Face-evoked thoughts

Xingchen Zhou[1], Rob Jenkins[1]*

[1]Department of Psychology, University of York, UK

*Corresponding author

Rob Jenkins,

Department of Psychology

University of York

York YO10 5DD

UK

rob.jenkins@york.ac.uk

+44 1904 323144

Abstract

The thoughts that come to mind when viewing a face depend partly on the face and partly on the viewer. This basic interaction raises the question of how much common ground there is in face-evoked thoughts, and how this compares to viewers' expectations. Previous analyses have focused on early perceptual stages of face processing. Here we take a more expansive approach that encompasses later associative stages. In Experiment 1 (free association), participants exhibited strong egocentric bias, greatly overestimating the extent to which other people's thoughts resembled their own. In Experiment 2, we show that viewers' familiarity with a face can be decoded from their face-evoked thoughts. In Experiment 3 (person association), participants reported *who* came to mind when viewing a face—a task that emphasises connections in a social network rather than nodes. Here again, viewers' estimates of common ground exceeded actual common ground by a large margin. We assume that a face elicits much the same thoughts in other people as it does in us, but that is a mistake. In this respect, we are more isolated than we think.

Key Words

metacognition; face perception; false consensus; egocentric bias

1 Introduction

What comes to mind when you see a face? To some extent, it depends on the face—not only its physical appearance, but also the person whose face it is, and everything that goes along with that person. However, it also depends on the viewer. This is partly because different faces are known to different viewers. Some people know who Arnold Schwarzenegger is, and other people do not. But even among those who do, experiences and preferences can differ widely. A politician might think of Arnold first and foremost as the Governor of California, whereas a cinema-goer might think of him primarily as *The Terminator* (and both incarnations divide opinion for different reasons). This interplay between face and viewer raises the question of how much common ground there is in face-evoked thoughts. A face is a natural entry point to a social network, but if the same node can lead different viewers in different directions, it is not clear how much of the network is really shared.

Depth of processing becomes important here. For some visual aspects of face processing, such as categorising faces as female or male, observers' responses are highly consistent (Bruce et al., 1993; Burton, Bruce & Dench, 1993). Yet even the seemingly objective task of characterising face shape reveals wide discrepancies between viewers (Towler, White, & Kemp, 2014). Judgements of photographic likeness are similarly idiosyncratic (Hay, Young, & Ellis, 1991), and yield little consensus among viewers. These visual face processing tasks hinge on physical information, meaning that observers need only consider the face as an image. As cognition proceeds beyond physical cues to the inferences we draw from them, opportunities for divergence multiply. In an influential analysis of facial attractiveness, Hönekopp (2006) showed that private preferences are roughly as powerful as shared preferences in determining judgements. This finding came as something of a surprise, as it overturned the prevailing view at the time that agreement on such judgements among

observers was high. But subsequent analyses have also concluded that private preferences, shaped by personal experience, are often the major determinant of attractiveness judgements (Germine et al., 2015; Hehman et al., 2017). For other trait inferences from faces—notably the cardinal dimensions of trustworthiness and dominance—agreement among viewers can be high (Oosterhof & Todorov, 2008; Todorov, Said, & Verosky, 2011), though recent studies have identified a role for personal experiences in explaining differences in viewers' impressions (Sutherland, Rhodes, Burton, & Young, 2020; Sutherland et al., 2020).

The preceding studies share some important features. All of them concern the first moments of face processing. Converging evidence from ERP measures (e.g. Mouchetant-Rostaing et al., 2000) and saccadic reaction times (e.g. Ramon, Sokhn, & Calder, 2019) indicate that female and male faces can be differentiated within 150 ms of stimulus onset. Impressions of trustworthiness, dominance, and attractiveness based on facial appearance also emerge quickly—as early as 100 ms post-stimulus (Willis & Todorov, 2006; Olivola & Todorov, 2010). But the cascade of cognition that a face sets in motion can last much longer. The influential Bruce & Young (1986) framework emphasises important differences between unfamiliar and familiar face processing (see Johnston & Edmonds, 2009 for a review). Evidence from a range of memory and perception tasks indicates that perception of an unfamiliar face is closely bound to the image, being dominated by picture-level information (*physical* information; Hancock, Bruce, & Burton, 2000; Burton & Jenkins, 2011). In contrast, perception of a familiar face is closely bound to the person (*non-physical* information; Jenkins & Burton, 2011; Burton, Jenkins, & Schweinberger, 2011), leading to mandatory recognition. For example, semantic information associated with that person appears to be automatically accessed, giving rise to priming at later tests (Ellis, Young, & Flude, 1990; McNeill & Burton, 2002). As access to non-physical information depends on

prior processing of physical information, we should expect to see it later in mental chronometry. That is the finding. Interaction with semantic and emotional processes appears to peak around 400 ms and remains clear until at least 600 ms post-stimulus (Wiese et al., 2019). As stimulus associations are often idiosyncratic and often chain together (at least in the word domain; Shapiro, 1966, De Deyne et al., 2019), it is precisely in these later processes that one would most expect individual viewers to diverge. The resulting heterogeneity of responses can make analyses unwieldy. Perhaps for that reason—and undoubtedly because of applied interest in early stages of face processing (Brewer & Wells, 2011; Schultz, 2012; Phillips et al., 2018)—later stages of face processing have received less attention in cognitive research. Those studies that have examined more associative processes in face perception tend to reveal idiosyncrasy in perceiver judgements (e.g. DeBruine 2004; Verosky et al., 2018).

This brings us to a second commonality among previous studies. Responses are typically constrained to a small set of options—for example, whether a face is female or male, whether or not faces match on some dimension, or a numerical rating of a particular attribute. A participant's first thought when seeing a face might be, "She looks like my primary school teacher!", but if the task is to rate the face for trustworthiness, that reaction will not be captured as part of the data. There are some exceptions where researchers have gathered free associations to face images (e.g. Oosterhof & Todorov, 2008; Sutherland et al., 2018). However, in those cases, free associations were not the main interest. Instead, they were used to compute dimensions of variation for first impressions from faces. Experimental participants then rated faces on those dimensions using Likert scales.

Constraining responses in this way makes sense when the focus is a specific psychological mechanism. Our focus here is rather different, as our questions concern networks of social cognition. When it comes to face-evoked thoughts, little is known about the extent of overlap among viewers. But without direct insight into the minds of others, our sense of common ground can not be based on the actual extent of overlap. It can only be based on our *impression* of overlap, and that depends on certain metacognitive assumptions. What occurs to other people when they see a particular face? How does that compare with one's own experience? These questions have not been addressed experimentally, although other areas of psychology offer some important clues.

Across a wide range of situations, we are inclined to assume that others think as we do, sharing our tastes, preferences, and understanding (false consensus effects; Ross, Greene, & House, 1977; Krueger & Clement, 1994), and generally seeing the world from our own perspective. For example, we tend to overestimate the extent to which others know the things that we know (Hinds, 1999; Gilovich, Medvec, & Savitsky, 2000) and make the choices that we make (Ross, Greene, & House, 1977; Wolfson, 2000). These failures of metacognitive insight are examples of egocentric bias—the tendency to understand others from our own perspective (Krueger & Clement, 1994). Recently, egocentric biases have been demonstrated in face perception. In one example, participants in an identity matching task predicted that the faces they themselves knew would be easy for other people to match (Ritchie et al, 2015). Zhou & Jenkins (2020) showed that, in matching tasks for identity, emotional expression, and gaze direction, high performing participants attributed higher performance to other people, and low performing participants attributed lower performance to other people.

