericity assumption of repeated measures ANOVA was violated, and so trend analysis was not possible. The multivariate equivalent of a one-way ANOVA only approached significance, Wilks's lambda = .80, $F(1, 22) = 2.73, p = .087$. However, this test does not take account of the ordering of the conditions and thus has substantially less power than trend analysis. A priori t tests showed that the ± 5 adapter produced a significantly smaller aftereffect than the ± 50 adapter, $t(23) = 2.36, p < .05$. The ± 20 adapter again fell in the middle and did not differ significantly from either the ± 5 adapter, $t(23) = 1.31, p > .05$, or the ± 50 adapter ($t < 1$). Overall, these results confirm that for symmetric distortions, adapters with more extreme deviations produced more adaptation, and for asymmetric distortions, they produced less adaptation.

We also compared amounts of adaptation for symmetric and asymmetric distortions at each adapter level. For the ± 50 adapters (within our proposed coding range for symmetric but very far outside it for asymmetric), the greater adaptation for symmetric than asymmetric distortions was significant, $t(23) = 4.90, p < .001$. The difference in the same direction for ± 20 adapters was not quite significant, $t(23) = 1.97, p = .06$, although it can be seen from Figure 11 that this was probably attributable to the abnormally large variability in the symmetric ± 20 condition.⁸ For ± 5 adapters (within our proposed coding range for both distortions), a quite different result was revealed. Here, symmetric showed only marginally more adaptation than asymmetric, and this difference was far from significant ($t < 1$). Overall, these results replicate the symmetric-asymmetric difference found in Experiment 1 with moderate and extreme adapters and also show that the effect is significantly modified by adapter position in the manner predicted by a two-pool model.

Finally, adaptation was compared with zero in each condition in Figure 9. For symmetric distortions, one-sample t tests showed significant adaptation for all three adapter levels (all p s $< .02$). For asymmetric distortions, the generally smaller effects were still all

Figure 9. Experiment 2: Interaction between distortion type (symmetric vs. asymmetric) and adapter position (± 5 , ± 20 , ± 50). Data are collapsed over identity and direction of adapter. Error bars show $\pm 1 SEM$.

⁷ Where the sphericity assumption of repeated measures ANOVA was violated, the multivariate approach was used.

⁸ The larger variability for ± 20 adapters may be a real finding. There are likely to be individual differences in the Fechner curves underlying perception across subjects, particularly in the width of the subthreshold region and the steepness of the log section immediately beyond this region. (These would correspond to slight differences in the shape and/or steepness of neural response curves presented in the model in Figure 1.) Because the ± 20 pixel adapter falls in the steep part of the log region (averaged across subjects) and is not far from the threshold, small individual differences in Fechner curve parameters could produce large differences in amount of adaptation. In contrast, the ± 5 and ± 50 adapters lie in flatter regions of the curve, where individual differences in exact curve parameters will not make as much difference to the response.

significant ($ps < .01$), including, somewhat surprisingly, the tiny (0.34 pixels) effect at ± 50 .

Discussion

The main finding of Experiment 2 was the very different influence of adapter position on the aftereffect for the two distortion types. For symmetric distortions, the more extreme the adapter was, the larger was the aftereffect. For asymmetric distortions, the more extreme the adapter was, the smaller was the aftereffect (although it remained significant even for ± 20 and ± 50 adapters). This also resulted in larger aftereffects for symmetric than for asymmetric distortions with more extreme adapters (± 20 and ± 50), replicating the results of Experiment 1.

The general pattern of adapter position effects was as predicted by the two-pool model. This model predicts increasing aftereffects as the deviation of the adapter becomes more extreme within the coding range. Assuming coding ranges of approximately -60 to $+60$ pixels for symmetric eye-height distortions but only -10 to $+10$ for asymmetric distortions, the two-pool model can then explain increased aftereffects across the ± 5 , ± 20 , and ± 50 adapter range for the symmetric distortions but decreased aftereffects beyond the ± 5 adapter for the asymmetric distortion. It is true that a multichannel model could also explain the pattern for the asymmetric condition (only one test position located around zero was used, and the aftereffect became weaker as the adapter was placed further from this position). However, a multichannel model could not explain the pattern observed for the symmetric condition; also recall that even for the asymmetric condition, this model could not explain the lack of repulsion away from the adapter position in Experiment 1.

The second important aspect of the results is that although adaptation was weaker outside the coding range for asymmetric distortions, it did not disappear completely; that is, the ± 20 asymmetric condition and the ± 50 asymmetric condition produced aftereffects that, while small, were significantly greater than zero. This result poses somewhat of a challenge for models to explain, as we consider further in the General Discussion.

Turning to more minor findings, with respect to identity effects, Experiment 2 again showed good generalization of adaptation over adapted-unadapted identity status. However, generalization results were not exactly as for Experiment 1. Previously, there was perfect generalization over identity for both symmetric and asymmetric distortions. In the present experiment, there was again perfect generalization for asymmetric distortions but less than complete generalization for symmetric distortions. One possible reason for the difference between the two experiments is that some individuals show greater identity generalization than others (see Figure 7), and it can be noted that Experiment 1 tested only 5 subjects whereas Experiment 2 tested 24. This suggestion is consistent with previous findings. Yamashita et al. (2005) found moderate to strong transfer across identity but also quite a lot of variability in the amount of transfer for the 4 individual subjects tested (Webster & MacLin, 1999, also found moderate to strong transfer but only presented averaged results for their 4 subjects).

