Benton, C. P., & Skinner, A. L. (2015). Deciding on race: A diffusion model analysis of race-categorisation. *Cognition*, *139*, 18-27. https://doi.org/10.1016/j.cognition.2015.02.011
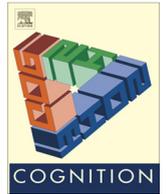
Deciding on race: A diffusion model analysis of race-categorisation

Christopher P. Benton [a,*], Andrew L. Skinner [a,b]

[a] Department of Experimental Psychology, University of Bristol, 12a Priory Road, Bristol BS8 1TU, UK
[b] MRC Integrative Epidemiology Unit, University of Bristol, 12a Priory Road, Bristol BS8 1TU, UK

ARTICLE INFO

ABSTRACT

It has long been known that a person's race can affect their decisions about people of another race; an observation that clearly taps into some deep societal issues. However, in order to behave differently in response to someone else's race, you must first categorise that person as other-race. The current study investigates the process of race-categorisation. Two groups of participants, Asian and Caucasian, rapidly classified facial images that varied from strongly Asian, through racially intermediate, to strongly Caucasian. In agreement with previous findings, there was a difference in category boundary between the two groups. Asian participants more frequently judged intermediate images as Caucasian and vice versa. We fitted a decision model, the Ratcliff diffusion model, to our two choice reaction time data. This model provides an account of the processes thought to underlie binary choice decisions. Within its architecture it has two components that could reasonably lead to a difference in race category boundary, these being *evidence accumulation rate* and *a priori bias*. The latter is the expectation or prior belief that a participant brings to the task, whilst the former indexes sensitivity to race-dependent perceptual cues. Whilst we find no good evidence for a difference in a priori bias between our two groups, we do find evidence for a difference in evidence accumulation rate. Our Asian participants were more sensitive to Caucasian cues within the images than were our Caucasian participants (and vice versa). These results support the idea that differences in perceptual sensitivity to race-defining visual characteristics drive differences in race categorisation. We propose that our findings fit with a wider view in which perceptual adaptation plays a central role in the visual processing of own and other race.

1. Introduction

The other-race effect has been a topic of scientific investigation for some 40 years or so; it's known by a number of other names, including the cross race effect and the own race advantage. These terms primarily refer to what is now a well established finding, namely that we have a tendency to better recognise the faces of members of our own race than those of another (Malpass & Kravitz, 1969; Meissner & Brigham, 2001). Clearly, this finding speaks to some deep societal concerns tied as it is to the issue of race relations, and with the obvious connections with the criminal justice system and its use of eyewitness testimony.

It is not therefore surprising that the body of research dealing with the other-race effect is considerable, and that it is comprised of contributions from a variety of different areas. Fortunately, over the past decade or so there have been a number of excellent reviews of the other race-effect, both in terms of results and underlying theories (Hugenberg, Young, Bernstein, & Sacco, 2010; Meissner & Brigham, 2001; Sporer, 2001; Young, Hugenberg,

* Corresponding author. Tel.: +44 (0)117 928 8542.
E-mail address: chris.benton@bristol.ac.uk (C.P. Benton).

Bernstein, & Sacco, 2012). In regards to the latter, there are two main strands of reasoning that have been used to explain the other-race effect, these being perceptual theories and social-cognitive theories. Social-cognitive theories are concerned with concepts such as motivation, attention and stereotyping (Rodin, 1987); whilst perceptual theories propose that it is differences in perceptual processing that underlie the other-race effect (Byatt & Rhodes, 2004; Tanaka & Farah, 1993).

A recent analysis makes the point that there is strong support for both social-cognitive and perceptual mechanisms underlying the other-race effect, and describes a general theory that combines elements of the two (Hugenberg, Wilson, See, & Young, 2013; Hugenberg et al., 2010; Young & Hugenberg, 2012). In this framework, an initial categorisation by race, accompanied by a motivation to treat races differently, allows differential perceptual processing which can then impact upon whether a group are well individuated or not. Clearly, one of the crucial parts of this theory, and the area addressed by the present study, is the role that the initial race categorisation plays.

In comparison to the primary other-race effect finding (the *same-race recognition advantage*), race categorisation has received substantially less attention. However two general findings have emerged from the race categorisation literature. The first is that people show an *other-race categorisation advantage* – when categorising by race, they categorise other-race faces more quickly and more accurately than own-race faces (Ge et al., 2009; Levin, 1996, 2000; Valentine & Endo, 1992). The second finding is that there is a shift in the category boundary such that a face morphed to be intermediate between two races is likely to be judged as other-race by both races (Webster, Kaping, Mizokami, & Dumahel, 2004).

In terms of the classification advantage, there are a couple of strong current theories. The first is the face-space approach. This relies on the idea that we represent faces using a multi-dimensional space (Lee, Byatt, & Rhodes, 2000; Valentine, 1991). The dimensions of this space are tuned by our visual diet of faces, such that commonly occurring sets of faces are better differentiated within this space. Consequently, other-race faces are seen as being grouped more closely together within face-space, whilst own-race faces are more widely separated from one another. Thus, other-race faces lie, on average, closer to the centre of their grouping than do own-race faces. From the observer's perspective, other-race faces are more prototypically characteristic of their race than are own-race faces and are therefore more readily categorised by race (Valentine, 1991; Valentine & Endo, 1992).

The other explanation offered is the idea that race is a feature (Levin, 1996, 2000). This idea ties back into visual search and the finding that the presence of a feature is detected more readily than its absence (Treisman & Gormican, 1988; Treisman & Souther, 1985). So in this case, other-race is signified by a feature (or set of features) whereas own-race is signified by a lack of other-race features. This idea can readily be used to explain the second race-categorisation finding, that shift in category boundary. If we consider a racially indeterminate face that contains both Caucasian and Asian features, Caucasian participants will be comparatively more sensitive to the Asian features and will therefore be more likely to classify the face as Asian (and vice versa for Asian participants).

The shift in category boundary has also been explained as a consequence of visual adaptation (Webster et al., 2004). In this view, our perceptual systems adjust themselves to their perceptual input. This is thought to occur primarily for a number of reasons, these being principally self calibration and efficiency of coding (Andrews, 1964; Clifford et al., 2007; Rhodes, Watson, Jeffery, & Clifford, 2010; Webster & MacLeod, 2011). The general heuristic for predicting the effect of adaptation is that the category boundary always moves towards the frequently viewed stimulus (the adapter). If we again make that reasonable assumption that Caucasians see more Caucasian faces than Asians (and vice versa) then this would account for the different category boundaries across the two populations.

Alternatively, the shift in category boundary might instead be driven by conservatism in our willingness to judge faces as own-race; we might simply require more evidence to reach an own-race conclusion than an other-race conclusion. Such a pattern is readily predicted by the in-group over-exclusion effect (Leyens & Yzerbyt, 1992; Rubin & Paolini, 2014) in which precisely this bias in required information is associated with in-group/out-group classification. Note that this bias in required information can also explain the faster responses for other-race classification – if less information needs to be accrued for an other-race decision, then presumably less time will be needed for that accrual process.

