



A face in a (temporal) crowd

Catrina M. Hacker^a, Emily X. Meschke^a, Irving Biederman^{a,b,*}

^a Program in Neuroscience, University of Southern California, United States

^b Department of Psychology, University of Southern California, United States

ARTICLE INFO

Number of Reviews = 1

Keywords:

Face recognition

RSVP

Prosopagnosia

Negative detection

Face familiarity

Celebrity recognition

Unconscious recognition

ABSTRACT

Familiar objects, specified by name, can be identified with high accuracy when embedded in a rapidly presented sequence of images at rates exceeding 10 images/s. Not only can target objects be detected at such brief presentation rates, they can also be detected under high uncertainty, where their classification is defined negatively, e.g., “Not a Tool.” The identification of a familiar speaker’s voice declines precipitously when uncertainty is increased from one to a mere handful of possible speakers. Is the limitation imposed by uncertainty, i.e., the number of possible individuals, a general characteristic of processes for person individuation such that the identifiability of a familiar face would undergo a similar decline with uncertainty? Specifically, could the presence of an unnamed celebrity, thus *any* celebrity, be detected when presented in a rapid sequence of unfamiliar faces? If so, could the celebrity be identified? Despite the markedly greater physical similarity of faces compared to objects that are, say, not tools, the presence of a celebrity could be detected with moderately high accuracy (~75%) at rates exceeding 7 faces/s. False alarms were exceedingly rare as almost all the errors were misses. Detection accuracy by moderate congenital prosopagnosics was lower than controls, but still well above chance. Given the detection of the presence of a celebrity, all subjects were almost always able to identify that celebrity, providing no role for a covert familiarity signal outside of awareness.

1. Introduction

Rapid serial visual presentation (RSVP) paradigms have been extensively employed to assess the temporal limits of object recognition. In a typical version of the task, observers search for a target, specified by name. The target is visually masked in both the forwards and backwards directions by the preceding and following images and may be present in only half the sequences, imposing high perceptual and attentional demands as the observer must maintain attentional scrutiny throughout the sequence until a target is detected or the sequence terminates without a target detected. Moreover, the observer is faced with high uncertainty; not only does the observer not know if there will be a target in the sequence and, if one is present, where in the sequence it will occur, but (typically) does not know what its exact instantiation might be, e.g., the specific shape and pose of the object.

Uncertainty can be greatly increased with a “negative detection” version of the RSVP task, first studied by Intraub (1981), in which all the images in the sequence are from a common category, say “tools,” and the observer is to detect an object that is *not* a tool. The set of objects that are not tools is, essentially, infinite. Intraub reported (1981) that at a duration of 114 ms/image, accuracy dropped from 71% when the target was specified by name, e.g., a “chair”, to 35% when

specified negatively, e.g., “Not a Tool.”

Can faces be recognized in the extremely high uncertainty of a negative detection RSVP task? Most studies of face recognition performance require a same-different response to a single unfamiliar face, perhaps where the faces are rotated in depth (or translated or varied in size) to assess invariance, and they are presented either simultaneously (as in a match-to-sample task) or several seconds earlier as in the Cambridge Face Memory Test (CFMT) (Duchaine & Nakayama, 2006). Given the apparent difficulty of the CFMT, where even normal subjects find the test challenging, it is not unreasonable to hypothesize that face recognition under high uncertainty and extremely short masked exposure durations is impossible. The instantiation of such a paradigm in the present investigation presented sequences of unfamiliar faces in which the face of a familiar celebrity was present in half the sequences. The subject’s task was to detect whether a celebrity headshot was present (i.e., to find the face that is *not* that of a *non*-celebrity) and, if so, to identify the celebrity. The issue of person recognition under high uncertainty is of some interest in that the ability to identify a familiar celebrity voice declines markedly as the number of possible celebrities is increased from one to only a handful (Legge, Grosman, & Pieper, 1984; Shilowich & Biederman, 2016; Xu et al., 2015). With an unlimited set of possible familiar celebrities, voice identification is almost

* Corresponding author at: Department of Psychology and Program in Neuroscience, University of Southern California, United States.
E-mail address: bieder@usc.edu (I. Biederman).

impossible. If the RSVP identification of faces shows a similar decline (compared to, say, identification of objects), then there may be a general decline in person perception under high uncertainty that is not evident when perceiving objects.