A final point is that the overall effect of adaptation in Experiment 2 was somewhat smaller than in Experiment 1. Comparing the data for the adapter level that appeared in both experiments (± 20), the mean shift in Experiment 2 was 3.07 pixels for sym-

metric and 0.74 pixels for asymmetric, and the mean shift in Experiment 1 was 4.98 pixels for symmetric and 1.44 pixels for asymmetric. We presume that the smaller effects in the adjust-to-most-like task than in the rating task can be attributed to the time required to make an adjustment (on average, subjects made approximately 8–11 adjustments taking 7–16 s per trial), reflecting some decay of adaptation in the period immediately after removal of each top-up adapter (see, e.g., Leopold et al., 2001) or during the lengthened presentation time of the test stimulus (Leopold et al., 2005).

General Discussion

The major findings of the present study were as follows. In Experiment 1, plotting full rating curves showed that the adaptation aftereffect took the form of a general shift in perception relative to the original norm, rather than perceptual repulsion away from the position of the intermediate value (± 20 pixel-per-eye) adapter. The size of the aftereffect (shift in face stimulus perceived as most normal) was also larger for the symmetric distortion of eye height than for the asymmetric distortion, again with the intermediate value adapter. In Experiment 2, the latter finding was replicated for the intermediate (± 20 pixel) and extreme (± 50 pixel) adapters, using a much larger number of subjects and upright faces only, showing that the result in Experiment 1 could not be attributed to exposure to inverted as well as upright faces. In addition, the size of the aftereffect did not differ for symmetric and asymmetric distortions with small-value (± 5 pixel) adapters, corresponding to a pattern in which adaptation for symmetric distortions increased as the distance of the adapter from the norm became more extreme (± 5 vs. ± 20 vs. ± 50 pixel-per-eye deviation), whereas adaptation for asymmetric distortions decreased. Finally, the aftereffects for asymmetric distortions with the more extreme adapters (± 20 and ± 50 pixels), though small, were still significantly greater than zero.

Other findings were as follows. There was no effect of adapter direction; for example, aftereffects arising from eyes-up adapters were of the same size as aftereffects arising from eyes-down adapters. Adaptation generalized strongly over identity changes, with adaptation for same-identity conditions as large or nearly as large as that for different-identity conditions (also see Webster & MacLin, 1999; Yamashita et al., 2005). With respect to orientation effects, assessed in Experiment 1, the aftereffect for inverted faces (I-I) took the same form (general shift in rating curve), was a similar size, and was affected by distortion type in the same manner as for upright faces (U-U). Transfer of aftereffects across orientation changes (U-I and I-U) was minimal. Inversion results were thus similar to previous findings with other distortion types (Leopold et al., 2001; Watson & Clifford, 2003; Webster & MacLin, 1999).

Our results have a number of important theoretical implications, considered in detail below. Briefly, the results are consistent with norm-based coding in face space. They argue that face-space trajectories associated with coding a greater range of a facial aspect (e.g., symmetric deviations of eye height) show very different patterns of aftereffects than face-space trajectories associated with coding a smaller range (e.g., asymmetric deviations of eye height). They support two-pool rather than multichannel models of adaptation aftereffects. Moreover, they indicate that norm-

based coding and two-pool models apply to inverted as well as upright faces; that is, neurons supporting adaptation for inverted faces behave in a very similar way to those supporting adaptation for upright faces, even though the poor transfer across orientations shows that the particular neurons involved must be largely different in each case (also see Rhodes et al., 2004).

Norm-Based Coding in Face Space

The concept of face space is widely considered to be a useful heuristic for understanding the representation of individual identity (e.g., Valentine, 1991; Valentine & Bruce, 1986). A long-standing question (e.g., Valentine, 1991) has been whether face space uses norm-based coding, in which individuals are coded in terms of their direction and level of deviation from the average, or exemplar-based coding, in which the center of face space has no particular special status other than that more exemplars tend to fall close to it than further from the center. In studies of adaptation aftereffects, it has been standard to presume that face space must use a norm-based coding, in order to explain the opponent behavior of the adaptation aftereffect (e.g., Leopold et al., 2001; Rhodes et al., 2003). This is illustrated in Figure 10, in which adaptation shifts the norm toward the adapter, along a trajectory describing the adapted facial aspects.

The present results provide further support for norm-based coding. Exemplar-based coding would be similar to a multichannel model in operation. As noted by Loffler, Yourganov, Wilkinson, and Wilson (2005; also see Freiwald, Tsao, Tootell, & Livingstone, 2005), an exemplar-based model would correspond to different neural populations responding to small subsets of individuals along a given trajectory in face space (i.e., to individuals lying along the same trajectory in face space but at different distances from the average); under these circumstances the effect of an adapter should always be to repulse perception away from the adapter. In contrast, norm-based coding would be similar to a two-pool model, in that all individuals falling along a given trajectory would activate one common neural population, allowing adaptation to shift perception with respect to a central neutral point, as we observed in the present study.