It is the process of race categorisation that forms the focus of the present study. In particular, we examine the shift in category boundary and attempt to see what processes underlie that shift. The novelty of the current study is that, from its inception, we set out to use a decision modelling approach. The modelling approach that we employ, the Ratcliff diffusion model (Ratcliff, 1978; Ratcliff & Rouder, 1998), allows us to determine whether race-categorisation differences are driven by differential sensitivity to the race-determining perceptual information, or to differences in the information required to reach a race-categorisation decision.

2. The diffusion model

Decision models seek to account for our behaviour during simple reaction time tasks; a good example of such a task being the rapid binary classification of faces as either Asian or Caucasian. Such tasks produce a series of responses (such as race judgements) along with their associated latencies. A traditional approach might take these responses and compress them down to a single number, such as the median reaction time for correct responses. In contrast, decision models are designed to account for the full latency distributions for both correct and incorrect responses. Thus, they provide a more comprehensive account than more traditional approaches. Importantly, decision models are formulated using plausible underlying psychological processes; fitting decision models to reaction time data makes a direct connection between the

pattern of the data and those underlying mechanisms. Such models have been successfully applied within a wide variety of experimental paradigms (Van Ravenzwaaij, Brown, & Wagenmakers, 2011), and are increasingly understood in terms of underlying neural mechanisms (Gold & Shadlen, 2007; Mulder, van Maanen, & Forstmann, 2014; Shadlen & Kiani, 2013).

The decision model employed in the current study, and the most influential to date, is the Ratcliff diffusion model (Ratcliff & Rouder, 1998) In general terms, the model works by proposing that evidence for a decision is accrued over time. Once the weight of evidence passes a decision threshold, then the corresponding response is made. The model has a number of parameters, which are outlined in Fig. 1 and are as follows. An *evidence accumulation rate* (v) is associated with a stimulus. This evidence accumulation rate reflects the average rate of information accrual for that type of stimulus. In the context of the present study, it is an amalgam of the amount of information conveyed by the stimulus and the participant's sensitivity to that information. There is variability in the rate of accumulation both within a trial (s) and between trials (s_v). The variability within a trial is the diffusion parameter (s) and it is this that allows the accumulated evidence to vary within a trial such that the trajectory of the evidence may cross either of the response boundaries – the model therefore accounts for both correct and incorrect responses. The wiggly dotted lines shown in Fig. 1 show three examples of trajectories that information accumulation may take within a trial. When the process reaches a boundary, then the relevant decision is made and the process terminates.

Evidence accrual starts from a certain level (βa) that lies between the two response boundaries which are separated by a distance, *a*. This latter denotes the amount of evidence separating the two decisions. As *a* (generally termed the boundary separation) increases, more evidence is needed before a decision may be made – boundary separation can therefore be thought of as a criterion (Wagenmakers, 2009). As it increases, decision latencies tend to lengthen and error rates decrease. The parameter β can be thought of as the *a priori bias* of the participants decisions (Voss, Rothermund, & Voss, 2004). For example, if β is close to zero then many rapid and incorrect Caucasian responses are made, whilst any Asian responses tend to be slow
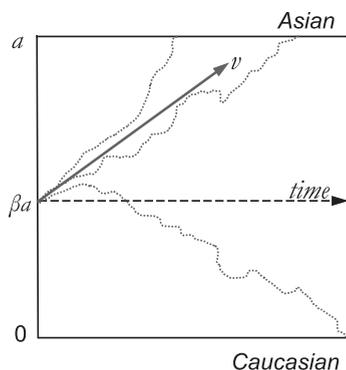
and accurate. In a similar manner to the evidence accumulation rate parameter, the starting point (βa) has a certain trial by trial variability associated with it ($s_{\beta a}$). Finally, the model contains a non-decisional component t_0 which incorporates, the time taken to make a physical response. This latter component has an associated variability, s_t. There are therefore a total of 8 parameters, however one of these, the diffusion parameter (s), is conventionally set to certain value (typically 0.1) and forms a scaling parameter.

Typically, such models are fitted to a participant's data so as to minimise the difference between the predicted and actual reaction time distributions for both the correct and incorrect and responses. Subsequent analysis of fitted parameters gives us a window onto the underlying psychological mechanisms. For example, a difference in evidence accumulation rate indicates either a difference in decision-diagnostic information in the stimulus, or a difference in sensitivity to that information.

Given a lack of the former (such as when we compare responses to the same stimulus), a difference in evidence accumulation rate in response to a visual stimulus indicates a difference in perceptual processing. Such a finding would be expected if, for example, other race categorisation was driven by a difference in visual features (Levin, 1996, 2000). In contrast, a priori bias controls the relative amounts of information necessary to reach the two different decisions. A difference in this parameter would accord with the theory that race-classification differences are driven by the in-group over-exclusion effect (Leyens & Yzerbyt, 1992; Rubin & Paolini, 2014). Note that we identify evidence accumulation rate as a perceptual factor and a priori bias as a cognitive factor; this is in the narrow sense that these are the obvious labels to attach to those components of the diffusion model in the context of the present study.

One reasonable question to ask is whether signal detection theory might provide a preferable alternative to the diffusion model. However, in the context of the current study, diffusion modelling offers important advantages. Firstly, it directly indexes the underlying psychological mechanisms because it *is* a model of those underlying psychological mechanisms (Wagenmakers, 2009); with signal detection theory, psychological processes need to inferred (Gold & Shadlen, 2007). Secondly, on a methodological point, in signal detection theory responses to target-present and target-absent stimuli are classified as correct or incorrect; calculation of criterion and sensitivity being based upon the proportion of hits and false alarms (Macmillan & Creelman, 2005). For such an analysis to occur, the experimenter needs to be able to determine whether responses are correct or incorrect. However, with the task used in the current study, the classification of such responses is far from straightforward. The problem is that the category boundary that we might wish to use to determine correctness of response is exactly the thing that our task seeks to measure. More concisely, signal detection theory applies to Type 1 "objective" tasks; ours is a Type 2 "subjective" task (Kingdom & Prins, 2009). The diffusion model is unaffected by this issue. It does not care whether

Fig. 1. Diagrammatic outline of the Ratcliff diffusion model (see text for details).

a response is correct or incorrect, it cares only which decision boundary has been crossed.

3. Method

3.1. Participants

There were two groups of participants, Asian and Caucasian. The majority (18) of the 20 Asian participants were students at the University of Bristol, the remaining two were in full-time employment. Age range was 19–30 years ($\mu = 22$, $\sigma = 3$), and 13 were female. Asian participants were not permanent residents of the UK and were volunteers recruited by approaching University of Bristol cultural societies. Six were Japanese, six were from Hong Kong, three were Singaporean, three were Malaysian, one was from Brunei, and one was from Thailand. Their mean period of stay in United Kingdom was 47 months. The 20 Caucasian participants were University of Bristol students and were each paid £5 to take part in the study. Their age range was 20–37 ($\mu = 24$, $\sigma = 4$), and 13 were female.