These findings demonstrate egocentric bias in early stages of face processing. Our scope here is deliberately more broad. Instead of focusing on early perceptual processes and fixed response options, we seek to capture whatever comes to mind when viewing the face. Importantly, this is not the same as establishing what the viewer knows about the seen person. At any moment, what comes to mind is only a subset of one's relevant knowledge. This is a critical distinction. It is what comes to mind, not what stays behind, that constitutes a train of thought. The current study addresses two related aspects of face-evoked thoughts— first, the degree of overlap among viewers, and second, how this overlap compares to viewers' expectations. Given the evidence of egocentric bias elsewhere in cognition, we predicted that viewers would overestimate the extent to which other people's thoughts resemble their own (a false consensus effect). We begin in Experiment 1 by asking viewers what comes to mind when they see a particular face. By focusing on the seen person, this question emphasises individual nodes in a social network. In Experiment 2, we test whether face-evoked thoughts differ for familiar and unfamiliar faces. Finally, in Experiment 3, we ask *who* comes to mind when they see a particular face. By focusing on related people, this question emphasises connections between nodes.

2 Experiment 1

The purpose of this experiment was twofold. First, we sought to quantify overlap among observers in face-evoked thoughts. Second, we sought to compare the extent of this overlap to observers' expectations. To this end, we devised a face association task comprising cognitive and metacognitive components. Participants were asked to write down whatever thoughts came to mind when they viewed a series of faces (cognitive component). Analysing these responses allowed us to quantify actual overlap among observers. We then asked the same participants to estimate how other participants' responses would compare to their own

(metacognitive component). Analysing these peer estimates allowed us to quantify expected overlap.

Our predictions concerned the numerosity, content, and order of face-evoked thoughts. At the cognitive level, we tested the following assumptions: (i) participants would differ in the number of thoughts they recorded, (ii) familiar faces would elicit more responses than unfamiliar faces, (iii) some responses to a particular face would be made by more than one participant, and (iv) salient associations would come to mind earlier than obscure associations. At the metacognitive level, we predicted the following signs of egocentric bias: (i) participants who produced a high number of responses would expect others to produce a high number of responses (and vice versa), (ii) participants would expect others to produce more responses for faces that they themselves knew, (iii) participants would overestimate the number of viewers who had the same thoughts that they themselves had, and (iv), this tendency to overestimate common ground would be strongest for thoughts that came to mind first.

2.1 Methods

2.1.1 Participants

Thirty UK students (25 female, 5 male; mean age: 19 years; age range 18–23 years) from the University of York took part in exchange for a small payment and task-related bonus. The experiments in this study were approved by the Ethics Committee at the University of York. All participants provided written informed consent.

2.1.2 Stimuli and apparatus

Ambient images of 8 familiar celebrities (4 female; 4 male) and 8 unfamiliar celebrities (public figures from outside of the UK, 4 female; 4 male) were downloaded from the internet. The names of these celebrities are listed in Appendix A. Each photo captured the whole face with no occlusions from a roughly frontal aspect. In pilot testing, 64 UK students, who did not participate in the main experiment, indicated whether or not they were familiar with each face. An independent t-test confirmed that the familiar celebrities were known to more respondents ($M = 89\%$) than the unfamiliar celebrities ($M = 10\%$) [$t(14) = 14.99$, $p < .0001$].

All photos were cropped to 570 pixels high × 380 pixels and colour printed at 72 dpi onto A4 sheets, which were collated into booklets. Pagination was randomised, so that familiar and unfamiliar faces were intermixed. To counteract possible order effects, page order was reversed for half of the participants.

2.1.3 Design

Each participant completed three separate tasks—a face association test, a metacognitive review, and a face familiarity check. In the face association test, the participant's task was to capture whatever came to mind (their 'points') as they viewed each face. Participants transcribed their thoughts into a personalised Microsoft Excel workbook, using a separate worksheet for each face. Each worksheet was headed with the prompt, "Points (Can you tell us any more?)", followed by a series of enumerated rows to accommodate separate points of information (see Figure 1). This was a self-paced task, and participants were encouraged to be as exhaustive as possible in recording their thoughts. To motivate participants to generate as many points as possible, we introduced a cash incentive of 1p per point in addition to standard participant payments. For example, generating an average of 13 points for each of the 16 faces would result in an additional £2.08 payment.
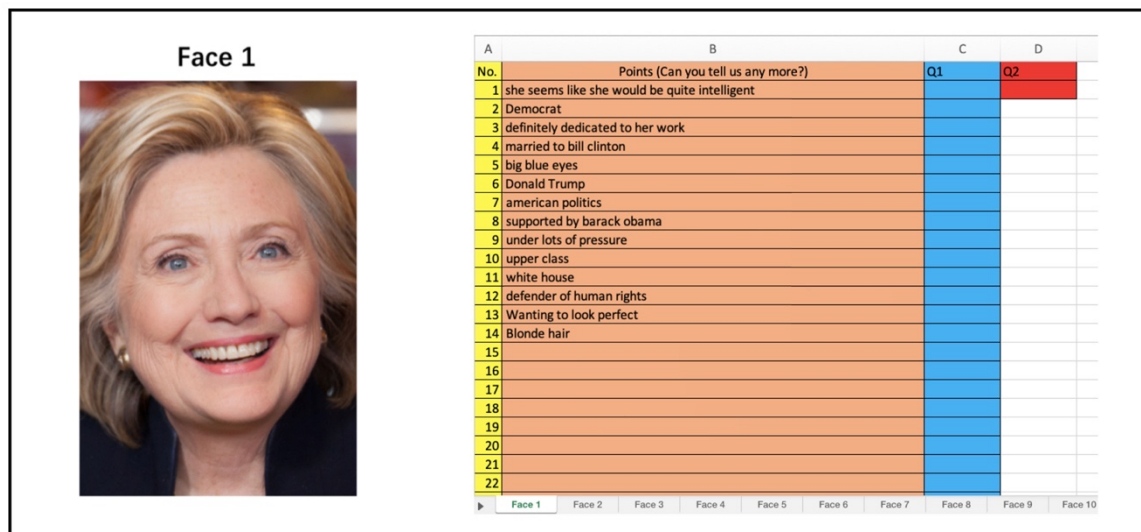
Figure 1. Example stimulus (*left*) and response sheet (*right*) from the free association task in Experiment 1. For each face, participants wrote down whatever came to mind. Example response sheet shows genuine responses from different participants for illustrative purposes.

In the metacognitive review, participants revisited their own responses from the completed face association test. First, for each point they had made, participants were asked to estimate how many other participants (out of 30) would make the same point. Participants were instructed that the point didn't have to be expressed in exactly the same words, but should express the same idea. Second, for each face they had seen, participants estimated how many points other participants would generate on average. These data allowed us to compare participants' actual performance against their estimates of peers' performance on the same association task.

The face familiarity check was a computer-based task that was used to define familiar and unfamiliar faces for each participant. Participants were presented with the 16 stimulus faces one at a time in a random order. For each face, the participant's task was to indicate whether or not the face was familiar (Yes/No response). Each face remained on screen until the

10

participant's keypress response, which initiated the next trial. Stimulus presentation and data collection were controlled by PsychoPy2 v1.82 (Peirce et al, 2019).

These three tasks gave rise to four types of data: (i) face associations—the thoughts that occurred to the participant upon seeing the face. These associations have numerosity, content, and sequential order; (ii) estimated overlap—for each point, the participant's estimate of how many other participants will make the same point; (iii) estimated numerosity—for each face, the participant's estimate of how many points participants will generate on average; (iv) the participant's own prior familiarity with each face.

2.1.4 Procedure

All participants completed the face association test, then the metacognitive review, then the face familiarity check in the same fixed order. All three tasks were self-paced, and participants could take breaks at any time. The entire test session took approximately 45 minutes to complete.