We assessed the ability of controls, moderate prosopagnosics, and two extreme prosopagnosics, one congenital (GJ) and the other acquired (MJH), to perform negative detection RSVP tasks with objects and celebrity and non-celebrity headshots. Studies of prosopagnosia typically compare individuals who, on the basis of some standardized tests, are at the extremes; either clearly prosopagnosic or not. However, the diagnostic tests for prosopagnosia yield graded scoring with individuals intermediate on a dimension of face recognition ability. If faces can be detected in a negative RSVP task and if the task engages the same processes that are deficient in prosopagnosia, then those classified as intermediate in face recognition might be expected to perform at a level intermediate between controls and those more extreme on the tests for prosopagnosia. We designate such an intermediate group as moderate congenital prosopagnosics (mCPs) and those at the more extreme as extreme prosopagnosics (xCPs). Although the term “Developmental Prosopagnosia” has been used to designate those individuals who are deficient in face recognition but who have no history of neurological insult or detectable lesions in face selective areas as distinct from those “Acquired Prosopagnosics” whose deficiency is a likely consequence of lesion or disease, we prefer the term “Congenital Prosopagnosia” as there is no evidence that early childhood experience can lead to prosopagnosia. Indeed, twin studies show a higher correlation in face recognition ability between monozygotic than dizygotic twins, suggesting a genetic linkage (Wilmer et al., 2010).

2. Method

2.1. Participants

Fifty-four students from the University of Southern California (52 right handed, 34 female, mean age of 20.3 years, range 18–38 years) participated for course credit or monetary compensation. From this distribution, 47 students served as controls, while six subjects (age range 19–21, three female) were classified as moderate congenital prosopagnosics (mCPs) given a) their unremarkable neurological history and b) a level of performance of one standard deviation below the mean on at least four of five diagnostic tests (Table 1, raw scores in appendix). The PI20 served as a self-report measure of face recognition ability. The CFMT assessed face perception, short-term face memory, and invariance to orientation (Duchaine & Nakayama, 2006). The USC Face Perception Test (USCFPT, link: testable.org/t/3732942e7), is a minimal match-to-sample task with a triangular display of three computer-generated faces, a sample (on top), with one of the two lower faces an exact match to the sample, the other being the distractor (Biederman, Margalit, Maarek, Meschke, & Shilowich, 2017). The subject indicates by key press whether the left or right face is the exact match. The display remains in view for 5 s or until the subject responds providing a relatively pure test of face perception, with virtually no

contribution of memory. The Famous Faces Test (<http://www.face-blind.org>) and USC Celebrity Test (<http://bit.ly/2Bd2dyP>) are celebrity recognition tasks, reflecting long-term memory for faces.

Additionally, two “extreme” prosopagnosic subjects, xPros, were run. GJ, a 33-year-old male, was classified as a congenital prosopagnosic, xCP, on the basis of self-report, an interview, his survey scores, and no evidence of neurological incidents. The other was MJH, a 53-year-old male who is an acquired prosopagnosic, xAP, with bilateral lesions to OFA and FFA suffered as a result of a fall at the age of 5 (reported in Xu & Biederman, 2014). While MJH’s simultagnosia and mild object agnosia reflect more generalized perceptual deficits than those presented by congenital prosopagnosics, his inclusion provides additional assessment of the RSVP task to differentiate those with face recognition difficulties from controls. All subjects reported normal or corrected-to-normal vision. The work was carried out in accordance with the Code of Ethics of the World Medical Association (Declaration of Helsinki). Informed consent was obtained for experimentation with human subjects.

2.2. Stimuli

Stimuli were colored photographic images of either objects or faces (headshots) obtained from a Google image search. In the object sequences, all the images, but possibly one, belonged to a single category: tools, animals, modes of transportation, plants, or articles of clothing. All the images in the face sequences, but possibly one, were of non-celebrities or contained one of the 50 most familiar celebrities (half female) as rated by USC undergraduates in prior studies of voice recognition. Non-celebrity images were taken from websites with headshots of aspiring actors and business executives. In the judgment of the experimenters and several other observers, there was no discernible difference in image quality between the celebrities and the non-celebrities. The backgrounds were removed from all images and replaced with a homogeneous gray before being scaled to 800 by 800 pixels. Images of faces were cropped to show the full face and top of the shoulders of each person and images of objects were scaled to fit, approximately, within the center of a 19.3° square given the typical distance of the subject to the screen.