Figure 10. Adaptation conceptualized as a simple shift in the position of the norm of a trajectory in face space, illustrated using the height of both eyes together. Black dots indicate distribution of individual faces along this trajectory. Some scatter is also included in the vertical direction to indicate that exemplars simultaneously vary along other trajectories (e.g., eye-height asymmetry).

Patterns of Adaptation Vary for Different Facial Attributes

Perhaps the most important empirical finding of the present article concerns the different patterns of adaptation effects for symmetric and asymmetric deviations of eye height. For symmetric deviations, the aftereffect became larger as the adapter was placed further from the original norm; for asymmetric deviations, it became smaller, leaving a larger aftereffect for symmetric than for asymmetric deviations at moderate-distortion and extreme-distortion adapters. This is the first report of a difference of this nature. Previous studies have not attempted to examine aftereffects to different facial attributes varied independently. Instead, manipulations such as global expansion and contraction (e.g., Rhodes et al., 2003; Watson & Clifford, 2003; Webster & MacLin, 1999) or identity adaptation (e.g., Leopold et al., 2001) have varied multiple aspects of the face at once. Our results indicate that varying facial attributes independently provides a promising avenue for future research.

We have argued that the different effects for symmetric and asymmetric deviations can be understood in terms of amount of variability of different facial attributes. The two distortion types were chosen because they correspond to different amounts of variability in natural face images. The height of both eyes together (symmetric deviation) varies substantially across different individuals and different up-down head rotations, whereas the height difference between eyes (asymmetric deviation) varies to a lesser extent across individuals and is not increased by three-dimensional head rotations. Thus, we suggest that facial attributes for which a large range of values needs to be coded produce adaptation effects that strengthen as the adapter is placed further from the norm, whereas facial attributes for which a small range needs to be coded produce adaptation effects that weaken as the adapter is placed further from the norm.

The symmetric-asymmetric difference also has other implications. First, if our interpretation in terms of coding range is correct, it implies that the coding of eye position must allow for the full range of up-down head rotations. If it were necessary to code only individual person variation in full-front images, then the amount of variation required to cover even symmetric deviations would be fairly small (e.g., perhaps ± 15 pixels or so on our scale). It is only by allowing for up-down head rotations that the coded amount of variation could include our ± 50 -pixel adapters (eyes approaching hairline or tip of nose), which our results indicate it does.

Second, the results indicate that eye height is coded in some way that takes into account the position of the other eye; it is not coded in terms of individual eye position, nor is it coded with respect to only one half of the face (i.e., left half or right half). This conclusion can be drawn because the symmetric and asymmetric distortions matched total amount of metric deviation from the original. Considering the right eye, for example, this means that a symmetric distortion of +20 pixels per eye shifted the right eye up by 20 pixels and an asymmetric distortion of +20 pixels per eye also shifted the right eye up by 20 pixels. Thus, if the position of features in the other side of the face was irrelevant, then the amount of adaptation for symmetric distortions (right eye up, left eye up) should have been the same as the amount of adaptation for asymmetric distortions (right eye up, left eye down).

A final question is whether symmetry might be the key factor producing the symmetric-asymmetric eye-height difference re-

ported here, rather than range of variability. We cannot rule out a role for symmetry per se in the present study, given that range of variability and symmetry were completely confounded. However, we have two reasons to believe that range of variability is the important factor. First, it provides a rational explanation of the effects of adapter position for each distortion type, whereas the idea that "symmetry" could explain the results does not make any obvious predictions about adapter position. Second, in a recent study in our lab (McKone & Edwards, 2005) we examined adaptation aftereffects for mouth shifts. With mouth shifts (unlike eye shifts), it is possible to make an asymmetric change that is still within the range of expected coding values allowing for three-dimensional head rotation: In particular, the mouth can be shifted sideways, destroying the midline symmetry of the face but keeping it in the normal range of positions relative to cheek outline, allowing for rotations around the vertical axis (to three quarter and profile views). Results showed that, for adapter positions 23 pixels from the norm, the aftereffect for left-right shifts was the same as for up-down shifts. This supports the idea that coding range determines the size of the aftereffects, rather than symmetry.

Explaining Our Key Results With a Two-Pool Model

We now revisit the two-pool and multichannel neural models described in the introduction and consider how our results might be best explained in terms of such models. We argue that the two-pool model is clearly superior, although there is one finding that it cannot easily explain.

One result to explain is the simple observation of an opposite perceptual aftereffect: That is, after adapting to eyes-up, an undistorted face appears to have its eyes shifted down (and vice versa), and after adapting to left-eye-up/right-eye-down, an undistorted face appears to have the right eye up and the left eye down (and vice versa). A two-pool model explains these types of effects in terms of the change in the balance of the output of the two pools. For our symmetric condition, for example, one pool would be preferentially excited by eyes-up stimuli, and the other by eyes-down stimuli (see Figure 1). In the unadapted state, the undistorted face is the stimulus perceived as most normal (most like the target individual in our task) because the eyes-up pool and the eyes-down pool are firing equally. After adaptation to an eyes-up face, the response of the eyes-up pool will be reduced more than that of the eyes-down pool, shifting the stimulus perceived as most normal (i.e., the crossover point of equal response from both pools) toward the eyes-up direction and making the original undistorted stimulus perceived as having its eyes down. The two-pool model for the asymmetric distortion would be similar, except that one pool codes for right-eye-up/left-eye-down and the other for left-eye-up/right-eye-down.