2.2 Results and discussion

2.2.1 Numerosity

Numerosity refers to the number of points participants generated in the face association task. To test for egocentric bias in estimates of other people's performance, we divided participants into three equal-sized performance groups (*Low*, *Middle*, *High*) according to their overall numerosity scores. We then compared peer estimates for these groups, separately for Familiar and Unfamiliar faces. Summary data for each condition are shown in Figure 2A.
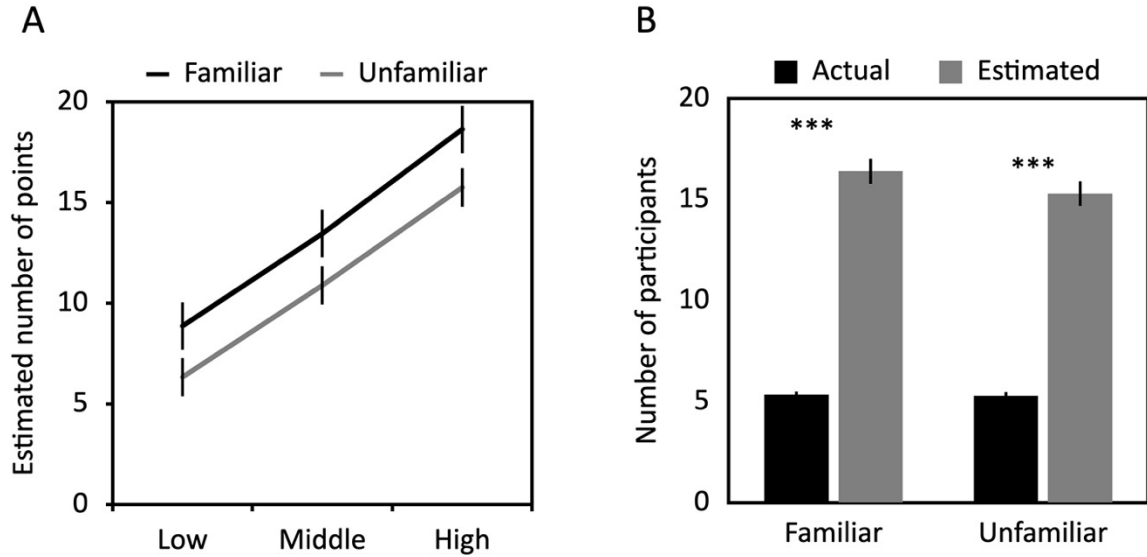
Figure 2. Egocentric bias and false consensus effects in the free association task in Experiment 1. (A) Peer estimates (condition means) from participants who themselves generated a low, middle, or high number of points, shown separately for familiar and unfamiliar faces. Peer estimates tracked participants' own performance. (B) Actual and estimated number of participants (condition means) who made the same point to the same face, shown separately for familiar and unfamiliar faces. Estimated overlap exceeded actual overlap by a factor of 3. Asterisks indicate $p < .001$. Error bars show SE.

Participants' estimates were submitted to a 2×3 mixed ANOVA with the within-subjects factor of *Familiarity* (*Familiar*, *Unfamiliar*) and the between-subjects factor of *Group* (*Low*, *Middle*, *High*). This analysis revealed a main effect of *Familiarity*, with higher estimates for *Familiar* faces ($M = 13.66$, $SE = .68$) than for *Unfamiliar* faces ($M = 11.00$, $SE = .55$) overall [$F (1,27) = 30.82$, $p < .001$, $\eta p^2 = .53$]. There was also a main effect of *Group*, with estimates increasing from the *Low* group ($M = 7.61$, $SE = .99$) through the *Middle* group ($M = 12.18$, $SE = .99$) to the *High* group ($M = 17.20$, $SE = .99$) [$F (2,27) = 23.51$, $p < .001$, $\eta p^2 = .64$]. There was no significant interaction between these factors [$F (1,27) = .06$, $p = .95$, $\eta p^2 < .01$]. Participants expected others to generate more points for faces that they themselves knew, and fewer points for faces that they themselves did not know. Moreover, for familiar and unfamiliar faces alike, high scoring participants produced high peer estimates, and low

scoring participants produced low peer estimates (Low < Middle < High in Figure 2A). Both of these findings point to egocentric bias in sizing up the face-evoked thoughts of other people.

2.2.2 Content

To quantify overlap among participants, we recruited two volunteer coders who categorised the face associations by content. We first grouped the data by face, pooling over participants, to create sixteen sets of associations (i.e. one set for each face). Each coder received all sixteen sets in a random order and independently worked through each set twice, blind to the familiarity of the faces to the participants. On the first pass, coders classified each point as either physical (relating to appearance; e.g. "red hair") or non-physical (other information; e.g. "famous actor"). This classification gives us an indicator of abstraction from the visual domain. Less than 1% of points were judged to contain both physical and non-physical information. Coders' judgements were highly correlated [κ (6297) = .869, p < .001]. On the second pass, coders grouped together points that they judged to convey the same meaning (e.g. "famous actor", "Hollywood star"), and assigned the size of the group to each point within the group. This procedure gives us the number of participants who made the same point to the same face, that is, the overlap among participants' associations. Coders' judgments were again highly correlated [r (6297) = .71, p <.001]. Any point on which the coders differed was assigned the mean of the two group sizes.

Overlap refers to the number of participants who made the same point to the same face. Summary data for each condition are shown in Figure 2B. The overlap analysis was similar to the numerosity analysis. Overlap scores were submitted to a 2 × 2 repeated measures ANOVA with the within-subjects factors of *Familiarity* (*Familiar*, *Unfamiliar*) and *Measure*

(*Estimated*, *Actual*). This analysis revealed a main effect of *Familiarity*, with higher scores for *Familiar* faces ($M = 10.88$, $SE = .32$) than for *Unfamiliar* faces ($M = 10.30$, $SE = .32$) overall [$F (1, 29) = 5.96$, $p < .05$, $\eta p^2 = .17$]. There was also a significant main effect of *Measure*, with *Estimated* scores ($M = 15.86$, $SE = .58$) exceeding *Actual* scores ($M = 5.32$, $SE = .16$) overall [$F (1, 29) = 310.12$, $p < .001$, $\eta p^2 = .92$]. These main effects were qualified by a significant interaction between *Familiarity* and *Measure* [$F (1, 29) = 5.18$, $p < .05$, $\eta p^2 = .15$]. Simple main effects confirmed that *Estimated* overlap exceeded *Actual* overlap in both the *Familiar* condition [$F (1, 29) = 295.14$, $p < .001$, $\eta p^2 = .91$] and the *Unfamiliar* condition [$F (1, 29) = 244.18$, $p < .001$, $\eta p^2 = .98$]. The simple main effect of *Familiarity* was significant for *Estimated* score [$F (1, 29) = 6.39$, $p < .05$, $\eta p^2 = .18$], but not for *Actual* score [$F (1, 29) = .10$, $p = .76$, $\eta p^2 < .01$]. The main message from this analysis is that participants overestimated the degree of overlap between their own face associations and those of other people—a false consensus effect. Participants imagined that other viewers would have the same thoughts that they themselves had upon seeing a particular face. Such convergences did occur, but less often than participants expected.

2.2.3 Order

Associations come to mind sequentially. For this analysis, we assumed that order of occurrence is a proxy for association strength: strong associations come to mind early, weak associations come to mind later (Collins & Loftus, 1975). We also assumed that order of reporting is related to order of occurrence, such that the list format of participants' responses conserves this ordinal information. To ensure that our analysis remained representative of the participant group as a whole (N = 30), we excluded data beyond list position 21, where the total number of associations across participants fell below 30 (that is, below one association per participant). Figure 3A summarises the ordinal data. Spearman's correlations confirmed