2.3. Design and Procedure

2.3.1. Familiarity ratings

Prior to the experimental trials, subjects rated their familiarity with the faces of 50 celebrities, listed by name, on a scale of 1–5 (unfamiliar, slightly familiar, moderately familiar, very familiar, most familiar). To assess whether providing the names of celebrities influenced their subsequent RSVP detection performance, eight additional subjects were run with their familiarity ratings made after their RSVP trials. As will be discussed later, there was no evidence of a detection benefit from rating the familiarity of the celebrities prior to their presence in the RSVP sequences.

Table 1
Mean Scores on Five Diagnostic Tests Distinguishing Controls, mCPs, and Two xPros.

Subject Classification	USC Face Perception Test ¹	Famous Faces Test	Cambridge Face Memory Test (CFMT)	USC Celebrity Test	PI20 ²
Controls	86.9% (72%–99%)	81.4% (27%–100%)	79.3% (57%–97%)	85.7% (52%–100%)	38.3 (23–72)
mCPs	71.2% (55%–89%)	38.3% (18%–65%)	56.2% (40%–65%)	50.3% (33%–77%)	65.3 (52–84)
xCP (GJ)	45%	18%	28%	25%	90
xAP (MJH)	52%	3%	38%	26%	83

¹ Chance on the USC Face Perception Test is 50%.

² All scores but those on the PI20 are percent correct. Scores on the PI20 are self-ratings with higher scores indicating greater difficulty in face recognition.




Fig. 1. Examples of the time course for both an object (A) and a face (B) trial. For sequence A, the prompt would be “Find the object that is not an item of clothing” or for sequence B, “Find the celebrity.” Note that the targets do not appear in the first or last six items in each trial to eliminate errors due to lack of attention, easier detection because of lack of masking at the end of the sequence or accidentally pressing through to the next sequence at the beginning of each sequence.

2.3.2. RSVP design and procedure

The main task consisted of 200 trials; 100 with faces and 100 with objects. Each trial consisted of a sequence of 32 images, with each sequence containing only objects or only faces. The sequences were presented at the center of a 27" iMac screen. In the face sequences, all the images (except possibly one) were of non-celebrities. Subjects were instructed to detect whether “a celebrity” was present in the sequence. In the object sequences, all the images (except possibly one) were from a single category, designated by category name on the screen prior to the presentation of the sequence. Subjects were to detect whether an object was presented that did not belong to the instructed category, e.g., Target: Not a tool (Fig. 1). Thus, for both faces and objects, participants performed negative detection: detecting either a celebrity among non-celebrities (i.e., a face that was *not* a non-celebrity) or an object that was not a member of the designated category. For both faces and objects, a target image was present on half the trials and never appeared in the first or last six images. All object sequences were shown at a rate of 83 ms/image with an equal number of trials for each object category (20 of each category for a total of 100 trials). The presentation rates for the face sequences were 133, 150 or 168 ms/image. Decisions about presentation time were based on previous results reported in Subramaniam, Biederman, and Madigan (2000) which used object images in an RSVP design. The decision to use different presentation times for objects and faces arose from a desire to minimize ceiling effects on the object trials where recognition was much easier. 32 non-celebrity images of various races and ages (16 female) were repeated throughout the face trials while the 50 target celebrity images each appeared only once. Subramaniam et al. (2000) reported that in RSVP experiments there was no memory, i.e., no improvement in detection, for non-target images that had been repeated up to 31 times prior to their becoming targets. Target object images were selected randomly and, by chance, a few target object images were presented as the target in two different sequences. Target face images were never repeated. Each subject viewed 200 sequences (100 object and 100 face), in one of four different orders with half the subjects viewing a given sequence in a forward order and the other in a backward order. Detection responses were made by depression of the space bar with the participant instructed to respond as quickly and as accurately as possible. The sequence ran through all 32 images independent of whether the subject made a keypress response before the sequence had finished.

