Why are familiar-only experiences more frequent for voices than for faces?

J. Richard Hanley
University of Essex, Colchester, U.K.
Jennifer M. Turner
University of Liverpool, Liverpool, U.K.

Hanley, Smith, and Hadfield (1998) showed that when participants were asked to recognize famous people from hearing their *voice*, there was a relatively large number of trials in which the celebrity's voice was felt to be familiar but biographical information about the person could not be retrieved. When a *face* was found familiar, however, the celebrity's occupation was significantly more likely to be recalled. This finding is consistent with the view that it is much more difficult to associate biographical information with voices than with faces. Nevertheless, recognition level was much lower for voices than for faces in Hanley et al.'s study, and participants made significantly more false alarms in the voice condition. In the present study, recognition performance in the face condition was brought down to the same level as recognition in the voice condition by presenting the faces out of focus. Under these circumstances, it proved just as difficult to recall the occupations of faces found familiar as it was to recall the occupations of voices found familiar. In other words, there was an equally large number of familiar-only responses when faces were presented out of focus as in the voice condition. It is argued that these results provide no support for the view that it is relatively difficult to associate biographical information with a person's voice. It is suggested instead that associative connections between processing units at different levels in the voice-processing system are much weaker than is the case with the corresponding units in the face-processing system. This will reduce the recall of occupations from voices even when the voice has been found familiar. A simulation was performed using the latest version of the IAC model of person recognition (Burton, Bruce, & Hancock, 1999) which demonstrated that the model can readily accommodate the pattern of results obtained in this study.

Hanley, Smith, and Hadfield (1998) have recently demonstrated that it is much more difficult to recognize a famous person from hearing their voice than from seeing their face. Participants in their study either were shown video recordings of famous people or

Requests for reprints should be sent to Prof. J. Richard Hanley, Department of Psychology, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, U.K. Email: rhanley@essex.ac.uk

The authors are indebted to Professor Mike Burton for conducting the simulations reported in this paper. They would also like to thank Professor Tim Valentine for his helpful comments on a previous draft of this paper.

http://www.tandf.co.uk/journals/pp/02724987.html

heard a recording of their voice. Hanley et al. found that fewer target voices than faces were correctly recognized, and more voices than faces of non-famous people were considered to be familiar. Such findings are consistent with research on episodic memory for previously unfamiliar faces and voices (e.g., Yarmey, Yarmey, & Yarmey, 1994), which shows that a recently seen face is much more likely to be correctly recognized than a recently heard voice. The problem with recognizing famous voices relative to famous faces is likely to be exacerbated by the fact that we probably encounter a famous person's face in the media more often than we hear their voice.

However, Hanley et al. (1998) did not merely demonstrate that familiar faces are easier to recognize than voices. Generally when we see a face that feels familiar, we are also able to retrieve information about the person concerned, such as their occupation or where they are typically encountered (Hanley & Cowell, 1988; Hay, Young, & Ellis, 1991; Young, Hay, & Ellis, 1985). For example, Hay et al. (1991) reported that if a famous face was found familiar, then participants could recall appropriate biographical information about the person on approximately 94% of occasions. Hanley et al. found a quite different pattern when they presented voices of famous people instead of their faces. Overall, the occupations of only 50–60% of voices found familiar could be recalled. There were therefore many more "familiar-only" responses to voices than to faces (i.e., situations in which the person feels familiar but no additional information about them can be retrieved). Similar results were obtained by Maylor (1997) in a study that investigated the effects of ageing on the identification of famous voices. Hanley et al. also observed a disproportionately large number of familiar-only responses even for voices that were deemed by participants to be very high in familiarity. They concluded that it is relatively difficult to associate the voice of a familiar person with biographical information about them.

One obvious extension of this line of research would be an investigation of exactly what it is about voices that appears to make them so difficult to associate with biographical information. However, as Hanley et al. (1998) acknowledged, the fact that overall performance was much lower in the voice condition than in the face condition was a cause for concern about the data they presented. Because the number of unfamiliar people that were falsely recognized was significantly higher in the voice condition, it is quite possible that the pool of familiar voices that was recognized contained a relatively large number of items to which the participants correctly responded "familiar" on the basis of noise or guesswork. It would not be surprising if the participants were unable to recall an appropriate occupation for "familiar" items of this kind.

It is therefore possible that poorer recall of occupations from voices found familiar than from faces found familiar is an artifact of the fact that the voices are associated with lower overall levels of familiarity than the faces. Hanley et al. (1998) attempted to circumvent this problem by comparing the performance of selected groups of participants from the face and voice conditions who were as closely matched as possible in terms of number of hits and false alarms. Even under these circumstances, fewer occupations of voices deemed familiar than those of faces deemed familiar were recalled. Nevertheless, post hoc analyses of this kind are far from optimal. They require that the results from a relatively large number of the participants are discarded, and they can introduce confounding variables into the experiment. It would therefore be reassuring if Hanley et al.'s findings could be replicated under conditions in which hit rate and false-alarm rate were similar in the voice and face conditions.