that early associations were more widely held than late associations, for *Familiar* faces [*r* (19) = -.94, *p* < .001] and *Unfamiliar* faces alike [*r* (19) = -.91, *p* < .001].
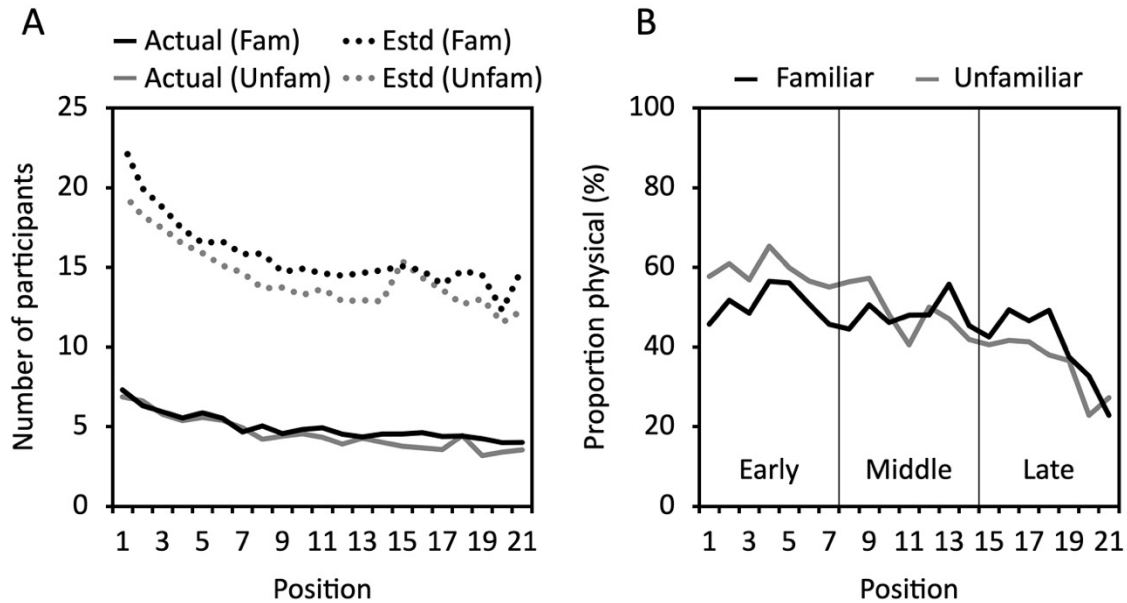

Figure 3. Ordinal analysis of consensus effects for the free association task in Experiment 1. (A) Actual and estimated (Estd) number of participants (condition means) who made the same point to the same face, plotted separately for familiar (Fam) and unfamiliar (Unfam) faces as a function of list position. Consensus was higher for earlier items than for later items. (B) Proportion of points that contained mainly physical information, plotted separately for familiar and unfamiliar faces as a function of list position.

Moreover, in keeping with egocentric bias, participants were more likely to attribute to others associations that they reported early, and less likely to attribute to others associations that they reported late (*Familiar* faces [*r* (19) = -.78, *p* < .001]; *Unfamiliar* faces [*r* (19) = -.84, *p* < .001]).

To test for qualitative differences in thoughts evoked by familiar versus unfamiliar faces, we next analysed participants' associations using the coders' classifications of content. Figure

3B shows the proportion of physical and non-physical associations as a function of list position, separately for *Familiar* and *Unfamiliar* faces.

To simplify the statistical analysis, we collapsed across list positions to form an *Order* factor with three levels—*Early* (positions 1–7), *Middle* (positions 8–14), and *Late* (positions 15–21). The proportion of points relating to physical information (facial appearance), as opposed to non-physical information (other associations), was analysed using a within-subjects ANOVA with the factors of *Familiarity* (*Familiar*, *Unfamiliar*) and *Order* (*Early*, *Middle*, *Late*). This analysis revealed no significant main effect of *Familiarity* [F (1, 6) = .77, p = .41, ηp^2 = .11]. However, there was a significant main effect of *Order*, with the highest proportion of physical points in *Early* responses (M = 54.80, SE = 1.42) followed by *Middle* responses (M = 48.55, SE = 1.43), and the lowest proportion in *Late* responses (M = 37.78, SE = 3.12) [F (1, 6) = 21.12, p < .001, ηp^2 = .78]. More importantly, there was a significant interaction between *Familiarity* and *Order* [F (1, 6) = 12.66, p < .01, ηp^2 = .68]. Simple main effects revealed a significant effect of *Order* for both *Familiar* faces [F (2, 24) = 7.16, p < .01, ηp^2 = .37] and *Unfamiliar* faces [F (2, 24) = 31.86, p < .001, ηp^2 = .73]. The simple main effect of *Familiarity* was significant for *Early* responses [F (1, 18) = 15.06, p < .01, ηp^2 = .46] and for *Late* responses (where the effect was reversed) [F (1, 18) = 4.83, p < .01, ηp^2 = .21], but not for *Middle* responses [F (1, 18) = .04, p = .89, ηp^2 < .01]. The content of face-evoked thoughts depends on the viewer's familiarity with the face. Physical information was more forthcoming for *Unfamiliar* faces than for *Familiar* faces. Conversely, non-physical information was more forthcoming for *Familiar* faces than for *Unfamiliar* faces. Thus, while the physical/non-physical categorisations reflect coders' judgements rather than an objective standard, they do appear to capture a meaningful distinction.

One feature of participants' responses that we did not anticipate was spontaneous mention of other people's names. Evidently, a presented face often brought to mind another specific individual. This occurred occasionally for *Unfamiliar* faces (10.0% of responses), but significantly more often for *Familiar* faces (30.4% of responses) [$t(29) = 3.07$, $p < .01$]. This observation suggests that person associations could be among the most salient associations evoked by faces. We return to this finding in Experiment 3.

3 Experiment 2

Experiment 1 revealed a qualitative difference between thoughts evoked by familiar and unfamiliar faces. Thoughts concerning physical appearance came to mind more readily for unfamiliar faces than for familiar faces. Thoughts concerning non-physical attributes came to mind more readily for familiar faces than for unfamiliar faces. To follow up this finding, we investigated regularities between face familiarity and face associations in an independent test. We reasoned that if familiarity affects the content of face associations, it should be possible to judge (from associations alone) whether the person who made the associations was looking at a familiar face or an unfamiliar face. Moreover, if the balance of physical content is the basis for such judgements, then sorting associations by physical content should be equivalent to sorting them by familiarity, such that the two sorting tasks give rise to similar outcomes.

To test this possibility, we administered two separate categorisation tasks in which sorters reviewed participants' response sheets from Experiment 1. In one task, sorters judged whether each sheet contained mainly *Physical* information or mainly *Non-Physical* information (focus sort). In the other task, sorters judged whether the viewer was looking at a *Familiar* or *Unfamiliar* face (familiarity sort). We expected that response sheets that were categorised as *Physical* in the focus sort would be categorised as *Unfamiliar* in the familiarity

sort. Conversely, response sheets that were categorised as *Non-Physical* in the focus sort should be categorised as *Familiar* in the familiarity sort.

3.1 Methods

3.1.1 Participants

Eighteen UK students (14 female, 4 male; mean age: 19 years; age range 18–25 years) from the University of York took part in exchange for course credit. None had participated in Experiment 1.

3.1.2 Stimuli and apparatus

The stimuli in this experiment were participants' response sheets from Experiment 1. Each of the 30 participants in Experiment 1 viewed 16 faces, resulting in a total of 480 response sheets. Screenshots of these response sheets were used as stimuli in the computer-based sorting tasks. Each image captured all of the associations that a single participant had made for a single face (see Figure 1). The same set of 480 images was used in the familiarity sorting task and the focus sorting task.

3.1.3 Design

To avoid carry-over effects, we randomly assigned the participants to one of two sorting tasks. Nine sorters categorised the response sheets by familiarity, and nine categorised them by focus. In the familiarity sort, their task was to judge, from the written associations on each sheet, whether the respondent was viewing a *Familiar* face or an *Unfamiliar* face. In the focus sort, their task was to judge whether the associations contained mainly *Physical* information or mainly *Non-Physical* information. These tasks allowed us to assign to each sheet two independent scores: (i) the number of times (out of nine) it was categorised as

Unfamiliar (vs *Familiar*), and (ii) the number of times (out of nine) it was categorised as *Physical* (vs *Non-Physical*). Each participant saw all 480 response sheets in a random order.

3.1.4 Procedure

Participants received instructions for either the familiarity sort or the focus sort before completing the prescribed task. Each trial consisted of a single response sheet presented at the centre of the screen until response. Participants categorised each sheet via keypress (Q or P), which immediately triggered the next trial. The categorisation task was self-paced, and participants could take breaks at any time. The entire test session took approximately 30 minutes to complete.