The next result to explain concerns the finding of an overall shift in rating curves following adaptation, rather than perceptual repulsion away from the adapter position. In Experiment 1, the adapter position was always ± 20 pixels per eye. Three regions are of interest: test faces that are more extreme in distortion level than the adapter (e.g., for a +20 adapter, these include the +21 to +50 stimuli), test faces that fall between the adapter and the undistorted face (0 to +19), and faces that are on the opposite side of the original norm to the adapter (0 to -50). The two-pool model predicts that perception in each of these regions will shift in the same direction. That is, after an eyes-up adapter, all test faces will

appear to have their eyes lower than before. This will mean that eyes-down test stimuli will look less normal than before (rating scores will go up) and eyes-up stimuli on both sides of the adapter will look more normal than before (rating scores will go down), producing a simple rightward shift of the curve (i.e., toward the adapter direction). This is the result that was obtained, most clearly for symmetric distortions.

In contrast, the multichannel model predicts perceptual repulsion away from the adapter. For example, after adapting to a +20 eyes-up stimulus, eyes-up stimuli between the adapter and the original norm (0 to +19) will appear to have their eyes moved lower (more normal), but eyes-up stimuli on the other side of the adapter (+21 to +50) will appear to have their eyes moved higher (less normal). This would produce a postadaptation bowing of the rating curves, which did not occur for either distortion type. In addition, presuming fairly narrow tuning of pools, the multichannel model would predict little or no adaptation on test faces a long way from the adapter. Again, this did not occur: Perception changed following adaptation even for faces on the opposite side of the original norm from the adapter.

The final result to explain concerns the effects of adapter position. The two-pool model predicts that the size of the aftereffect will increase as the adapter becomes more extreme within the range covered by the two pools (presuming that the tuning curves continue to increase over that full range, rather than flattening out). The change in the relative output of the two pools due to adaptation—and thus the amount of shift in the crossover point determining the norm—will be greatest when the adapter is located in a region producing very different initial responses from each pool (e.g., very strong response from the eyes-up pool, very weak from the eyes-down pool).

For symmetric deviations of eye height, our results can then be explained with a two-pool model that makes the assumptions illustrated in the left panel of Figure 11 ("large range two-pool model"). Specifically, we assume (a) a large coding range of -60 to +60 pixels per eye (i.e., eyes from tip of nose to hairline); (b) neural response curves of approximately the shape illustrated, particularly including the fairly slow rate of growth as deviation of the stimulus is increased away from 0; and (c) that the +5, +20, and +50 adapters are all within the primary range covered by the two pools and that location relative to the neural response curves is approximately as marked. For a 5-pixel eyes-up (+5) adapter, the shift in the norm after adaptation will then be relatively small, because the eyes-up and eyes-down pools are initially firing at only moderately different levels, meaning that the reduction in firing rate for the eyes-up pool will be only a small amount more than for the eyes-down pool. For a +20 adapter, the shift in norm will be larger, because the initial difference in firing rate between the two pools is larger. For a +50 adapter, the initial difference is larger again, and the firing of the eyes-down pool would barely be affected at all whereas the eyes-up pool would be strongly affected, leading to the largest possible shift in the norm.⁹ In the case of the symmetric distortion, this "large range two-pool model" can

⁹ For adapters beyond ± 60 , the normal "first-order" configuration of the face (i.e., eyes below hairline and above nose) would be broken. We suspect that adapters outside of this range (e.g., an adapter with the eyes above the top of the head) would not produce an adaptation aftereffect.

Figure 11. Large range two-pool model for symmetric distortions and small range two-pool model for asymmetric distortions. The positions of our three adapters relative to the coding range are indicated (only positive adapters are shown; the negative adapters would be exactly reversed).

then fully explain our results, in which the aftereffect increased with increasing adapter distortion.

For the asymmetric distortion, it is possible to explain most aspects of our results (although not all) in terms of a model that simply covers a smaller range of stimulus values than for symmetric. This "small range two-pool model" is illustrated in the right panel of Figure 11. This model assumes the primary coding range of asymmetric eye-height values is approximately -10 to $+10$ pixels per eye (as justified in the introduction to Experiment 2). The $+5$ pixel adapter will then be within the coding range of the two pools, whereas both the $+20$ and $+50$ adapters will fall outside it, in the approximate locations illustrated relative to the neural response curves. This model correctly predicts the finding of clear adaptation for asymmetric distortions with a ± 5 adapter. It also correctly predicts that the adaptation effect will be smaller for ± 20 and ± 50 adapters than for ± 5 . However, it is able to predict this only because the 20 or 50 adapters activate neither pool (i.e., essentially, the neural response drops immediately to zero outside the primary coding range), and under this circumstance, the model would produce no shift in the norm at all (i.e., no adaptation). This prediction is inconsistent with the finding of small but significant aftereffects for asymmetric distortions with ± 20 and ± 50 adapters.