3.2. Materials

The stimuli for the race perception task comprised nine images of faces taken from a race morph sequence (Fig. 2). In order to create the stimuli, 25 neutral Caucasian male faces and 25 neutral Caucasian female faces were combined to form an average Caucasian face; 25 neutral Asian male faces and 25 neutral Asian female faces were combined to form an average Asian face (Stephan et al., 2005). The facial shape and texture was morphed from the average Caucasian face to the average Asian face, using the program PsychoMorph (Tiddeman, Burt, & Perrett, 2001). Extrapolating beyond the ends of the morph sequence allowed us to create hyper-Caucasian and hyper-Asian faces. We use 0% morph to indicate the average Caucasian face and 100% morph to indicate the average Asian face. Our 9 image sequence ranged between −50% morph (a hyper-Caucasian) to 150% (a hyper-Asian) in steps of 25%. All images were converted to greyscale and the edges of the faces were blurred to display mean luminance.

Stimuli were presented on an Iiyama MA203DTD 22″ CRT monitor with display resolution 1280 × 1024 pixels and refresh rate 75 Hz. The average interocular distance of the pictures used was 144 pixels, screen resolution was 15.3 pixels per cm. Stimulus display and data collection was accomplished using DMDX (Forster & Forster, 2003), responses were gathered using a Microsoft serial mouse (Plant, Hammond, & Whitehouse, 2003). Participants were seated at a viewing distance of

approximately 1 m from the monitor, which was the only major light source in an otherwise darkened room.

3.3. Procedure

Participants performed a simple binary choice reaction time task in response to facial pictures from our morph sequence which varied from strongly Caucasian through racially indeterminate to strongly Asian. Participants held the mouse in both hands and pressed the left mouse button if they thought that the face appeared to be Asian and the right mouse button if the face appeared to be Caucasian. Each image was presented until the participant responded. There was a 150 ms inter stimulus interval between the gathering of a response and the presentation of the next stimulus. participants were instructed to respond quickly and accurately. Prior to starting the blocks, participants were given a short practice block (18 trials, 2 of each of the 9 levels in random order) in order to familiarise themselves with the task.

The decision model that we used fits the reaction time distributions of the correct and incorrect responses (Vandekerckhove & Tuerlinckx, 2008). In order to apply reaction time modelling to our data, we needed to gather a substantial amount of data for each participant. To this end, participants completed six blocks of 180 trials with each block containing all 9 stimulus levels presented 20 times each in a completely random order. Each block took about 5 min to complete, subjects took short self-regulated breaks between blocks.

4. Results

The manner in which we have carried out the analyses described below is partly motivated by recent criticisms of psychological research that have raised the issue of researcher degrees of freedom (Simmons, Nelson, & Simonsohn, 2011). This latter refers to the possibility that researchers can potentially manipulate their choices of covariates, statistical tests reported, outlier identification and so on, in order to obtain interesting statistically significant results. So as to give a detailed description of our data, where appropriate, we use standard boxplots (Tukey, 1977) in which each box shows the lower quartile, median and upper quartile. Individual points plotted in this graphs show outliers, defined as those points more than 1.5 times the interquartile range beyond the box limits. The whiskers extend to the position of the most extreme non-outliers. Where appropriate, we report results both including and excluding boxplot outliers. Furthermore, as described below, for our computational modelling we report tests for a variety of model architectures. In taking this

Fig. 2. Images used in the study; a morphed sequence running from hyper-Caucasian to hyper-Asian.

approach, our aim has been to demonstrate the generality and robustness of our results to the reader. Effect sizes (Cohen's *d*) were calculated by using the standard deviations of the measures within conditions rather than their paired values (Dunlap, Cortina, Vaslow, & Burke, 1996). Where differences in simple binary comparisons are reported in the text below, the numbers in brackets show the 95% confidence limits of the difference.

4.1. Initial analysis

First we looked at where the category boundary lies for each participant by fitting, an S-shaped curve (a cumulative normal distribution) to the proportion of Asian responses as a function of percent morph (Wichmann & Hill, 2001). Here we are simply treating our task as a psychophysical method of constants task and taking a standard curve fitting approach to estimate the point of subjective equality (PSE). In this case the PSE is the point on the morph sequence at which a participant is equally like to judge an image as either Asian or Caucasian. These balance points are plotted, for the two groups of participants, in Fig. 3. In agreement with Webster et al. (2004), there is a clear difference in balance points (all data, $t(38) = 2.83$, $p < .01$, $d = .92$; without boxplot outliers $t(36) = 2.31$, $p = .027$, $d = .77$) with those for the Caucasian participants being closer to the Caucasian end of the morph sequence. The mean Asian balance point was 44.4%, mean Caucasian balance point was 34.7%, giving a difference of 9.73% (2.78–16.68). With outliers removed the mean Asian balance point remained at 44.4%, Caucasian was 37.2%, difference was 7.27% (.88–13.66).

4.2. Computational modelling

Having reproduced Webster et al.'s (2004) category boundary difference, the next stage is to model the data. To this end we use the Diffusion Model Analysis Toolbox
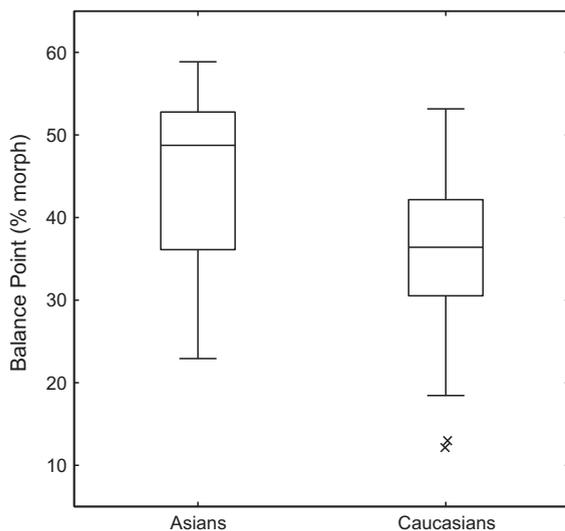

Fig. 3. Boxplot showing balance points for the Asian and Caucasian participants.