3.2 Results and discussion

For each item in the categorisation tasks (i.e. each response sheet), we recorded the number of times it was categorised as *Unfamiliar* (resulting in a familiarity score 0–9) and the number of times it was categorised as *Physical* (resulting in a focus score 0–9). Figure 4 shows how many items received each combination of scores. As can be seen from the figure, familiarity scores and focus scores were strongly correlated [r (478) = .74, p < .001]. Specifically, *Unfamiliar* judgements cleaved with *Physical* judgements, and *Familiar* judgements cleaved with *Non-Physical* judgements. These regularities suggest that abstraction away from facial appearance is taken as evidence of familiarity.
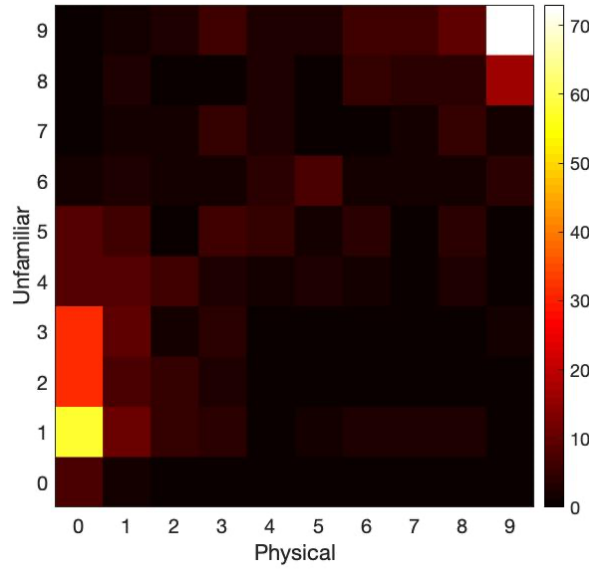
Figure 4. Analysis of face-evoked thoughts in Experiment 2. Responses that were deemed to be physical in content were deemed to refer to an unfamiliar face; responses that were deemed to be non-physical in content were deemed to refer to a familiar face. Colours indicate frequency.

To gauge the accuracy of these inferences, we next compared sorters' categorisations of familiarity (based on their reading the response sheets), to participants' actual familiarity with the faces concerned (familiarity checks in Experiment 1). The overall accuracy rate was 62%, significantly higher than chance performance of 50% [$z = 16.02$, $p < .001$]. With moderate accuracy, we can decode a person's familiarity with a face by reading what occurred to them when they saw it. In the final experiment, we focus on person associations evoked by faces, that is, connections between nodes in social networks.

4 Experiment 3

In Experiment 1, common ground among participants was smaller than participants expected. In some respects, the lack of common ground may not be surprising. After all, the face association task was entirely unconstrained, and we would expect some thoughts, such as episodic memories, or one's like or dislike of the depicted individual, to be idiosyncratic.

However, not all face-evoked thoughts are idiosyncratic in this sense. One interesting aspect of Experiment 1 was participants' inclusion of personal names as associations with the seen face. Rather often, looking at the face led the viewer to think of someone else. This observation suggests that social associations are among the most salient thoughts that come to mind when viewing a face.

Semantic priming effects attest to the strength of such social associations. Viewers are typically faster to identify a known face when it is immediately preceded by a related person than when it is preceded by an unrelated person (Young et al., 1994). This phenomenon indicates that people who co-occur or share semantic information become closely associated in memory, such that activating the representation of one person activates the representation of related people (Burton et al., 1990; Wiese & Schweinberger, 2015). We suggest that this rings true when reflecting on everyday social encounters. Utterances such as, "Have you seen Mary?" or "How are the kids?" are common in conversation.

For present purposes, the key distinction is that social networks of co-occurrence and shared semantics are not idiosyncratic; they are objective features of the world. Although participants, by definition, do not share the same idiosyncratic associations, they do inhabit the same external world, albeit a particular corner of that world. This basic contrast raises the question of whether the pattern seen in Experiment 1 (overestimation of common ground) persists even when the association task is constrained to external connections between individuals.

If the observed pattern does persist, it would imply a more thoroughgoing egocentric bias: viewers wrongly assume that facts about social networks that occur to them also occur to

others. Recognising that one's own perspective on the world is limited requires a leap of metacognitive insight. Recognising that another person's perspective will be different requires a further leap; and the process can fail at either stage. If the pattern seen in Experiment 1 is eliminated, this would imply a more limited egocentric bias: participants wrongly assume that face-evoked opinions that occur to them also occur to others, but they do not make the same mistake about face-evoked facts. Such a finding would suggest that a common frame of reference (the external world) allows better calibration of peer estimates.

To distinguish between these possibilities, we ran a modified version of the face association test in which associations were constrained to social relations, that is, connections between nodes in social networks. Instead of asking *what* comes to mind when viewing a face, we asked *who* comes to mind when viewing a face.

4.1 Methods

4.1.1 Participants

Thirty UK students (28 female, 2 male; mean age: 19 years; age range 18–21 years) from the University of York took part in exchange for a small payment and task-related bonus. None had participated in the preceding experiments.

4.1.2 Stimuli and apparatus

The stimuli and response booklets were the same as in Experiment 1, except that the instructions were modified to reflect the change in task.

4.1.3 Design

The study design was the same as in Experiment 1, except for the following changes. The face association test now called for associated people specifically, rather than any thoughts that came to mind. Accordingly, the header on each worksheet was modified to read, "Names (Can you tell us any more?)" (see Figure 5). Participants were instructed not to write the name of the person whose face was presented. Given that names can be difficult to recall, we accepted individuating semantic descriptions (e.g. "he played Harry Potter's friend") in cases where the name could not be retrieved or was never known. For simplicity, we include such entries as names in the rest of this paper. To discourage spurious responses, participants were also asked to supply a reason that the named individual was associated with the presented face.

Figure 5. Example stimulus (*left*) and response sheet (*right*) from the person association task in Experiment 3. For each face, participants wrote down whoever came to mind. Example response sheet shows genuine responses from different participants for illustrative purposes.

In the metacognitive review, participants were now asked to estimate, for each name, how many other participants (out of 30) would mention the same person, and for each face, how

many names other participants would generate on average. These data allowed us to compare

participants' actual performance against their estimates of peers' performance on the same

association task.

4.1.4 Procedure

The procedure was the same as in Experiment 1. All participants completed the face

association test, the metacognitive review, and the face familiarity check in that order.

4.2 Results and discussion

4.2.1 Numerosity

To test for egocentric bias, we again divided participants into three performance groups (*Low*,

Medium, *High*) according to their overall numerosity scores. We then compared peer

estimates for these groups, separately for *Familiar* and *Unfamiliar* faces. Summary data for
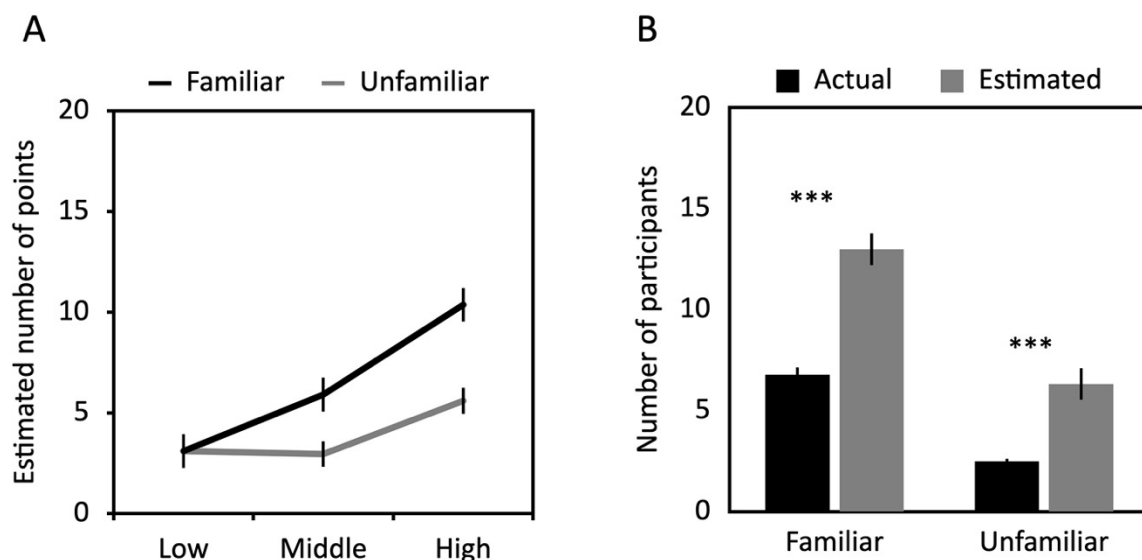
each condition are shown in Figure 6A.