Prior to the experimental trials, subjects performed 10 practice trials; 5 with faces and 5 with objects. For both types of stimuli, the first two practice trials presented the images for 140 ms and the final three

showed them for 105 ms to familiarize participants with the task at varied frame rates. Unlike the actual experimental trials, in the practice trials with faces, subjects were given the names and faces of the three possible target celebrities before the practice trials began.

After each trial, subjects verbally noted whether they had detected a target and named or described the target image to the experimenter who recorded the verbal report and scored it for accuracy. The scoring scheme awarded one point for accurately detecting whether or not a target was present and, when judged to be present, correctly naming or describing the target image, e.g., “actor in the Titanic”. Half points were given for answers that did not uniquely identify the target, such as “actor” or “tool”.

3. Results

3.1. Familiarity ratings

Controls were familiar with a larger proportion of the celebrities than the mCPs. On the pre-test familiarity scale with 5 being the highest rating, the average familiarity rating was 4.03 for Controls and 3.35 for mCPs, $t(52) = 2.90$, $p < 0.01$, Cohen's $d = 1.37$.

3.2. Overall detection accuracy

Mean accuracy scores of positive and negative trials (combined) for both objects and faces are shown in Fig. 2. Trials which contained a target celebrity that a subject had rated at the two lowest familiarity values, i.e., a 1 or a 2, were not included. The number of face trials included in the detection rates ranged from 85 to 100 for Controls and 74 to 96 for mCPs. Overall accuracy (hit and misses combined) was 74.6% for Controls and 68.8% for the mCPs, a highly reliable difference, $F(1,52) = 9.04$, $p < 0.01$, Cohen's $d = 1.24$. Mean A' scores for celebrity detection were 0.86 for Controls, 0.80 for mCPs, 0.38 for xAP MJH, and 0.67 for xCP GJ.

There were only minimal and non-reliable variations of accuracy in celebrity detection with image duration, $F(2,104) < 1.00$, ns , and an unreliable Group X Duration interaction with controls showing a slight but insignificant increase in accuracy with longer durations and the mCPs a slight decrease in accuracy as durations increased, $F(2,104) = 3.41$, ns .

Accuracy for both Controls and mCPs was substantially lower for positive trials (those which contained a celebrity headshot) than for negative trials (those without a celebrity). That is, misses were far more




Fig. 2. Overall accuracy (positive and negative trials combined) of Controls (square), mCPs (circle), xCP GJ (triangle), and xAP MJH (inverted triangle) as a function of image presentation duration on both object and face trials. Error bars are S.E. of mean. For the Controls, the error bars are encompassed within the data points.

Table 2
Percent correct on positive (celebrity present) and negative (celebrity absent) face trials for the different subject classifications.

Group	Mean Accuracy on Positive Trials	Mean Accuracy on Negative Trials
Controls	50.0%	95.0%
mCPs	29.0%	96.5%
xCP (GJ)	11.0%	90.9%
xAP (MJH)	4.0%	75.6%

frequent than false alarms as shown in Table 2. Accuracy on positive trials was 50.0% for Controls and 29.0% for mCPs, a reliable difference, $t(52) = 2.94, p < 0.01$ Cohen's $d = 1.91$. As shown in Fig. 3, a high proportion of these misses were for celebrities that were less familiar to the participant. For Controls, faces given a familiarity rating of a “3” were detected on 25.0% of their occurrences; faces rated a “4” or a “5” were detected on 47.8% and 55.5% of their occurrences, respectively. Controls and mCPs were highly accurate on negative trials, with mean scores of 95.0% and 96.5%, respectively. False alarms were thus infrequent and did not reliably differ between the two groups, $t(52) < 1.0, ns$, Cohen's $d = 0.386$.

3.3. xPros

GJ was familiar with 33 of the 50 celebrities, rating them with a mean of 4.18, while MJH was familiar with 41 celebrities and gave them a mean rating of 4.60. xCP GJ scored 60.8% ($A' = 0.67$) which was not reliably above chance, $z_{prop} < 1.00, ns$ and xAP MJH scored 46.7% ($A' = 0.38$) which was below chance (50%).