(Vandekerckhove & Tuerlinckx, 2008). Before applying the model to the data, we first look at the distribution of all reaction times gathered over the entire experiment. The aim here is to apply conservative criteria to trim obvious reaction time outliers (Ratcliff, 1993). We first apply a lower threshold of 200 ms, discarding those 33 faster responses (out of a total of 43,200 responses). This threshold marks the lower limit of the obvious reaction time distribution; the discarded responses range as low as 15 ms and are likely the result of random rapid button presses. We next apply an upper threshold of 2000 ms, discarding those 1467 responses (some 3.4% of the total number of responses) that fall above that threshold. This upper threshold marks a point on the tail of the distribution where the distribution's height has reduced to about 1.5% of its maximum and where the distribution appears to be increasingly dominated by late non-characteristic responses that range up to 10 s. Please note that we report results for both trimmed and non-trimmed reaction time distributions.

We fit a separate set of model parameters for each participant. Within each participant all levels of the stimulus have a common starting point (βa) and boundary separation (a). In fact, the only parameter that we allow to vary across stimulus levels is the evidence accumulation rate (v) as it is this that should be associated directly with changes in stimulus strength. We show the results of two sets of models, the first where we fit all 9 levels of the stimulus, the second where we fit only a subset of these. For the latter we were motivated by the suggestion that some model parameters (such as boundary separation) may change as a function of evidence accumulation rate (Hanks, Mazurek, Kiani, Hopp, & Shadlen, 2011). Any such effect should be greatly reduced by our choice of subset, which comprised the middlemost three levels falling around the category boundary (so 25%, 50% and 75% morph).

We set the model to fit the reaction time distributions by binning the reaction times using percentiles (10%, 30%, 50%, 70% and 90%). The DMAT toolbox provides two methods of model fitting (chi squared and maximum likelihood) – we used both of these methods. For the sake of comprehensibility results from only one model (all stimulus levels with chi-squared fit applied to trimmed RTs) are plotted and reported in the main text. These options present a good choice for the accurate recovery of parameter estimates (Ratcliff & Tuerlinckx, 2002). However, for pertinent findings, we tabulate the results of our statistical tests on the outputs of all models applied to both trimmed and untrimmed RTs. Furthermore, we also tabulate statistical outputs both with and without the inclusion of boxplot outliers. Note that all graphs show comparisons between participant groups (Asians and Caucasians) on their *x*-axes.

4.3. Modelling results

We are interested in attempting to explain the difference in category boundary between two groups of participants. In terms of the decision model that we have used, there are two potential elements of the model that could readily give rise to that difference. The first is the initial a

priori bias of the participants (β); the second is evidence accumulation rate. Looking at the first of these, we find no good evidence for a difference in a priori bias (β) between the two groups; the full set of t tests, differences and effect sizes is shown in Table 1, data are plotted in Fig. 4a.

Fig. 4b shows a plot of evidence accumulation rates across the two groups. To analyse these data we conducted a mixed ANOVA with one between participant factor (participant race) and one within participant factor (stimulus level, with 9 levels). Mauchly's test indicated that the assumption of sphericity had been violated ($X^2(35) = 484$, $p < .001$) therefore, where appropriate, degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity ($\varepsilon = .163$). Our results show, an entirely predictable main effect of stimulus level ($F(1.30,49.4) = 291$, $p < .001$), however there is no good evidence for an interaction between stimulus level and participant race ($F(1.30,49.4) = .53$, $p = .515$). Importantly, there is a main effect of participant race ($F(1,38) = 7.93$, $p < .01$) which we plot if Fig. 4c by averaging over stimulus level. When expressed in this form the between-participant comparison in our ANOVA can be repeated using a simple t test. The results of such tests applied to various model architectures and data treatments are shown in Table 2. Our findings show that evidence accumulation rates are reliably higher for Asian participants than for Caucasian participants. This indicates that our Asian participants appear comparatively more sensitive to the Caucasian information within the images, whilst our Caucasian participants appear comparatively more sensitive to the Asian information in those images.

The clear conclusion to draw is that the difference in category boundary that we find between our two groups of participants is caused by the difference in evidence accumulation rate. To take a concrete example, if we consider our intermediate 50% morph (see Fig. 2), this stimulus provides more information about Asian-ness to our Caucasian participants than it does to our Asian participants. Consequently, this intermediate stimulus is more readily classified as Asian by our Caucasian participants than our Asian participants. This means that the category boundary for our Caucasian participants should lie closer to the Caucasian end of our morph sequence than it does for our Asian Participants – which it does.

5. Discussion

We looked at the race categorisation of faces. Our participants performed a simple binary choice reaction time task in response to facial pictures from a morph sequence that varied from strongly Caucasian through racially indeterminate to strongly Asian. In agreement with previous research (Webster et al., 2004), we found a difference in category boundary between our Caucasian participants and our Asian participants. In comparison to Caucasians, the category boundary for Asian participants lay closer to the Asian end of the morph sequence.

We modelled participant responses using a process model of binary choice reaction times – the Ratcliff diffusion model (Ratcliff, 1978). Within this model there are two components that could readily give rise to the differences in category boundary, these being *a priori bias* and *evidence accumulation rate*. The a priori bias parameter controls the relative balance of the amount of information that is needed to make a decision. So for example, if more information is required to reach a "this face is Asian" decision than a "this face is Caucasian" decision, then, all things being equal, a Caucasian decision is more likely.

In contrast, evidence accumulation rate is an amalgam of the amount of information conveyed by a stimulus and the participant's sensitivity to that information. So, with a racially intermediate face, comprising a mix of Caucasian and Asian information, if a person is more sensitive to the Caucasian information then a Caucasian decision becomes more likely. In the context of the present

Table 1

Results of t tests applied to biases generated using various models. First column indicates whether the reaction times were trimmed prior to modelling (1) or not (0). Second column indicates whether responses to all 9 stimulus levels were modelled (1) or whether only the middle 3 levels were used (0). Third column indicates whether the model was fitted by minimising Chi squared (1) or maximum likelihood (0). Fourth column indicates whether outliers were stripped out (1) or not (0) prior to statistical analysis. Outliers were determined using the boxplot technique applied to the mean drift rates. Remaining columns show, from left to right, Asian mean, Caucasian mean, difference (Caucasian mean minus Asian mean), lower 95% confidence limit of the difference, upper 95% confidence limit, number of degrees of freedom, t statistic, p value and effect size (Cohen's d).