Figure 6. Egocentric bias and false consensus effects in the person association task in Experiment 3.
(A) Peer estimates (condition means) from participants who themselves generated a low, middle, or

high number of names, shown separately for familiar and unfamiliar faces. Peer estimates tracked participants' own performance. (B) Actual and estimated number of participants (condition means) who mentioned the same name to the same face, shown separately for familiar and unfamiliar faces. Estimated overlap exceeded actual overlap by at least a factor of 2. Asterisks indicate $p < .001$. Error bars show SE.

The analysis took the same form as in Experiment 1. Participants' estimates were submitted to a 2 × 3 mixed ANOVA with the within-subjects factors of *Familiarity* (*Familiar*, *Unfamiliar*) and the between-subjects factor of *Group* (*Low*, *Middle*, *High*). This analysis revealed a main effect of *Familiarity*, with higher estimates for *Familiar* faces ($M = 6.47$, *SE* = .48) than for *Unfamiliar* faces ($M = 3.89$, *SE* = .37) overall [$F (1,27) = 35.69$, $p < .001$, $\eta p^2 = .57$]. There was also a main effect of *Group*, with estimates increasing from the *Low* group ($M = 3.11$, *SE* = .65) through the *Middle* group ($M = 4.44$, *SE* = .65) to the *High* group ($M = 7.99$, *SE* = .65) [$F (2,27) = 15.28$, $p < .001$, $\eta p^2 = .53$]. These main effects were qualified by a significant *Familiarity* × *Group* interaction [$F (2,27) = 10.40$, $p < .001$, $\eta p^2 = .44$]. Simple main effects showed that *Familiar* estimates exceeded *Unfamiliar* estimates in the *Middle* group [$F (1,27) = 15.66$, $p < .001$, $\eta p^2 = .37$] and the *High* group [$F (1,27) = 40.83$, $p < .001$, $\eta p^2 = .60$], but not in the *Low* group [$F (1,27) < .01$, $p = 1.00$, $\eta p^2 < .001$]. The simple main effect of *Group* was significant for both *Familiar* faces [$F (2,27) = 19.15$, $p < .001$, $\eta p^2 = .59$] and *Unfamiliar* faces [$F (2,27) = 5.38$, $p < .05$, $\eta p^2 = .29$]. Overall, participants expected others to generate more names for faces that they themselves knew, and fewer names for faces that they themselves did not know. In addition, high scoring participants produced high peer estimates, and low scoring participants produced low peer estimates. The overall pattern is again indicative of egocentric bias, this time in estimating how many people will occur to other viewers when they see a particular face.

4.2.2 Content

In this analysis of content, overlap refers to the number of participants who mentioned the same name to the same face. Given that names are so constrained, matching responses were unambiguous. As such, the coding step in Experiment 1 was not necessary here. Summary data for each condition are shown in Figure 6B. As with the numerosity scores, overlap scores were submitted to a 2 × 2 repeated measures ANOVA with the within-subjects factors of *Familiarity* (*Familiar*, *Unfamiliar*) and *Measure* (*Estimated*, *Actual*). This analysis revealed a main effect of *Familiarity*, with higher scores for *Familiar* faces ($M = 9.89$, $SE = .51$) than for *Unfamiliar* faces ($M = 4.41$, $SE = .41$) overall [$F (1, 29) = 115.58$, $p < .001$, $\eta p^2 = .80$]. There was also a significant main effect of *Measure*, with *Estimated* scores ($M = 9.65$, $SE = .67$) exceeding *Actual* scores ($M = 4.64$, $SE = .21$) overall [$F (1, 29) = 62.75$, $p < .001$, $\eta p^2 = .68$]. These main effects were qualified by a significant interaction between *Familiarity* and *Measure* [$F (1, 29) = 10.65$, $p < .01$, $\eta p^2 = .27$].

Simple main effects confirmed that *Estimated* overlap exceeded *Actual* overlap in both the *Familiar* condition [$F (1, 29) = 80.73$, $p < .001$, $\eta p^2 = .74$] and the *Unfamiliar* condition [$F (1, 29) = 24.88$, $p < .001$, $\eta p^2 = .46$]. The simple main effect of *Familiarity* was significant for both *Estimated* score [$F (1, 29) = 66.81$, $p < .001$, $\eta p^2 = .70$] and *Actual* score [$F (1, 29) = 156.00$, $p < .001$, $\eta p^2 = .84$]. Once again, participants overestimated the degree of overlap between their own associations and those of other participants, this time, for social associations specifically. Participants imagined that other viewers would think about the same people that they themselves thought about upon seeing a particular face. In fact, the overlap was smaller than they expected.

4.2.3 Order

As with the free associations in Experiment 1, we analysed the order in which social associations were generated as a proxy for association strength. To ensure that our analysis remained representative of the participant group as a whole (N = 30), we excluded data beyond list position 5, where the total number of associations across participants fell below 30 (that is, below one per participant). Figure 7 summarises the resulting ordinal pattern.
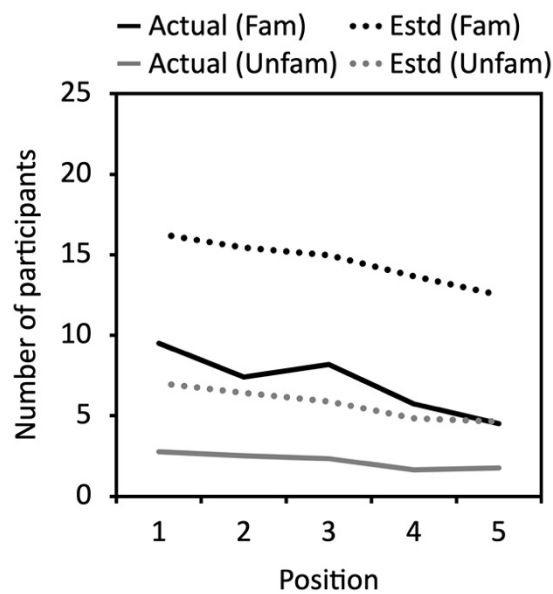

Figure 7. Ordinal analysis of consensus effects for the person association task in Experiment 3. Actual and estimated (Estd) number of participants (condition means) who mentioned the same name to the same face, plotted separately for familiar (Fam) and unfamiliar (Unfam) faces as a function of list position. Consensus was higher for earlier items than for later items.

Spearman's correlations confirmed that earlier associations were more widely held than later associations, for *Familiar* faces [$r(3) = -.90$, $p < .05$] and *Unfamiliar* faces alike [$r(3) = -.90$, $p < .05$]. As expected, participants were also more likely to attribute to others associations that came to mind early, and less likely to attribute to others associations that came to mind late (*Familiar* faces [$r(3) = -.99$, $p < .001$]; *Unfamiliar* faces [$r(3) = -.99$, $p < .001$]).

5 General Discussion

We set out to quantify overlap in face-evoked thoughts, and to compare the actual overlap with participants' expectations. This comparison revealed a consistent egocentric bias: across experiments, viewers overestimated the extent to which other people's thoughts resembled their own. In this respect, we are more isolated than we think.