The false positive rate for xCP GJ was 9.1%, which was not significantly different from that of the Controls, Crawford's $t(46) < 1.00, ns$. xAP MJH had a false positive rate of 24.4%, which was significantly higher than Controls, Crawford's $t(46) = 4.24, p < 0.001$, Cohen's $d = 6.07$. All subjects were told that a target would be present on half the trials. As a participant in many prior experiments, MJH was keenly aware of his deficit and it is possible that he adopted a more liberal criterion even when he had not detected a celebrity knowing that there was a high probability that one was present.

GJ accurately identified that a target was present on only four of the 33 trials which contained a celebrity headshot with whom he was familiar. After the experiment, he reported that he was able to make some of these identifications not because he recognized the face, but because he “recognized the picture.” For example, he noted the use of Obama's iconic presidential portrait. His apparent decline with longer presentation durations appeared to be a consequence of those instances of




Fig. 3. Percent correct detection on positive trials only of a celebrity face among non-celebrity foils as a function of rated familiarity of the celebrity's face for mCPs and Controls. Error bars are S.E. of mean.

a familiar “picture” tending to occur at the shorter presentation durations.

3.4. Object trials

All subjects, save xAP MJH, performed well on the object trials with overall accuracy rates of 89.1% (range 74%–98%) for Controls ($A' = 0.94$), 86.17% (range 67%–95%) for mCPs ($A' = 0.93$) and 83.50% for xCP GJ ($A' = 0.93$). Differences between Controls and mCPs were non-significant, $t(52) = 1.08, ns$, Cohen's $d = 0.37$ and GJ's score was not significantly different from mCPs, Crawford's $t(6) < 1.0, ns$. xAP MJH scored 58.00% on the object trials ($A' = 0.74$), which was significantly worse than mCPs, Crawford's $t(6) = -2.6, p < 0.05$, Cohen's $d = 2.80$. This is consistent with his more general perceptual deficit which extends to objects.

3.5. Celebrity Identification

Fig. 3 shows the accuracy of the detection of a celebrity's face on positive trials, i.e., trials that contained a celebrity headshot, as a function of the pre-experimental rated familiarity of that face. For both Controls and mCPs accuracy increased monotonically with higher familiarity ratings. At every familiarity level, detection accuracy of Controls exceeded that of mCPs.

For both groups, detection accuracy hovered near 0% at the lowest familiarity levels (Fig. 3) suggesting that there was no extraneous photographic quality or feature in the images that might have distinguished celebrity from non-celebrity headshots. The absence of an extraneous cue distinguishing celebrity from non-celebrity headshots is also supported by the near chance responding in the face sequences by both extreme prosopagnosics.

Fig. 4 shows that when a Control or an mCP made a detection response on a positive trial, they were almost always able to identify who that celebrity was. This near-ceiling identification accuracy following celebrity detection leaves only minimal opportunity for the employment of an outside-of-awareness familiarity signal, as has been indexed by SCRs and EEG in acquired prosopagnosic subjects when viewing faces (Renault, Signoret, Debruille, Breton, & Bolgert, 1989; Tranel & Damasio, 1985). Correct detection was accompanied by correct identification on 97.5% and 100% of the object trials and 96.4% and 100% of the face trials, for Controls and mCPs respectively (Fig. 4). The main effect of both subject group, $F(1, 105) = 1.92, p > 0.5$, and stimulus type, $F(1, 105) = 1.63, p > 0.5$, fell short of significance.

This doesn't mean that an unconscious familiarity signal does not exist, merely that in the present study, conscious identification almost always accompanied detection leaving little opportunity for an




Fig. 4. Percent of trials in which correct detection of the presence of a target (celebrity or object) was followed by a correct identification of that target. Error bars are S.E. of mean.

unconscious process to have an effect on detection or identification. Although the Tranel & Damasio and Ranault et al. studies did not employ a 2AFC (2 Alternative Forced Choice Task) which would have provided a more sensitive behavioral measure of whether their prosopagnosics could consciously distinguish familiar from unfamiliar faces, other studies with CPs (e.g., Avidan & Behrmann, 2008; Rivalta, Palermo, Schmalzl, & Coltheart, 2012) have reported that unrecognized familiar faces can, nonetheless, be processed more efficiently than unfamiliar faces.