Trimmed	All	Chi	Outliers	Asian	Cauc	Diff	Lower	Upper	df	t stat	p	d
1	1	1	0	0.51	0.53	0.021	−0.012	0.054	38	1.29	.205	0.418
1	1	1	1	0.51	0.52	0.014	−0.016	0.045	37	0.93	.357	0.307
1	1	0	0	0.51	0.53	0.022	−0.008	0.052	38	1.47	.150	0.477
1	1	0	1	0.51	0.52	0.016	−0.012	0.044	37	1.13	.267	0.371
1	0	1	0	0.49	0.48	−0.006	−0.047	0.035	38	−0.31	.762	0.099
1	0	1	1	0.49	0.49	0.003	−0.035	0.040	37	0.15	.878	0.051
1	0	0	0	0.50	0.49	−0.011	−0.060	0.038	38	−0.46	.647	0.150
1	0	0	1	0.48	0.49	0.005	−0.030	0.041	35	0.30	.764	0.102
0	1	1	0	0.50	0.53	0.026	−0.003	0.056	38	1.81	.078	0.587
0	1	1	1	0.50	0.52	0.020	−0.007	0.048	37	1.49	.144	0.491
0	1	0	0	0.50	0.53	0.027	−0.002	0.056	38	1.89	.066	0.613
0	1	0	1	0.50	0.52	0.021	−0.006	0.048	37	1.58	.123	0.519
0	0	1	0	0.49	0.49	0.008	−0.027	0.042	38	0.45	.657	0.145
0	0	1	1	0.49	0.49	0.008	−0.027	0.042	38	0.45	.657	0.145
0	0	0	0	0.49	0.50	0.007	−0.029	0.043	38	0.39	.701	0.126
0	0	0	1	0.49	0.49	0.000	−0.034	0.035	37	0.02	.984	0.007
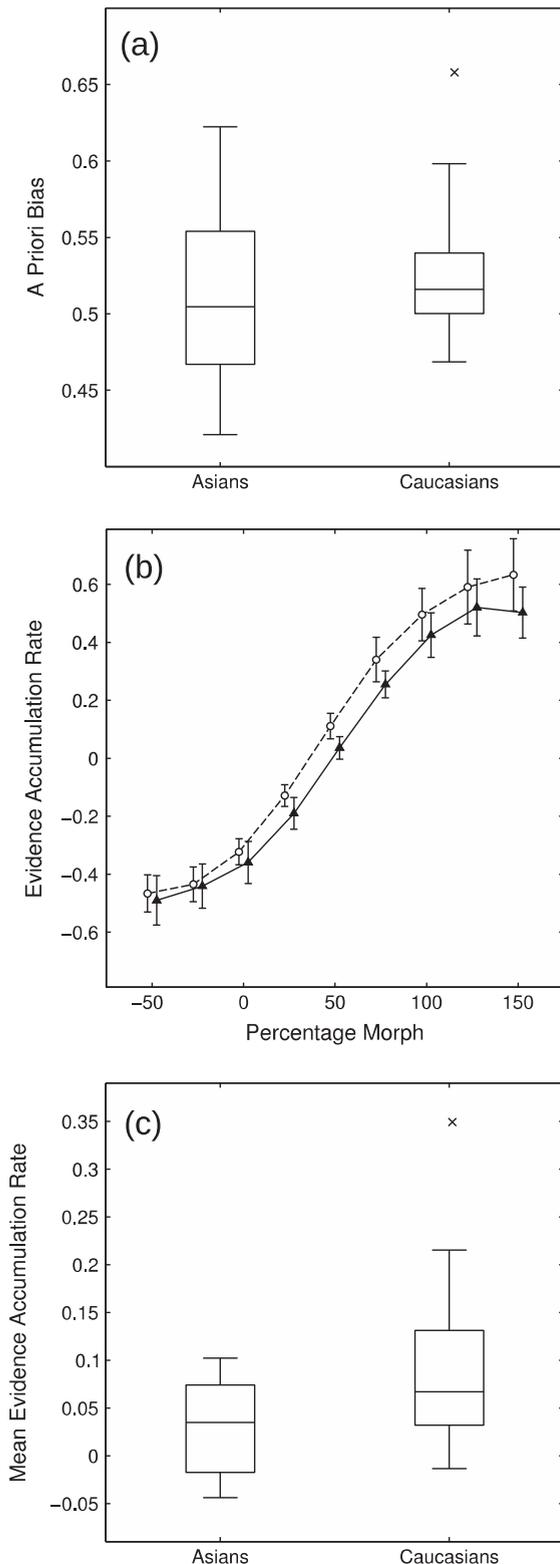
Fig. 4. Boxplots showing (a) bias, and (c) mean drift rate for Asian and Caucasian participants, (b) shows drift rates for Asian (filled triangles) and Caucasian (unfilled circles) participants with error bars showing 95% confidence limits; the two data sets are horizontally offset from one another for clarity.

study, evidence accumulation rate indexes sensitivity to the perceptual information contained within the stimulus, whilst a priori bias is about how that information is used. Bias is present before the stimulus is presented, being part of the framework for decision making, whilst evidence accumulation rate is intrinsically related to the perceptual information conveyed by the stimulus.

We found no strong evidence for a difference in a priori bias between our two groups of participants. The idea that race categorisation differences arise as a consequence of differing evidence requirements for own-race and other-race, such as predicted by Leyens & Yzerbyt's (1992) in-group over-exclusion effect, is therefore not supported by our evidence. However we did find good evidence for a difference in evidence accumulation rate. Of course, the most reasonable conclusion to draw from this finding is that the shift in category boundary is primarily caused by a difference in evidence accumulation rate. The participants in our study were more sensitive to the race-classifying perceptual information in the other race images, than they were to that of their own race. How does such a finding accord with those other mechanisms proposed to explain race categorisation differences?

We turn first to the notion of face space, which has previously been used to explain the race-categorisation advantage (Lee et al., 2000; Valentine, 1991). In this model there is no obvious means of accounting for the category boundary shift. Furthermore, it is difficult to think how concepts such as information accrual and decision boundaries might work within the face-space framework. Face-space deals with information already arrived, the information being represented by a point within face-space. As it currently stands, face-space does not therefore provide a ready framework for our findings; however the idea can likely be extended to incorporate decision processes. It is certainly possible that any such extension might incorporate an explanation for race categorisation differences.

We turn next to Levin's (1996) feature-selection model. In this, other-race is defined by a visual feature whilst own race is not. This idea does fit well with our drift-rate finding. Our stimuli were drawn from a morph sequence that ran from Caucasian to Asian. This means that you gradually get a shift in balance from Caucasian to Asian features with an intermediate stimulus containing elements of both. If Asians are comparatively more sensitive to Caucasian race-defining features, and if greater sensitivity means more rapid evidence accumulation, then this would clearly result in the shift in evidence accumulation rates curves that we describe in our results.

The idea, of being more sensitive to the race-defining features of other-races, sits well with the perceptual adaptation account of the race-categorisation boundary shift (Webster & MacLeod, 2011). This latter can be outlined as follows. Let us say that you have two pools of neurons, one better tuned to Asian faces and one better tuned to Caucasian faces. Let us also say that you have a mixed race face intermediate between Asian and Caucasian such that the two pools produce equal amounts of activity – this image therefore lies on the category boundary. Adaptation works by reducing the sensitivity of neurons that are repeatedly stimulated. So if you now view a

Table 2

Results of *t* tests applied to mean evidence accumulation rates generated using various models. First column indicates whether the reaction times were trimmed prior to modelling (1) or not (0). Second column indicates whether responses to all 9 stimulus levels were modelled (1) or whether only the middle 3 levels were used (0). Third column indicates whether the model was fitted by minimising Chi squared (1) or maximum likelihood (0). Fourth column indicates whether outliers were stripped out (1) or not (0) prior to statistical analysis. Outliers were determined using the boxplot technique applied to the mean evidence accumulation rates. Remaining columns show, from left to right, Asian mean, Caucasian mean, difference (Caucasian mean minus Asian mean), lower 95% confidence limit of the difference, upper 95% confidence limit, number of degrees of freedom, *t* statistic, *p* value and effect size (Cohen's *d*).