These findings expand on previous work in a number of ways. First, they take an expansive view of face processing that runs from early perceptual stages through to late associative stages. In so doing, they shed new light on differences between familiar and unfamiliar face processing, contrasting the thoughts that they elicit in the viewer. Second, they encompass cognitive and metacognitive measures from the same participants. This approach illuminates the same processes from two different perspectives, and extends egocentric bias and false-consensus effects to a new area of social cognition.

Our cognitive measures conformed to our initial assumptions, providing a secure basis for comparison. For both free associations (Experiment 1) and person associations (Experiment 3), we observed large individual differences in the number of thoughts that participants recorded. In addition, points that participants mentioned early were more likely to be mentioned by others. These quantitative differences were accompanied by qualitative differences in content. Physical information was especially forthcoming for unfamiliar faces, and non-physical information was especially forthcoming for familiar faces (Experiment 1). Naive observers were apparently sensitive to these regularities. In Experiment 2, sorters who categorised viewers' responses according to inferred familiarity with the face, and sorters who categorised the same responses according to their focus on physical information, produced similar solutions.

Our metacognitive measures revealed egocentric bias in both numerosity and content of face-evoked thoughts. For numerosity, participants who generated many responses expected other viewers to generate many responses, and vice versa. This egocentric bias tracked not only individual differences in participants' response rates, but also their familiarity with individual faces (Ritchie et al., 2015). For content, participants overestimated the number of viewers whose face-evoked thoughts matched their own—and by a large margin (cf. Bui, 2012). Peer estimates exceeded actual counts by a factor of 3 in Experiment 1, and by at least a factor of 2 in Experiment 3. These false-consensus effects gave rise to especially high estimates for thoughts that participants reported first.

All of these patterns were evident in a free association task that emphasised the seen person, corresponding to an individual node in a social network (Experiment 1). They were also evident in a person association task that emphasised related people, corresponding to connections between nodes in a social network (Experiment 3). Egocentric bias and false-consensus effects at both levels indicate that we overestimate common ground in face-evoked thoughts.

So far, we have discussed these findings in relative terms—estimated overlap exceeded actual overlap. However, it is also striking how small actual overlap was in absolute terms. Any given point was mentioned by only around 20% of participants (5 out of 30) on average, and even those associations that were reported first did not command a majority. Moreover, 10% of free associations in Experiment 1 and 40% of person associations in Experiment 3 were unique, being generated by only participant. At the time of writing, there is much speculation about social media and the segmentation of society into bubbles of like-minded people

(Nikolov, Oliveira, Flammini, & Menczer, 2015; Nguyen, 2020). For this particular aspect of social cognition (face-evoked thoughts), false consensus appears so at odds with true consensus, it threatens to condemn each of us to a bubble of one. It may seem obvious that communication allows us to escape this fate. After all, the whole purpose of communication is to improve insight into the minds of others (Fernbach & Sloman, 2017). However, cognitive biases are often deeply engrained and difficult to change (Kahneman, 2011). It is worth noting that our participants each brought to the experiment 20 years of experience in social cognition. Evidently, this everyday experience was not enough to quash egocentric bias in the social cognition tasks presented here. There is some evidence that egocentric bias diminishes with age (Yinon, Mayraz, & Fox, 1994; Hayashi & Nishikawa, 2019). Future experiments could test whether older adults become better attuned to those around them following their additional exposure. Indeed, the idiosyncratic nature of face-evoked thoughts suggests that an individual differences approach will be especially fruitful. For example, the degree of focus on physical versus non-physical information may vary with the demographic background of the viewer or the seen face (Xie, Flake, & Hehman, 2019); and expectations about other people's focus may be subject to similar influences (Zhou, Burton, & Jenkins, 2021).

One way in which individuals differ is in the faces that they know. In the current experiments, all of the faces were faces of public figures, and half of them were public figures from outside of UK mainstream media. Thus, while participants expected others to say more about famous faces than about non-famous faces, their view of which ones were famous could only be based on whether or not they themselves recognised those faces. An interesting extension of this familiarity manipulation would be to include the faces of personally familiar people such as friends and family members in the stimulus set (Wiese et

al., 2019). We are confident that participants would grasp that their own friends and family members are not widely known to others. What is more interesting is whether this insight would rein in the egocentric bias seen here.

The observed false consensus effects for face-evoked thoughts raise some interesting questions about own-face processing. People tend to be highly selective about photos of themselves (Hancock & Toma, 2009). Recent studies of photographic likeness and face identification suggest the operation of egocentric bias in selection of own-face photos. Specifically, skewed representations of self interfere with our ability to judge which photographs faithfully capture our own facial appearance (White, Burton, & Kemp, 2016; White, Sutherland, & Burton, 2017). Such findings concern primarily perceptual aspects of own-face processing. In light of the current findings, it would be interesting to test whether false consensus effects also emerge in associative aspects of own-face processing. If so, then people should expect their own face to evoke in others the same thoughts it evokes in them (Gilovich, Medvec, & Savitsky, 2000).

Although we focus on face-evoked thoughts in these experiments, we do not claim that our findings are face specific. Indeed, a promising avenue for future research would be to expand the repertoire of stimulus types presented to participants. To return to our example from the introduction, the printed name "Arnold Schwarzenegger" provides access to stored semantic information without presenting a facial image. Different images of Arnold Schwarzenegger—perhaps from the *Terminator* movie or his Gubernatorial term—may evoke distinct thoughts, even though they show the same person.

Author contributions

Xingchen Zhou: Conceptualization, Methodology, Formal analysis, Data curation, Visualisation, Writing – original draft preparation. Rob Jenkins: Conceptualization, Methodology, Supervision, Writing – Reviewing and Editing.

Acknowledgements

Appendix A. Name list

Familiar celebrities. Andrew Lincoln, Avril Lavigne, David Beckham, Hillary Clinton, Mark Zuckerberg, Rupert Grint, Taylor Swift, Theresa May.

Unfamiliar celebrities. Alexander Becht, Daniele Pecci, Helen Dalley, Mathias Lauridsen, Mélanie Laurent, Minka Kelly, Stuart Bellamy, Yana Marinova.

References

Brewer, N., & Wells, G. L. (2011). Eyewitness identification. *Current Directions in Psychological Science, 20*(1), 24–27.

Bruce, V., Burton, A. M., Hanna, E., Healey, P., Mason, O., Coombes, A., Fright, R. & Linney, A. (1993). Sex discrimination: how do we tell the difference between male and female faces? *Perception, 22*(2), 131–152.

Bruce, V., & Young, A. (1986). Understanding face recognition. *British journal of Psychology*, *77*(3), 305–327.

Bui, N. H. (2012). False consensus in attitudes toward celebrities. *Psychology of Popular Media Culture, 1*(4), 236–243.

Burton, A. M., Bruce, V., & Dench, N. (1993). What's the difference between men and women? Evidence from facial measurement. *Perception, 22*(2), 153–176.

Burton, A. M., & Jenkins, R. (2011). Unfamiliar face perception. In A. J. Calder, G. Rhodes, M. Johnson, & J. Haxby (Eds.), *Oxford handbook of face perception* (pp. 287–306). New York: Oxford University Press.

Burton, A. M., Jenkins, R., & Schweinberger, S. R. (2011). Mental representations of familiar faces. *British Journal of Psychology, 102*(4), 943–958.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*(6), 407–428.

DeBruine, L. M. (2004). Facial resemblance increases the attractiveness of same–sex faces more than other–sex faces. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *271*(1552), 2085–2090.

De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2019). The "Small World of Words" English word association norms for over 12,000 cue words. *Behavior Research Methods, 51*(3), 987–1006.

Ellis, A. W., Young, A. W., & Flude, B. M. (1990). Repetition priming and face processing: Priming occurs within the system that responds to the identity of a face. *The Quarterly Journal of Experimental Psychology Section A*, *42*(3), 495–512.

Fernbach, P., & Sloman, S. (2017). The knowledge illusion. Penguin.