To assess whether judging the familiarity of the faces of the celebrities provided prior to the experimental trials facilitated their detection in the RSVP sequences, a separate group of eight subjects rated the familiarity of the celebrities' faces after their RSVP trials. The average detection accuracy of these subjects was 77.7% which is comparable to the 74.6% accuracy of the subjects rating the celebrities prior to the experimental trials. Moreover, celebrity detection was accompanied by correct identification on 97.0% of the trials in these additional subjects (vs. 97.5% of the trials for subjects who rated the celebrities prior to the RSVP trials), indicating that having the familiarity ratings at the beginning of the task did not inflate target detection or identification.

4. Discussion

Given the conditions of enormous uncertainty in the present experiment as to what the target image might be on a given trial ("any celebrity" allowing at least a thousand individuals) with drastic limitations on time and attentional capacity in detecting an image masked by preceding and subsequent images, it was not at all clear that any subjects would be able to detect an unspecified celebrity face when embedded in an RSVP sequence of highly similar headshots of unfamiliar individuals. The presence of a familiar celebrity could not only be detected in such presentations but was almost always identified when detected. The vast majority of the errors were misses on positive trials. Rarely did control or moderate congenitally prosopagnosic subjects false alarm on negative trials when a celebrity's headshot was not in the sequence.

This ability to detect and recognize brief, masked presentations of familiar faces from a set size likely in the thousands presents a stark contrast to the effect of the number of possible individuals in voice recognition. Three experiments have shown that increasing the possible number of voices beyond a handful, either in celebrity identification (Shilowich & Biederman, 2016; Xu et al., 2015) or old-new recognition

of newly learned voices (Legge et al., 1984), leads to a marked decrease in accuracy.

Compared to Control subjects, mild prosopagnosics were less accurate than Controls but were well above chance. The two extreme prosopagnosics did not differ reliably from chance. The ordering of the subject groups, Controls, mCPs, and xPros, on the basis of celebrity face detection accuracy documents graded perceptual performance predictable on the basis of non-perceptual assessments of face recognition ability such as the PI20. The lower detection accuracy of the mCPs and xPros speaks to the relevance of the RSVP task to general face recognition processes. That the extreme prosopagnosics were at chance and the near zero detection rate for faces of low familiarity also suggest that there was not some inadvertent non-face cue that might have signaled the presence of a celebrity.

Given the much greater physical similarity of target celebrity faces to the non-celebrity foils, it is not surprising that, despite the shorter presentation duration of 83 ms, the negative detection accuracy of objects was markedly higher than the face detection rate.

4.1. Differences in celebrity familiarity

An inadvertent finding is that the mCPs, even though they were from the same undergraduate population as the controls, judged themselves to be less familiar with the faces of celebrities than the Controls. The mCPs' familiarity ratings were, by and large, as predictive of detection accuracy as the Controls although their detection accuracy was lower than the controls at every level of rated familiarity. For both Controls and mCPs, detection accuracy monotonically increased with ratings. At this point, we can only speculate on the cause of this difference in familiarity with celebrity faces. It is not rare for prosopagnosics to remark that they have difficulty in distinguishing characters in a movie or a TV production. Indeed, item #14 of the PI20 is, "I sometimes find movies hard to follow because of difficulties recognizing characters." Given that a sizable proportion of the celebrities were entertainers it is perhaps not surprising that people with difficulty in individuating faces might not engage entertainment venues as frequently as others without such difficulties, causing them to be less familiar with the faces of entertainers.

4.2. Role of an unconscious familiarity signal?

Two prior studies, Tranel and Damasio (1985) and Renault et al. (1989), with two and one APs, respectively, reported that EEG and SCR (skin conduction level) measures could provide an unconscious physiological signal for familiar faces in acquired prosopagnosics. Two subsequent studies by Avidan and Behrmann (2008) and Rivalta et al. (2012) with CPs showed that a familiar face that was not recognized in an explicit test, nonetheless, could exert a facilitative effect on subsequent face processing. The participants in the present study had much higher accuracy on negative trials than on positive trials. This high criterion for a positive response indicates that subjects had no need to employ such a signal to guide detection. When CPs and Controls correctly indicated that a target face was present, nearly 100% of the time this detection was accompanied by identification of that celebrity. The basis of detection was thus not limited to a general feeling of familiarity in the absence of knowing who the celebrity was, but the actual identification of the celebrity.