Overall, the two-pool model is consistent with most of our empirical findings, whereas the multichannel model cannot explain the key results of overall shift and increasing adaptation for more extreme adapters for symmetric distortions. However, the two-pool model does not obviously predict the small but significant aftereffects for asymmetric distortions found for adapters of ± 20 and ± 50 pixels.

How Can We Explain Nonzero Adaptation at ± 20 and ± 50 for the Asymmetric Distortion?

How can we then explain this phenomenon? One possibility is a multichannel model in which the multiple channels are broadly tuned (i.e., respond to a large range of eye-height values) rather than narrowly tuned (respond to a small range of eye-height values). In a model of this nature, the size of the aftereffect will

decrease as the adapter is shifted further from the norm, and if the overlap between channels is broad enough, it may be that an adapter even as far away from the norm as ± 50 pixels could produce a small shift at that test position. However, such a model would be somewhat ad hoc in that it could not explain the increasing effects of adapter position for the symmetric condition, and implausible in that it could not explain the lack of perceptual repulsion away from the adapter for both distortion types.

A second possibility involves a modification of the small range two-pool model. We have, so far, not considered what happens to the tuning profile of the two pools outside the primary coding range (e.g., once the eyes reach outside the head for symmetric or outside the natural range of individuals for asymmetric). Do these simply drop immediately to zero? Or might they turn around at larger stimulus values and return more slowly to zero? We have illustrated the latter idea in the left panel of Figure 12, where we have assumed that maximum firing to asymmetric distortions is reached quite close to the norm and that, beyond this range, the curves fall off sharply at moderate deviation levels but have very extended tails. This model is able to explain the small but nonzero adaptation for ± 50 adapters, because preadaptation output from the two pools is slightly different. Also, if the tuning curves had the shape indicated in the figure, the model can successfully predict that adaptation for ± 20 should be greater than for ± 50 (because the preadaptation difference in pool output is larger) and less than for ± 5 (because the preadaptation difference is smaller). Unfortunately, however, the "turnaround" model also has one very serious drawback. Regardless of whether a difference readout or a ratio readout from the two pools is assumed, the model falsely predicts the existence of metamers (right two panels of Figure 12), namely, different physical stimuli that produce the same percept. (The term *metamer* is derived from color vision, where, for example, the percept of white can be produced by multiple combinations of three wavelengths.) Any two-pool model that allows response curves to decrease, rather than consistently increasing, will predict metamers. If we were to assume that response curves took the specific shape drawn in the top panel of Figure 12, then the

Figure 12. A two-pool model in which response curves turn around beyond the primary coding range, shown with adapter positions as for asymmetric distortions. The problem with this model is that it produces metamers (i.e., two deviation levels that produce the same output values and thus the same percept), as shown for difference and ratio readouts from the two pools.

difference readout (top right panel of Figure 12) indicates that the -45 stimulus and the -5 stimulus should be perceived as identical. These face metamers clearly do not exist: Two very different eye-height deviations do not look the same to observers.

A third possibility is that a small amount of adaptation in the asymmetric condition comes from midlevel vision rather than face-level processing and so is not directly relevant to our models. In particular, if there were a generic "equal height" or "symmetry" mechanism that could be adapted (e.g., adapting to $_ _$ makes $_ _$ appear somewhat like $_ _$), then a small aftereffect for asymmetric eye height outside its coding range in face space could potentially be explained. As far as we are aware, no one has looked for the existence of such a mechanism in midlevel vision. However, some potentially related evidence comes from Wolfe (1987), who demonstrated an opposite aftereffect to vernier offset using vertical gratings that did not depend on strict fixation. If the same effect were to exist for horizontal gratings, the result could perhaps arise from a generic mechanism suitable for explaining our results.

Overall, the origin of the aftereffect for ± 20 and ± 50 adapters in the asymmetric condition is unresolved. If the effect comes from midlevel vision, then it does not require any alterations to the two-pool model successfully used to explain the rest of our results. On the other hand, if it arises from face-level processes, then it will require modifications to our preferred model, but the nature of the modifications needed remains unclear.

Predictions of a Two-Pool Model for the Tuning Functions of Individual Neurons

We have argued that two critical findings in the present study—the overall left–right shift in rating curves rather than perceptual repulsion away from the adapter in Experiment 1, and the larger aftereffect with ± 50 adapters than with ± 5 adapters for symmetric distortions in Experiment 2—support a two-pool model of face distortion aftereffects. Note that what we are in fact suggesting is not one two-pool model but many two-pool arrangements describing all of the various aspects of faces that vary across individuals and face images and that are coded in some way within face space. Thus, for example, there may be two opponent neural pools tuned maximally to eyes up and eyes down (with average eye height as the neutral point of equal response), two opponent pools tuned maximally to eyes in and eyes out (with average eye separation as the neutral point), two opponent pools tuned maximally to narrow face width and wide face width (with average face width as the neutral point), and so on.¹⁰

¹⁰ Note again that the actual dimensions coded in face space are unknown. The examples given are not necessarily core facial attributes; they might be, but alternatively they could be constructed from their projection onto other representations that form the cardinal axes (e.g., the eigenfaces of principal-components analysis).