Germine, L., Russell, R., Bronstad, P. M., Blokland, G. A., Smoller, J. W., Kwok, H., Anthony, S. E., Nakayama, K., Rhodes, G., & Wilmer, J. B. (2015). Individual aesthetic preferences for faces are shaped mostly by environments, not genes. *Current Biology, 25*(20), 2684–2689.

Gilovich, T., Medvec, V. H., & Savitsky, K. (2000). The spotlight effect in social judgment: an egocentric bias in estimates of the salience of one's own actions and appearance. *Journal of Personality and Social Psychology, 78*(2), 211–222.

Hancock, P. J. B., Bruce, V., & Burton, A. M. (2000). Recognition of unfamiliar faces. *Trends in Cognitive Sciences, 4,* 330–337.

Hancock, J. T., & Toma, C. L. (2009). Putting your best face forward: The accuracy of online dating photographs. *Journal of Communication, 59*(2), 367–386.

Hay, D. C., Young, A. W., & Ellis, A. W. (1991). Routes through the face recognition system. *The Quarterly Journal of Experimental Psychology Section A, 43*(4), 761–791.

Hayashi, H., & Nishikawa, M. (2019). Egocentric bias in emotional understanding of children and adults. *Journal of Experimental Child Psychology, 185*, 224–235.

Hehman, E., Sutherland, C. A., Flake, J. K., & Slepian, M. L. (2017). The unique contributions of perceiver and target characteristics in person perception. *Journal of Personality and Social Psychology, 113*(4), 513–529.

Hinds, P. J. (1999). The curse of expertise: The effects of expertise and debiasing methods on prediction of novice performance. *Journal of Experimental Psychology: Applied, 5*(2), 205–221.

Hönekopp, J. (2006). Once more: is beauty in the eye of the beholder? Relative contributions of private and shared taste to judgments of facial attractiveness. *Journal of Experimental Psychology: Human Perception and Performance, 32*(2), 199–209.

Jenkins, R., & Burton, A. M. (2011). Stable Face Representations. *Philosophical Transactions of the Royal Society B: Biological Sciences, 366*(1571), 1671–1683.

Johnston, R. A., & Edmonds, A. J. (2009). Familiar and unfamiliar face recognition: A review. *Memory, 17,* 577–596.

Kahneman, D. (2011). Thinking, fast and slow. Macmillan.

Krueger, J., & Clement, R. W. (1994). The truly false consensus effect: an ineradicable and egocentric bias in social perception. *Journal of Personality and Social Psychology, 67*(4), 596–610.

McNeill, A., & Burton, A. M. (2002). The locus of semantic priming effects in person recognition. *The Quarterly Journal of Experimental Psychology Section A*, *55*(4), 1141–1156.

Mouchetant-Rostaing, Y., Giard, M. H., Bentin, S., Aguera, P. E., & Pernier, J. (2000). Neurophysiological correlates of face gender processing in humans. *European Journal of Neuroscience, 12*, 303–310.

Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme, 17*(2), 141–161.

Nikolov, D., Oliveira, D. F., Flammini, A., & Menczer, F. (2015). Measuring online social bubbles. *PeerJ Computer Science, 1*, e38.

Olivola, C. Y., & Todorov, A. (2010). Elected in 100 milliseconds: Appearance-based trait inferences and voting. *Journal of Nonverbal Behavior, 34*(2), 83–110.

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, 105*(32), 11087–11092.

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203.

Phillips, P. J., Yates, A. N , Hu, Y., Hahn, C. A., Noyes, E., Jackson, K., Cavazos J. G.,

    Jeckeln, G., Ranjan, R., Sankaranarayanan, S., Chen, J, Castillo, C. D., Chellappa, R.,

    White, D., & O'Toole, A. J. (2018). Face recognition accuracy of forensic examiners,

    superrecognizers, and face recognition algorithms. *Proceedings of the National*

    Academy of Sciences, 115(24), 6171–6176.

Ramon, M., Sokhn, N., & Caldara, R. (2019). Decisional space modulates visual

    categorization–Evidence from saccadic reaction times. *Cognition, 186*, 42–49.

Ritchie, K. L., Smith, F. G., Jenkins, R., Bindemann, M., White, D., & Burton, A. M. (2015).

    Viewers base estimates of face matching accuracy on their own familiarity: Explaining

    the photo-ID paradox. *Cognition*, *141*, 161–169.

Ross, L., Greene, D., & House, P. (1977). The "false consensus effect": An egocentric bias in

    social perception and attribution processes. *Journal of Experimental Social Psychology*,

    13(3), 279–301.

Schultz, R. T. (2005). Developmental deficits in social perception in autism: the role of the

    amygdala and fusiform face area. *International Journal of Developmental*

    Neuroscience, 23(2-3), 125–141.

Shapiro, S. S. (1966). Word association norms: Stability of response and chains of

    association. *Psychonomic Science*, *4*(6), 233–234.

Sutherland, C. A., Burton, N. S., Wilmer, J. B., Blokland, G. A., Germine, L., Palermo, R.,

    Collova, J. R., & Rhodes, G. (2020). Individual differences in trust evaluations are

shaped mostly by environments, not genes. *Proceedings of the National Academy of Sciences, 117*(19), 10218–10224.

Sutherland, C. A., Liu, X., Zhang, L., Chu, Y., Oldmeadow, J. A., & Young, A. W. (2018). Facial first impressions across culture: Data-driven modeling of Chinese and British perceivers' unconstrained facial impressions. *Personality and Social Psychology Bulletin, 44*(4), 521–537.

Sutherland, C. A., Rhodes, G., Burton, N. S., & Young, A. W. (2020). Do facial first impressions reflect a shared social reality? *British Journal of Psychology, 111*(2), 215–232.

Towler, A., White, D., & Kemp, R. I. (2014). Evaluating training methods for facial image comparison: The face shape strategy does not work. *Perception, 43*(2-3), 214–218.

Verosky, S. C., Porter, J., Martinez, J. E., & Todorov, A. (2018). Robust effects of affective person learning on evaluation of faces. *Journal of Personality and Social Psychology*, *114*(4), 516–528.

Wiese, H., Tüttenberg, S. C., Ingram, B. T., Chan, C. Y., Gurbuz, Z., Burton, A. M., & Young, A. W. (2019). A robust neural index of high face familiarity. *Psychological Science*, *30*(2), 261-272.

White, D., Burton, A. L., & Kemp, R. I. (2016). Not looking yourself: The cost of self-selecting photographs for identity verification. *British Journal of Psychology, 107*(2), 359–373.

White, D., Sutherland, C. A., & Burton, A. L. (2017). Choosing face: The curse of self in profile image selection. *Cognitive research: principles and implications, 2*(1), 1–9.

Wiese, H., Tüttenberg, S. C., Ingram, B. T., Chan, C. Y., Gurbuz, Z., Burton, A. M., & Young, A. W. (2019). A robust neural index of high face familiarity. *Psychological Science, 30*(2), 261–272.

Willis, J., & Todorov, A. (2006). First impressions: Making up your mind after a 100-ms exposure to a face. *Psychological Science, 17*(7), 592–598.

Wolfson, S. (2000). Students' estimates of the prevalence of drug use: Evidence for a false consensus effect. *Psychology of Addictive Behaviors, 14*(3), 295–298.

Yinon, Y., Mayraz, A., & Fox, S. (1994). Age and the false-consensus effect. *The Journal of Social Psychology, 134*(6), 717–725.

Xie, S. Y., Flake, J. K., & Hehman, E. (2019). Perceiver and target characteristics contribute to impression formation differently across race and gender. *Journal of Personality and Social Psychology, 117*(2), 364–385.

Zhou, X., & Jenkins, R. (2020). Dunning–Kruger effects in face perception. *Cognition*, *203*, 104345.

Zhou, X., Burton, A. M., & Jenkins, R. (2021). Two factors in face recognition: Whether you know the person's face and whether you share the person's race. *Perception*, 03010066211014016.