5. Conclusions

Under what could be maximum uncertainty as to the set of possible familiar faces—likely in the thousands—and the very brief, masked image presentation durations with high attentional load induced by the negative detection RSVP task, the current experiment demonstrated that reasonably accurate detection and identification of a familiar face is possible. This is in stark contrast to the effect of uncertainty on voice

recognition, suggesting that the striking cost of uncertainty in voice recognition is not a general phenomenon when individuating people. As almost every detection of a celebrity on positive trials was accompanied by the accurate identification of that celebrity, no opportunity was available for the employment of an unconscious familiarity signal.

Appendix A.

Raw Scores on Five Diagnostic Tests for Six Subjects Classified as Moderate Congenital Prosopagnosics.

Subject #	USC Face Perception Test	Famous Faces Test	Cambridge Face Memory Test	USC Celebrity Test	PI20
1	82%	18%	65%	42%	56
2	55%	44%	65%	57%	52
3	70%	33%	57%	33%	73
4	70%	48%	54%	77%	84
5	89%	22%	56%	46%	68
6	61%	65%	40%	47%	59

References

Avidan, G., & Behrmann, M. (2008). Implicit familiarity processing in congenital prosopagnosia. *Journal of Neuropsychology*, 2(1), 141–164.

Biederman, I., Margalit, E., Maarek, R. S., Meschke, E. X., & Shilowich, B. S. (2017). What is the nature of the perceptual deficit in developmental prosopagnosia? *Journal of Vision*, 17(10), <http://dx.doi.org/10.1167/17.10.619> 619–619.

Duchaine, B., & Nakayama, K. (2006). The Cambridge Face Memory Test: Results for neurologically intact individuals and an investigation of its validity using inverted face stimuli and prosopagnosic participants. *Neuropsychologia*, 44, 576–585.

Intraub, H. (1981). Rapid Conceptual Identification of Sequentially Presented Pictures. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 604–610.

Legge, G. E., Grosman, C., & Pieper, C. M. (1984). Learning unfamiliar voices. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 10, 298–303.

Renault, B., Signoret, J., Debrulle, B., Breton, F., & Bolgert, F. (1989). Brain potentials reveal covert facial recognition in prosopagnosia. *Neuropsychologia*, 27, 905–912.

Rivolta, D., Palermo, R., Schmalzl, L., & Coltheart, M. (2012). Covert face recognition in

Acknowledgments

We thank Rafael Maarek and Jordan Juarez for their assistance in running subjects, writing some of the code, and preparing the stimuli. Supported by NSF BCS 0617699 and the Dornsife Research Fund.

congenital prosopagnosia: A group study. *Cortex*, 48(3), 344–352.

Shilowich, B. E., & Biederman, I. (2016). An estimate of the prevalence of developmental phonagnosia. *Brain & Language*, 159, 84–91.

Subramaniam, S., Biederman, I., & Madigan, S. A. (2000). Accurate identification but no priming and chance recognition memory for pictures in RSVP sequences. *Visual Cognition*, 7, 511–535.

Tranel, D., & Damasio, A. R. (1985). Knowledge without awareness: an automatic index of facial recognition by prosopagnosics. *Science*, 228, 1453–1455.

Wilmer, J. B., Germine, L., Chabris, C. F., Chatterjee, G., Williams, M., Loken, E., et al. (2010). Human face recognition ability is specific and highly heritable. *PNAS*, 107, 5238–5241.

Xu, X., & Biederman, I. (2014). Neural correlates of face detection. *Cerebral Cortex*, 24, 1555–1564. <http://dx.doi.org/10.1093/cercor/bht005>.

Xu, X., Biederman, I., Shilowich, B. E., Herald, S. B., Amir, O., & Allen, N. E. (2015). Developmental phonagnosia: Neural correlates and a behavioral marker. *Brain & Language*, 149, 106–117.