This idea makes a clear prediction, namely that at some level within the visual stream processing faces, there should exist individual neurons whose pattern of responses to variation in a given facial aspect falls into two groups. Each group should have broad tuning (e.g., produce at least some response to a wide range of values of the given facial aspect) and should have a response profile that increases monotonically as the facial aspect is varied from one end on the range to the other (e.g., from eyes very close together to eyes extremely far apart) and is at a maximum for facial aspects at an extreme value (e.g., eyes maximally far apart). Of the two groups of neurons, one group should be tuned to one end of the range (e.g., eyes maximally close together) and the other group to the other end of the range (e.g., eyes maximally far apart). This prediction is very different from that of a multichannel model, which would predict multiple groups of neurons, each having a Gaussian-like response profile and peaking at a wide range of different stimulus values (e.g., some peaking to eyes 1 pixel apart, some peaking to 5 pixels apart, some to 10 pixels apart, and so on).

It is important to note that we have suggested neurons with a two-pool profile should exist at some level of the visual system. Presumably, this would be the level responsible for coding facial identity in terms of a face space representing how individuals differ from the average on multiple aspects of faces. This is not to say that there might not be neurons that respond to facial attributes in different ways at earlier levels of the processing stream: We make no claims about the sort of lower level information from which an eventual two-pool (norm-based) coding mechanism might be derived.¹¹

Do neurons with the relevant two-pool profile exist? Very recent evidence suggests they do. Recording in inferotemporal cortex of Rhesus monkeys, Leopold, Bondar, and Giese (2006) showed that the response of many individual neurons increased with increasing identity strength (distance from the average face). Similarly, in a conference abstract, Freiwald et al. (2005) reported examining the response of cells in the "middle face patch" of inferotemporal cortex to cartoon faces that independently varied 11 different facial attributes (e.g., face aspect ratio, pupil size, interocular distance). Individual cells generally responded to more than one of these attributes. Tuning curves were determined, with response of the cell plotted as a function of values of an attribute to which they were sensitive. For most cells, tuning curves were monotonic rather than peaked, with maximum responses occurring at extreme attribute values. For interocular distance, for example, this corresponds to maximal response being either to eyes very close together or to eyes very far apart (and not to eyes an average distance apart or eyes slightly further apart than average, etc). The pattern in these two studies—peak responses at extreme values rather than at the average value—is exactly that expected from a two-pool coding model.

Finally, it is interesting to speculate on where the relevant cells in humans might be located. (As with any behavioral study of adaptation, our own findings cannot tell us this directly.) For adaptation to upright faces, probably the most likely region is the fusiform face area (FFA; Kanwisher, McDermott, & Chun, 1997; McCarthy, Puce, Gore, & Allison, 1997). For upright faces, the FFA appears to be involved in both recognizing facial identity (e.g., Mazard, Schiltz, & Rossion, 2006; Yovel & Kanwisher, 2005) and representing a norm-based face space (the BOLD re-

sponse increases as a face is moved along a trajectory away from the average face; Loffler et al., 2005). Also, the extreme-value-tuned monkey cells of Freiwald et al. (2005) were in a region that apparently corresponds to the human FFA (the middle face patch was located in monkeys using the same fMRI procedures usually used to locate the FFA in humans; Tsao, Freiwald, Tootell, & Livingstone, 2006). With respect to inverted faces, evidence that the particular cells involved in aftereffects for inverted faces are largely different from the cells involved for upright faces is neutral with respect to the question of whether the cells are located in a different cortical region. Aftereffects for both upright and inverted faces could potentially be coming from cells located within this general area. On the other hand, the FFA has shown poor ability to discriminate inverted faces (Mazard et al., 2006; Yovel & Kanwisher, 2005), and no imaging data are available on norm-based coding for inverted faces, leaving open the possibility that effects for inverted faces come from some quite different cortical region. Ultimately, questions of location must be resolved by imaging of adaptation aftereffects (i.e., norm shifts) for upright and inverted faces.

Conclusion

The present study has provided the first report that the distortion level of an adapting face can affect the size of the perceptual aftereffect. We also showed that the direction of this effect (increasing or decreasing) could differ for different facial attributes, a finding we suggest is attributable to the coding range required to represent these facial attributes, given their variability in natural face images. We have also begun the process of developing neural models of face aftereffects, with key results favoring a two-pool broadband model (consistent with a norm-based representation of face space). We do not claim that the model presented here will necessarily survive in its simple form in the face of future evidence. In many areas of low-level vision, initially simple models of neural coding (and associated explanations of perceptual aftereffects) have rapidly become more complex in the light of additional empirical findings. We have, however, at least started the process.

¹¹ For example, in color vision, an opponent-coding mechanism that produces perceptual aftereffects at one level is derived from a lower level representation of wavelength that is not opponent in nature and cannot, by itself, explain opposite aftereffects. Similarly, in motion perception, perceptual opponency with a neutral point is created from pooling positive and negative inputs from multiple narrow-band mechanisms coding the full range of motion directions.