In the present study, therefore, we attempted to match overall levels of recognition in the face and voice conditions by reducing the level of performance in the face condition. This was achieved by presenting subjects with faces slightly out of focus. The critical question was whether it would remain easier to recall the occupations of faces found familiar than the occupations of voices found familiar.

Method

Participants

A total of 75 people, the majority of whom were undergraduate students at Liverpool University, took part in the experiment. Their ages ranged from 18 to 40 years.

Design

A between-subjects design was employed; 25 participants were randomly allocated to the voices condition, 25 were allocated to the standard-face condition, and 25 were allocated to the blurred-faces condition.

Stimuli

A videotape containing extracts from television interviews with 34 celebrities and 17 non-famous people was used to present the stimuli in this experiment. Each extract lasted between 10 and 15 s, and included the celebrity saying something neutral that gave no clues as to their identity or occupation. Each extract was presented once to each participant. The order of the extracts on the videotape was randomized. The celebrities comprised approximately equal numbers of politicians, sportsmen/women, actors, TV personalities, comedians, and popstars. The tape was a shortened version of the one used by Hanley et al. (1998).

Procedure

In the blurred-face condition, a video projector was used to present moving images of the faces onto a large white screen. Participants sat approximately 4 metres away from the screen. The focus of the slide projector was set such that all of the faces were presented out of focus. Results from a pilot study had previously revealed a similar proportion of correct familiarity decisions in this condition as in the voice condition in Hanley et al.'s (1998) study. The volume control was set to zero so that participants were unable to hear what the person was saying. Conditions in the standard-face condition were the same except that the faces were presented in focus. Participants in the voices condition heard a series of voices one at a time through a speaker, but were unable to see the screen on which the faces were displayed. All participants were exposed to the same extracts, thus ensuring that time exposure to each person was equated in the three conditions.

Participants in all three groups were asked first of all to indicate whether they found the face or voice unfamiliar or familiar. If the latter, they were asked to rate the degree of familiarity on a scale of 1–3 where "3 signifies high familiarity, 2 medium familiarity, and 1 low familiarity". A rating scale containing only 3 points was used in an attempt to ensure that each participant would produce a reasonably large number of responses at each level of familiarity. They were then asked to recall the occupation and the name of the celebrity.

Results

Mean scores on the different performance measures can be seen in Table 1. The alpha level for all statistical analyses was set at .05. An analysis of variance (ANOVA) revealed a significant difference in the number of correct familiarity decisions, $F(2, 72) = 40.39$, and in the number of false positives, $F(2, 72) = 31.63$, in the three conditions. Newman-Keuls post hoc tests revealed significantly more correct familiarity decisions and significantly fewer false positives in the standard-face condition than in the blurred-face conditions, which did not differ significantly from each other. An analysis by items also revealed a significant main effect of condition on the number of correct familiarity decisions, $F(2, 66) = 33.66$. Newman-Keuls post hoc tests once again revealed significantly more correct familiarity decisions in the standard-face condition than in the blurred-face and voice conditions, which did not differ significantly from each other.

Analyses were also performed using the standard signal detection measures of d' (sensitivity) and beta (bias). There was a significant difference in d' scores in the three conditions, $F(2, 72) = 118.14$. Post hoc tests revealed that d' was significantly higher in the standard-face condition than in the blurred-face and voice conditions, which did not differ significantly from each other. This demonstrates greater ability to distinguish famous from non-famous people by participants in the standard-face condition. There was, however, no significant difference in beta scores in the three conditions ($F < 1$). This shows that there is no evidence of subjects in the voice or blurred-face conditions adopting a more or less conservative criterion than participants in the standard-face condition.

We were also interested in whether or not the voice and blurred-face conditions were matched in terms of the number of familiarity decisions that were made at each of the three levels of familiarity. The relevant data are summarized in Table 2. An ANOVA

TABLE 1
Mean performance in the face, blurred-face, and voice conditions

	Condition					
	Face		*Blurred face*		*Voice*	
	M	*SE*	*M*	*SE*	*M*	*SE*
Found familiar	29.76	0.78	20.64	0.90	19.28	1.00
False alarms	1.04	0.30	4.64	0.28	4.88	0.52
d′	2.67	0.11	0.92	0.06	0.80	0.11
beta	1.37	0.19	1.17	0.07	1.30	0.11
Occupations recalled	27.12	1.05	11.08	0.28	11.12	0.52
Occupations conditionalized on familiarity	0.91	0.02	0.55	0.03	0.56	0.05
Names recalled	20.56	1.41	8.48	0.79	7.80	1.01
Names conditionalized on familiarity	0.69	0.04	0.41	0.04	0.38	0.05