Trimmed	All	Chi	Outliers	Asian	Cauc	Diff	Lower	Upper	df	t stat	p	d
1	1	1	0	0.029	0.091	0.062	0.017	0.107	38	2.82	.0077	0.91
1	1	1	1	0.029	0.077	0.049	0.012	0.085	37	2.72	.0098	0.90
1	1	0	0	0.031	0.082	0.051	0.012	0.090	38	2.67	.0110	0.87
1	1	0	1	0.031	0.082	0.051	0.012	0.090	38	2.67	.0110	0.87
1	0	1	0	0.033	0.128	0.095	0.045	0.146	38	3.84	.0004	1.25
1	0	1	1	0.033	0.128	0.095	0.045	0.146	38	3.84	.0004	1.25
1	0	0	0	0.035	0.203	0.167	0.010	0.325	38	2.15	.0378	0.70
1	0	0	1	0.035	0.131	0.095	0.034	0.156	37	3.16	.0031	1.04
0	1	1	0	0.035	0.094	0.059	0.010	0.108	38	2.44	.0193	0.79
0	1	1	1	0.035	0.080	0.046	0.004	0.088	37	2.21	.0336	0.73
0	1	0	0	0.034	0.089	0.055	0.014	0.095	38	2.73	.0095	0.89
0	1	0	1	0.034	0.089	0.055	0.014	0.095	38	2.73	.0095	0.89
0	0	1	0	0.052	0.118	0.066	0.011	0.121	38	2.42	.0205	0.78
0	0	1	1	0.040	0.118	0.079	0.028	0.129	37	3.18	.0030	1.05
0	0	0	0	0.054	0.127	0.073	0.022	0.123	38	2.91	.0060	0.94
0	0	0	1	0.054	0.127	0.073	0.022	0.123	38	2.91	.0060	0.94

particular diet of faces (say Caucasian faces) then the responsiveness of the "Caucasian neurons" will be reduced so that that mixed-race face will now be judged as Asian. If you want to find the new post-adaptation category boundary you will need a facial image that is more Caucasian than that lying on the pre-adaptation category boundary. The reasonable assumption, that increased neural sensitivity means an increased evidence accumulation rate, allows us to connect our diffusion model finding to an explanation of race categorisation through perceptual adaptation.

If perceptual adaptation can explain the race-categorisation boundary shift, can it explain that complementary effect, the race-categorisation advantage? In our above description of two populations of neurons, Asian-tuned and Caucasian-tuned, the sensitivity of neurons tuned to own-race is reduced. A race-categorisation decision based upon whichever population of neurons signals the highest output, or reaches a certain threshold first, will result in faster other-race categorisation. More generally, increased sensitivity for the detection of novel or rare stimuli has been proposed as a key component of adaptation (Barlow, 1990; McDermott, Malkoc, Mulligan, & Webster, 2010), clearly we can conceive of other-race as being comparatively novel or rare.

If adaptation explains the two race-categorisation effects can it also, as recently suggested (Rhodes, Jeffery, Taylor, Hayward, & Ewing, 2014), explain that better-known other race finding, the same-race recognition advantage? As well as the standard suppressive effect of adaptation, which we have used to account for the race-categorisation effects, there is evidence of another consequence of adaptation, this being an enhancement of discrimination or individuation around the adapting stimulus (Clifford, 2002). This enhanced discrimination may be analogous to the greater dispersion of own-race faces in the face-space framework (see below); which does raise the possibility that we can couch face-space within the wider framework of perceptual adaptation. In the

face-adaptation domain, a number of studies provide good evidence for enhancement discrimination around the adapting stimulus (Keefe, Dzhelyova, Perrett, & Barraclough, 2013; Oruç & Barton, 2011; Rhodes et al., 2010; Yang, Shen, Chen, & Fang, 2011), critically, at least in the context of the other-race effect, Rhodes et al. (2010) showed that adaptation to a prototypical Caucasian or Asian face led to increased discriminability for faces of the adapted race compared to those of the unadapted race. Thus, adaptation does appear to induce the same-race recognition advantage that we find in the other-race effect.

Perceptual adaptation can therefore provide a unified, parsimonious account of classification and discriminability. Other approaches also span classification and discriminability. For example, in face space theory, the other-race classification advantage is driven by other-race faces being more closely grouped together (and therefore more race-prototypical) than same-race faces. This difference in dispersion has been used to account for the same-race recognition advantage, increased dispersion leading to greater distinguishability (Valentine, 1991). In contrast, a mechanism such as motivation to individuate is an account of discriminability; its effects may be made manifest after an initial categorisation process (Hugenberg et al., 2010). While such a multi-mechanism framework may seem unparsimonious, it does readily account for other-race differences in the absence of perceptual cues to categorisation (Hourihan, Fraundorf, & Benjamin, 2013). The occurrence of the latter succinctly makes the point that perceptually-driven accounts cannot provide a complete account of the other-race effect.

Above, we have made the points that the general framework of perceptual adaptation can accommodate our evidence accumulation rate finding and can account for the category boundary shift, the other-race categorisation advantage and the same-race recognition advantage. Furthermore, the finding, from perceptual adaptation, that

repeated stimulation of neural mechanisms reduces their response, provides a mechanism by which differences in sensitivity to features can arise; in other words, perceptual adaptation can readily incorporate Levin's (1996) feature-selection model. Of course, this means moving away from any idea of a race-feature as something that is all-or-nothing. It requires instead the idea that we can be differentially sensitive to different sets of race-features. Put in this way, the feature and the adaptation accounts become one and the same, with our description of adaptation being an account of neural responses, the features being the properties to which those neurons are sensitive.

Theoretical approaches see perceptual adaptation as functional, its purpose being to optimise the use of our neural mechanisms to better encode our perceptual input (Clifford, 2005). In terms of a general model of the other-race effect perceptual adaptation is therefore an attractive idea – it is motivated by an idea that makes real intuitive sense, that being the optimisation of our expensive-to-build and expensive-to-maintain neural machinery. In turn, the concepts that link that motive to adaptation's behavioural effects, are again intuitively attractive. Ideas such as perceptual recalibration, contrast gain control and redundancy reduction make real sense from a fundamental ''what's a good way to engineer this neural system'' viewpoint. Furthermore, adaptation offers a ready mechanism for experience-dependent plasticity in our perception of other-race faces (Tanaka, Heptonstall, & Hagen, 2013). However it is worth noting the following point. One of the weaknesses of perceptual adaptation is that it has been hard to connect the readily observable laboratory effects of adaptation to adaptation in the real world (Ditye, Javadi, Carbon, & Walsh, 2013). Whilst it makes sense to assume that effects such as the other-race effect, other-age effect (Kuefner, Macchi Cassia, Picozzi, & Bricolo, 2008), other-gender effect (Wright & Sladden, 2003) and the other species effect (Pascalis, 2011), are all manifestations of long term adaptation, that connection still needs to be made.