References

- Atick, J. J., Griffin, P. A., & Redlich, A. N. (1996). The vocabulary of shape: Principal shapes for probing perception and neural response. *Network: Computation in Neural Systems*, 7, 1–5.
- Bartlett, J. C., & Searcy, J. (1993). Inversion and configuration of faces. *Cognitive Psychology*, 25, 281–316.
- Blakemore, C., & Campbell. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology*, 203, 237–260.

- Carbon, C.-C., & Leder, H. (2006). The Mona Lisa effect: Is "our" Lisa fame or fake? *Perception*, 35, 411–414.
- Clifford, C., Wenderoth, P., & Spehar, B. (2000). A functional angle on some after-effects in cortical vision. *Proceedings of the Royal Society London, B*, 267, 1705–1710.
- Cohen, J., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: An interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavior Research Methods, Instruments, & Computers*, 25, 257–271.
- Farkas, L. G., Hreczko, T. A., & Katic, M. J. (1994). Craniofacial norms in North American Caucasians from birth (one year) to young adulthood. In L. G. Farkas (Ed.), *Anthropometry of the head and face* (2nd ed., pp. 241–335). New York: Raven Press.
- Fechner, G. (1966). *Elements of psychophysics* (Vol. 1). New York: Holt, Rinehart & Winston. (Original work published 1860)
- Freiwald, W. A., Tsao, D. Y., Tootell, R. B. H., & Livingstone, M. S. (2005). Single-unit recording in an fMRI-identified macaque face patch: II. Coding along multiple feature axes [Abstract]. Washington, DC: Society for Neuroscience.
- Gibson, J. J. (1937). Adaptation, after-effect and contrast in the perception of tilted lines: II. Simultaneous contrast and the areal restriction of the aftereffect. *Journal of Experimental Psychology*, 20, 553–569.
- Hreczko, T. A., & Farkas, L. G. (1994). Norms of the craniofacial asymmetries in North American Caucasians. In L. G. Farkas (Ed.), *Anthropometry of the head and face* (2nd ed., pp. 359–380). New York: Raven Press.
- Hsu, S.-M., & Young, A. W. (2004). Adaptation effects in facial expression recognition. *Visual Cognition*, 11, 871–899.
- Hubel, D. H., & Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology*, 158, 267–293.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17, 4302–4311.
- Kevin Raner Software. (1993). MacCurveFit (Version 1.5.5 for PowerMacs) [Computer software]. Mt. Waverley, Victoria, Australia: Author.
- Kourtzi, Z., & Kanwisher, N. (2001, August 24). Representation of perceived object shape by the human lateral occipital complex. *Science*, 293, 1506–1509.
- Leder, H., & Bruce, V. (1998). Local and relational aspects of face distinctiveness. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 51(A), 449–473.
- Leopold, D. A., Bondar, I. V., & Giese, M. A. (2006, August 3). Norm-based face encoding by single neurons in the monkey inferotemporal cortex. *Nature*, 442, 572–575.
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level after effects. *Nature Neuroscience*, 4, 89–94.
- Leopold, D. A., Rhodes, G., Muller, K., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Proceedings of the Royal Society of London, Series B*, 272, 897–904.
- Little, A. C., DeBruine, L. M., & Jones, B. C. (2005). Sex-contingent face after-effects suggest distinct neural populations code male and female faces. *Proceedings of the Royal Society of London, Series B*, 272, 2283–2287.
- Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H. R. (2005). fMRI evidence of the neural representation of faces. *Nature Neuroscience*, 8, 1386–1390.
- MacLin, O. H., & Webster, M. A. (2001). Influence of adaptation on the perception of distortions in natural images. *Journal of Electronic Imaging*, 10, 100–109.
- Maddess, T., McCourt, M. E., Blakeslee, B., & Cunningham, R. B. (1988). Factors governing the adaptation of cells in Area-17 of the cat visual cortex. *Biological Cybernetics*, 59, 229–236.
- Mather, G. (1980). The movement aftereffect and a distribution-shift model for coding the direction of visual movement. *Perception*, 9, 379–392.
- Mazard, A., Schiltz, C., & Rossion, B. (2006). Recovery from adaptation to facial identity is larger for upright than inverted faces in the human occipito-temporal cortex. *Neuropsychologia*, 44, 912–922.
- McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*, 9, 605–610.
- McKone, E., Aitkin, A., & Edwards, M. (2005). Categorical and coordinate relations in faces, or Fechner's law and face-space instead? *Journal of Experimental Psychology: Human Perception and Performance*, 31, 1181–1198.
- McKone, E., & Edwards, M. (2005). [Coding range and not symmetry determines aftereffect size for mouth shift]. Unpublished raw data.
- Morikawa, K. (2005). Adaptation to asymmetrically distorted faces and its lack of effect on mirror images. *Vision Research*, 45, 3180–3188.
- Movshon, J. A., & Lennie, P. (1979, April 26). Pattern-selective adaptation in visual cortical neurons. *Nature*, 278, 573–599.
- O'Toole, A. J., Abdi, H., Deffenbacher, K. A., & Valentin, D. (1995). A perceptual learning theory of the information in faces. In T. Valentine (Ed.), *Cognitive and computational aspects of face recognition* (pp. 159–182). London: Routledge.
- Over, R. (1971). Comparison of normalization theory and neural enhancement explanation of negative aftereffects. *Psychological Bulletin*, 75, 225–243.
- Perrett, D. I., Smith, P. A. J., Potter, D. D., Mistlin, A. J., Head, A. S., Milner, A. D., & Jeeves, M. A. (1985). Visual cells in the temporal cortex sensitive to face view and gaze direction. *Proceedings of the Royal Society of London B*, 223, 293–317.
- Regan, D., & Hamstra, S. J. (1992). Shape-discrimination and the judgment of perfect symmetry: Dissociation of shape from size. *Vision Research*, 32, 1845–1864.
- Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W. G., & Nakayama, K. (2003). Fitting the mind to the world: Face adaptation and attractiveness aftereffects. *Psychological Science*, 14, 558–566.
- Rhodes, G., Jeffery, L., Watson, T. L., Jaquet, E., Winkler, C., & Clifford, C. W. (2004). Orientation-contingent face aftereffects and implications for face coding mechanisms. *Current Biology*, 14, 2119–2123.
- Rhodes, G., Robbins, R., Jaquet, E., McKone, E., Jeffery, L., & Clifford, C. (2005). Adaptation and face perception: How aftereffects implicate norm based coding of faces. In G. Rhodes & C. Clifford (Eds.), *Fitting the mind to the world: Adaptation and aftereffects in high level vision* (pp. 213–240). Oxford, England: Oxford University Press.
- Rogers, B., & Graham, M. (1985). Motion parallax and the perception of three-dimensional surfaces. In D. J. Ingle, M. Jeannerod, & D. N. Lee (Eds.), *Brain mechanisms and spatial vision* (pp. 95–111). Dordrecht, the Netherlands: Martinus Nijhoff.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24, 1193–1216.
- Smith, A. T., & Edgar, G. K. (1994). Antagonistic comparison of temporal frequency filter outputs as a basis for speed perception. *Vision Research*, 34, 253–265.
- Suzuki, S. (2001). Attention-dependent brief adaptation to contour orientation: A high-level aftereffect for convexity? *Vision Research*, 41, 3883–3902.
- Suzuki, S. (2003). Attentional selection of overlapped shapes: A study using brief shape aftereffects. *Vision Research*, 43, 549–561.
- Tsao, D. Y., Freiwald, W. A., Tootell, R. B., & Livingstone, M. S. (2006, February 3). A cortical region consisting entirely of face-selective cells. *Science*, 311, 670–674.
- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology: A*, 43, 161–204.