TABLE 2
The mean number of decisions on familiar items made at each level of familiarity in the face, blurred-face, and voice conditions

	Level of Familiarity					
	1		*2*		*3*	
	M	*SE*	*M*	*SE*	*M*	*SE*
Face	2.60	0.53	5.32	0.91	21.84	1.64
Blurred Face	6.72	0.52	6.44	0.64	7.48	0.70
Voice	5.00	0.68	4.92	0.60	9.36	1.20

revealed a significant main effect of level of familiarity, $F(2, 72) = 60.82$, and of condition, $F(2, 144) = 25.29$. Of particular interest was the presence of a significant interaction between these factors, $F(4, 144) = 27.80$. In order to compare directly the number of familiarity decisions, made in the blurred-face and voice conditions, this interaction was further investigated by performing a two-way ANOVA after the data from the normal face condition had been removed. This revealed no effect of condition ($F < 1$), a significant effect of level of familiarity, $F(2, 96) = 7.97$, and a significant interaction, $F(2, 96) = 3.47$. Tests of simple main effects revealed a significant effect of level of familiarity in the voice condition, $F(2, 96) = 10.95$, but not in the blurred face condition ($F < 1$). There was, however, no significant difference in the number of responses made in the blurred-face and voice conditions at any of the three levels of familiarity. The voice condition and the blurred-face condition were therefore matched for both the number and the strength of the familiarity decisions that participants made.

There were significant differences in the number of occupations recalled, $F(2, 72) = 86.45$, and in the number of names recalled, $F(2, 72) = 42.62$, in the three conditions. An analysis by items also revealed a significant main effect of condition on the number of occupations recalled, $F(2, 66) = 81.69$. Post hoc tests revealed that significantly more occupations and significantly more names were recalled in the standard-face condition than in the blurred-face and voice conditions, which did not differ significantly from each other. An analysis by items revealed no significant difference in the number of occupations or names recalled in the voice and blurred-face conditions ($F < 1$).

An analysis also investigated the number of occupations recalled conditionalized on finding the face or voice familiar. This showed a significant difference in the proportion of occupations recalled in the three conditions by subjects, $F(2, 72) = 31.88$, and by items, $F(2, 66) = 54.05$. Post hoc tests revealed that a significantly higher proportion of occupations was recalled in the face condition than in the blurred-face and voice conditions. There was no significant difference in the number of occupations recalled conditionalized on familiarity in the voice and blurred-face conditions.

Figure 1 shows the proportion of occupations recalled when a face or voice was recognized at each of the three levels of familiarity. There was a significant effect of condition on the proportion of occupations recalled, $F(2, 72) = 9.98$. Post hoc analyses revealed significantly higher performance in the normal-face condition than in the

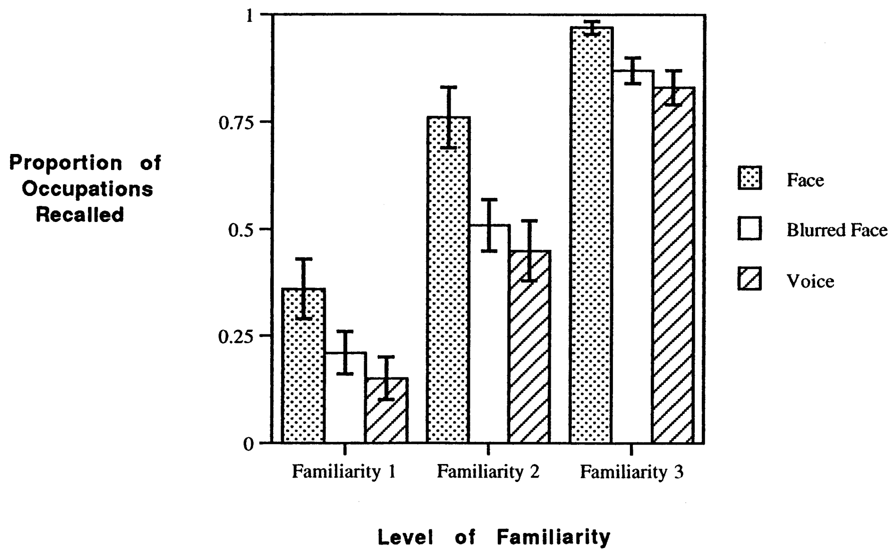

Figure 1. The proportion of occupations recalled as a function of level of familiarity in the face, blurred-face, and voice conditions (standard error bars included).

blurred-face or voice conditions, which did not differ from each other. There was also a significant effect of level of familiarity, $F(2, 144) = 160.20$. Post hoc analyses revealed that significantly more occupations were recalled following Level three than following Level two familiarity responses, and significantly more occupations recalled following Level two than following Level one familiarity responses. The interaction between condition and level of familiarity was not significant, $F(4, 144) = 1.01$.