Our results offer a clear story in which race categorisation differences are driven by differences in perceptual sensitivity to race-defining visual characteristics. However, we wish to emphasise our earlier point, that perceptually driven accounts cannot provide a complete account of the other race effect. It is undoubtedly the case that any such perceptual mechanisms may be paralleled, influenced or caused by factors from the social-cognitive domain. For example, in the perceptual adaptation framework, motivation to attend to own-race over other-race faces could well mediate adaptation, with greater attention causing greater adaptation and hence greater discriminability (Rhodes et al., 2011). Moreover, the wider other race-effect, which we take to mean the constellation of findings describing race-dependent response differences, is clearly a complex phenomenon in which perceptual, social and cognitive factors interact (Hugenberg et al., 2013). Further, the work reported here lies in the psychological domain. A full understanding of the wider other race effect must go well beyond such explanations to take into account the social and historical contexts within which the other race effect is so clearly embedded.

References

Andrews, D. (1964). Error-correcting perceptual mechanisms. *Quarterly Journal of Experimental Psychology, 16*(2), 104–115. http://dx.doi.org/10.1080/17470216408416355.

Barlow, H. B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. In C. Blakemore (Ed.), *Vision: Coding and Efficiency* (pp. 363–375). Cambridge, UK: Cambridge University Press. Retrieved from <http://dx.doi.org/10.1017/CBO9780511626197.034>.

Byatt, G., & Rhodes, G. (2004). Identification of own-race and other-race faces: Implications for the representation of race in face space. *Psychonomic Bulletin & Review, 11*(4), 735–741. http://dx.doi.org/10.3758/BF03196628.

Clifford, C. W. G. (2002). Perceptual adaptation: Motion parallels orientation. *Trends in Cognitive Sciences, 6*(3), 136–143.

Clifford, C. W. G. (2005). Functional ideas about adaptation applied to spatial and motion vision. In G. Rhodes & C. W. G. Clifford (Eds.), *Fitting the mind to the world: Adaptation and after-effects in high-level vision* (pp. 47–82). Oxford: Oxford University Press.

Clifford, C. W. G., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., et al. (2007). Visual adaptation: Neural, psychological and computational aspects. *Vision Research, 47*, 3125–3131.

Ditye, T., Javadi, A. H., Carbon, C.-C., & Walsh, V. (2013). Sleep facilitates long-term face adaptation. *Proceedings of the Royal Society B – Biological Sciences, 280*(1769), 20131698. http://dx.doi.org/10.1098/rspb.2013.1698.

Dunlap, W. P., Cortina, J. M., Vaslow, J. B., & Burke, M. J. (1996). Meta-analysis of experiments with matched groups or repeated measures designs. *Psychological Methods, 1*, 170–177.

Forster, K. I., & Forster, J. C. (2003). DMDX: A windows display program with millisecond accuracy. *Behavior Research Methods, Instruments, & Computers, 35*(1), 116–124. http://dx.doi.org/10.3758/BF03195503.

Ge, L., Zhang, H., Wang, Z., Quinn, P. C., Pascalis, O., Kelly, D., et al. (2009). Two faces of the other-race effect: Recognition and categorisation of Caucasian and Chinese faces. *Perception, 38*, 1199–1210. http://dx.doi.org/10.1068/p6136.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. In *Annual review of neuroscience* (Vol. 30, pp. 535–574). Palo Alto: Annual Reviews.

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *The Journal of Neuroscience, 31*(17), 6339–6352. http://dx.doi.org/10.1523/JNEUROSCI.5613-10.2011.

Hourihan, K., Fraundorf, S., & Benjamin, A. (2013). Same faces, different labels: Generating the cross-race effect in face memory with social category information. *Memory & Cognition, 41*(7), 1021–1031. http://dx.doi.org/10.3758/s13421-013-0316-7.

Hugenberg, K., Wilson, J. P., See, P. E., & Young, S. G. (2013). Towards a synthetic model of own group biases in face memory. *Visual Cognition, 21*(9–10), 1392–1417. http://dx.doi.org/10.1080/13506285.2013.821429.

Hugenberg, K., Young, S. G., Bernstein, M. J., & Sacco, D. F. (2010). The categorization-individuation model: An integrative account of the other-race recognition deficit. *Psychological Review, 117*(4), 1168–1187. http://dx.doi.org/10.1037/a0020463.

Keefe, B. D., Dzhelyova, M. P., Perrett, D. I., & Barraclough, N. E. (2013). Adaptation improves face trustworthiness discrimination. *Perception Science, 4*, 358. http://dx.doi.org/10.3389/fpsyg.2013.00358.

Kingdom, F. A. A., & Prins, N. (2009). *Psychophysics: A practical introduction* (1st ed.). London: Academic Press.

Kuefner, D., Macchi Cassia, V., Picozzi, M., & Bricolo, E. (2008). Do all kids look alike? Evidence for an other-age effect in adults. *Journal of Experimental Psychology. Human Perception and Performance, 34*(4), 811–817. http://dx.doi.org/10.1037/0096-1523.34.4.811.

Lee, K., Byatt, G., & Rhodes, G. (2000). Caricature effects, distinctiveness, and identification: Testing the face-space framework. *Psychological Science, 11*, 379–385.

Levin, D. T. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology-Learning Memory and*

Cognition, 22, 1364–1382. http://dx.doi.org/10.1037//0278-7393.22.6.1364.

Levin, D. T. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. Journal of Experimental Psychology. General, 129(4), 559–574.

Leyens, J.-P., & Yzerbyt, V. Y. (1992). The ingroup overexclusion effect: Impact of valence and confirmation on stereotypical information search. European Journal of Social Psychology, 22(6), 549–569. http://dx.doi.org/10.1002/ejsp.2420220604.

Macmillan, N. A., & Creelman, C. D. (2005). Detection theory: A user's guide. London: Lawrence Erlbaum Associates.

Malpass, R., & Kravitz, J. (1969). Recognition for faces of own and other race. Journal of Personality and Social Psychology, 13(4), 330. http://dx.doi.org/10.1037/h0028434.

McDermott, K. C., Malkoc, G., Mulligan, J. B., & Webster, M. A. (2010). Adaptation and visual salience. Journal of Vision, 10(13), 17. http://dx.doi.org/10.1167/10.13.17.