- Valentine, T., & Bruce, V. (1986). The effects of distinctiveness in recognising and classifying faces. *Perception*, 15, 525–535.
- Watson, T., & Clifford, C. W. G. (2003). Pulling faces: An investigation of the face distortion aftereffect. *Perception*, 32, 1109–1116.
- Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004, April 1). Adaptation to natural facial categories. *Nature*, 428, 557–561.
- Webster, M. A., & MacLin, O. H. (1999). Figural aftereffects in the perception of faces. *Psychonomic Bulletin & Review*, 6, 647–653.
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5, 598–604.
- Wilson, H. R., Ferrera, V. P., & Yo, C. (1992). A psychophysically motivated model for two-dimensional motion perception. *Visual Neuroscience*, 9, 79–97.
- Wolfe, J. M. (1987). The vernier aftereffect. *Perception*, 16, 593–597.
- Yamashita, J. A., Hardy, J. L., De Valois, K. K., & Webster, M. A. (2005). Stimulus sensitivity of figural aftereffects for faces. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 420–437.
- Yovel, G., & Kanwisher, N. (2004). Face perception: Domain specific, not process specific. *Neuron*, 44, 889–898.
- Yovel, G., & Kanwisher, N. (2005). The neural basis of the behavioural face-inversion effect. *Current Biology*, 15, 2256–2262.
- Zhao, L., & Chubb, C. (2001). The size-tuning of the face-distortion after-effect. *Vision Research*, 41, 2979–2994.

Received March 18, 2005

Revision received September 11, 2006

Accepted September 25, 2006 ■

Correction to Schlaghecken et al. (2006)

In the article “Dissociating Local and Global Levels of Perceptuo-Motor Control in Masked Priming” by Friederike Schlaghecken, Howard Bowman, and Martin Eimer (*Journal of Experimental Psychology: Human Perception and Performance*, 2006, Vol. 32, No. 3, pp. 618–632), Figure 5 was inadvertently duplicated in the production process and was incorrectly substituted in place of the original Figure 6 submitted by the authors. The correct figure and caption that should have appeared for Figure 6 are listed below:

Figure 6. Results of Experiment 2. Reaction times (RTs; line graphs) and error rates (bar graphs) in compatible, neutral, and incompatible trials under two-alternative choice task (2ACT) conditions (black circles, black bars) and under four-alternative choice task (4ACT) conditions (white diamonds, white bars). ISI = interstimulus interval.