Meissner, C. A., & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces – A meta-analytic review. Psychology Public Policy and Law, 7, 3–35. http://dx.doi.org/10.1037/1076-8971.7.1.3.

Mulder, M. J., van Maanen, L., & Forstmann, B. U. (2014). Perceptual decision neurosciences – A model-based review. Neuroscience, 277, 872–884. http://dx.doi.org/10.1016/j.neuroscience.2014.07.031.

Oruç, I., & Barton, J. J. S. (2011). Adaptation improves discrimination of face identity. Proceedings of the Royal Society of London B: Biological Sciences. http://dx.doi.org/10.1098/rspb.2010.2480. rspb20102480.

Pascalis, & Wirth (2011). Recognizing the faces of other species: what can a limited skill tell us about face processing? In Rhodes, Calder, Johnson, & Haxby (Eds.), Oxford Handbook of Face Perception (pp. 719–730). Oxford, UK: Oxford University Press. Retrieved from http://www.oxfordhandbooks.com/10.1093/oxfordhb/9780199559053.001.0001/oxfordhb-9780199559053-e-037.

Plant, R. R., Hammond, N., & Whitehouse, T. (2003). How choice of mouse may affect response timing in psychological studies. Behavior Research Methods, Instruments, & Computers, 35(2), 276–284. http://dx.doi.org/10.3758/BF03202553.

Ratcliff, R. (1978). Theory of memory retrieval. Psychological Review, 85, 59–108.

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. Psychological Bulletin, 114(3), 510–532.

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. Psychological Science, 9(5), 347–356. http://dx.doi.org/10.1111/1467-9280.00067.

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. Psychonomic Bulletin & Review, 9(3), 438–481. http://dx.doi.org/10.3758/BF03196302.

Rhodes, G., Jeffery, L., Evangelista, E., Ewing, L., Peters, M., & Taylor, L. (2011). Enhanced attention amplifies face adaptation. Vision Research, 51(16), 1811–1819. http://dx.doi.org/10.1016/j.visres.2011.06.008.

Rhodes, G., Jeffery, L., Taylor, L., Hayward, W. G., & Ewing, L. (2014). Individual differences in adaptive coding of face identity are linked to individual differences in face recognition ability. Journal of Experimental Psychology-Human Perception and Performance, 40(3), 897–903. http://dx.doi.org/10.1037/a0035939.

Rhodes, G., Watson, T. L., Jeffery, L., & Clifford, C. W. G. (2010). Perceptual adaptation helps us identify faces. Vision Research, 50(10), 963–968. http://dx.doi.org/10.1016/j.visres.2010.03.003.

Rodin, M. (1987). Who is memorable to whom – A study of cognitive disregard. Social Cognition, 5(2), 144–165. http://dx.doi.org/10.1521/soco.1987.5.2.144.

Rubin, M., & Paolini, S. (2014). Out-group flies in the in-group's ointment: Evidence of the motivational underpinnings of the in-group overexclusion effect. Social Psychology, 45, 265–273. http://dx.doi.org/10.1027/1864-9335/a000171.

Shadlen, M. N., & Kiani, R. (2013). Decision making as a window on cognition. Neuron, 80(3), 791–806. http://dx.doi.org/10.1016/j.neuron.2013.10.047.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis

allows presenting anything as significant. Psychological Science, 22(11), 1359–1366. http://dx.doi.org/10.1177/0956797611417632.

Sporer, S. L. (2001). Recognizing faces of other ethnic groups: An integration of theories. Psychology, Public Policy, and Law, 7(1), 36–97. http://dx.doi.org/10.1037/1076-8971.7.1.36.

Stephan, C. N., Penton-Voak, I. S., Perrett, D. I., Tiddeman, B. P., Clement, J. G., & Henneberg, M. (2005). Two-dimensional computer generated average human face morphology and facial approximation. In M. Marks & J. Clement (Eds.), Computer-graphic facial reconstruction. Elsevier Academic Press.

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. The Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology, 46(2), 225–245.

Tanaka, J. W., Heptonstall, B., & Hagen, S. (2013). Perceptual expertise and the plasticity of other-race face recognition. Visual Cognition, 21(9–10), 1183–1201. http://dx.doi.org/10.1080/13506285.2013.826315.

Tiddeman, B. P., Burt, D. M., & Perrett, D. I. (2001). Prototyping and transforming facial texture for perception research. IEEE Computer Graphics Applications, 21, 42–50.

Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. Psychological Review, 95(1), 15–48.

Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. Journal of Experimental Psychology. General, 114(3), 285–310.

Tukey, J. W. (1977). Exploratory data analysis. Addison-Wesley Publishing Company.

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. Quarterly Journal of Experimental Psychology Section A – Human Experimental Psychology, 43, 161–204.

Valentine, T., & Endo, M. (1992). Towards an exemplar model of face processing – The effects of race and distinctiveness. Quarterly Journal of Experimental Psychology Section A – Human Experimental Psychology, 44, 671–703.

Van Ravenzwaaij, D., Brown, S., & Wagenmakers, E.-J. (2011). An integrated perspective on the relation between response speed and intelligence. Cognition, 119(3), 381–393. http://dx.doi.org/10.1016/j.cognition.2011.02.002.

Vandekerckhove, J., & Tuerlinckx, F. (2008). Diffusion model analysis with MATLAB: A DMAT primer. Behavior Research Methods, 40, 61–72.

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. Memory & Cognition, 32, 1206–1220.

Wagenmakers, E. J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. European Journal of Cognitive Psychology, 21, 641–671.

Webster, M. A., Kaping, D., Mizokami, Y., & Dumahel, P. (2004). Adaptation to natural face categories. Nature, 428, 558–561.

Webster, M. A., & MacLeod, D. I. A. (2011). Visual adaptation and face perception. Philosophical Transactions of the Royal Society B – Biological Sciences, 366(1571), 1702–1725. http://dx.doi.org/10.1098/rstb.2010.0360.

Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. Perception & Psychophysics, 63, 1293–1313.

Wright, D. B., & Sladden, B. (2003). An own gender bias and the importance of hair in face recognition. Acta Psychologica, 114(1), 101–114. http://dx.doi.org/10.1016/S0001-6918(03)00052-0.

Yang, H., Shen, J., Chen, J., & Fang, F. (2011). Face adaptation improves gender discrimination. Vision Research, 51(1), 105–110. http://dx.doi.org/10.1016/j.visres.2010.10.006.

Young, S. G., & Hugenberg, K. (2012). Individuation motivation and face experience can operate jointly to produce the own-race bias. Social Psychological and Personality Science, 3(1), 80–87. http://dx.doi.org/10.1177/1948550611409759.

Young, S. G., Hugenberg, K., Bernstein, M. J., & Sacco, D. F. (2012). Perception and motivation in face recognition: A critical review of theories of the cross-race effect. Personality and Social Psychology Review, 16(2), 116–142. http://dx.doi.org/10.1177/1088868311418987.