Discussion

It is clear from Tables 1 and 2 that the manipulation of distorting faces by presenting them out of focus has proved successful in reducing performance in the blurred-face condition to a level similar to that observed in the voice condition, and well below that of the standard-face condition. Crucially, the number of hits and the number of false positives are very similar in the blurred-face and voice conditions (see Table 1), and these two groups are matched for both sensitivity (d') and criterion (beta). There are also an equivalent number of familiarity responses in the voice and blurred-face conditions at all three levels of familiarity (Table 2).

The critical question now becomes the ability to recall occupations from faces and voices found familiar. Hanley et al. (1998) found superior recall of occupations in the standard-face condition than in the voice condition, and the present study has replicated those findings. Does it, however, remain easier to recall occupations from familiar faces than from familiar voices when familiarity level is matched? The unequivocal answer to this question is that it does not. It is just as difficult to recall an occupation from a blurred face deemed familiar as it is to recall an occupation from a voice deemed familiar. Furthermore, there was no difference between the proportion of occupations recalled in the voice and blurred-face conditions at any of the three levels of familiarity (see Figure 1). It therefore appears that the reason that it is relatively easy to recall occupations from normally presented faces that are found familiar is the lower overall level of performance in the voice condition than in the standard-face condition. When steps are taken to reduce familiarity level in the face condition to a level equivalent to that in the voice condition, then voices and faces appear to behave in exactly the same way.

It is striking that equating performance in the face and voice conditions by blurring the faces led to a quite different outcome from the technique that Hanley et al. (1998) used to equate performance in the face and voice conditions. Hanley et al. compared the highest scoring subjects in the voice condition with the lowest performing subjects in the face condition, and found poorer recall of occupations to voices found familiar than to faces found familiar. We are unable to offer a definitive explanation for this discrepancy. It seems possible, however, that the best performing subjects in the voice condition may in fact have been the subjects who were the "luckiest" guessers when making familiarity decisions. Although their hit rate and false alarm rate were similar to those of the subjects in the voice condition, the pool of famous people that they found "familiar" may therefore have contained a relatively large number of people recognized on the basis of relatively weak levels of familiarity. Whether or not this account is correct, the contrasting results show how misleading it can be to equate performance in two groups by post hoc subject selection rather than by an experimental manipulation.

How easily can the results obtained in this experiment be explained in terms of the models of person recognition put forward by Bruce and Young (1986) and Burton, Bruce, & Johnston (1990)? According to Burton et al.'s interactive activation and competition (IAC) model, familiar voices and faces have their own voice recognition units (VRUs) and face recognition units (FRUs), but these units do not themselves signal familiarity. The role of these units is to pass activation on to a modality-free person identity node (PIN), which represents the point at which the face and voice recognition systems converge. It is the level of activation at the PIN that determines whether or not a person is found familiar. A similar view about the role of PINs can be found in Bredart, Valentine, Calder, and Gassi (1995). Such an account is slightly different from that put forward previously by Bruce and Young (1986). Bruce and Young argued that activation at the level of the FRU indicated familiarity of a face, and activation of a VRU indicated that a voice was familiar.

In terms of these models, one possibility is that the familiar-only responses that normal subjects make under laboratory conditions arise because of a failure of the FRU or VRU to fully activate the PIN. According to Bruce and Young (1986), this will directly reduce the number of occupations that are recalled because in their model semantic information is stored at the level of the PINs. In terms of Burton et al. (1990), a failure of the FRU or VRU to fully activate the PIN could create a familiar-only response if there was enough activation for the PIN to reach threshold and allow the person to be identified as familiar, but insufficient activation to permit recall of semantic information units (SIUs) from the semantic pool.

Attempted simulations of familiar-only responses

In an attempt to confirm that the IAC model could indeed successfully simulate the results obtained in the voice and blurred-face conditions, new computer simulations were performed using the latest version of the IAC model (Burton, Bruce, & Hancock, 1999; see pages 26–27 for technical details of the simulations). This is by far the most complex and powerful version of the IAC model to date; it can recognize and retrieve semantic information about the faces of 50 different people. The model does not include a pool of VRUs. For the purposes of the simulation of performance in the voice condition, therefore, the FRUs in the model were taken to represent the VRUs. We refer to them henceforth as VRUs.

In order to simulate familiar-only performance in the voice condition, all of the associative links between VRUs and PINs were weakened. This was achieved by setting them to 40% of the strength of the connection between the FRUs and the PINs in Burton et al.'s (1999) model. An individual VRU was then activated (input level = 1.0), and the model was allowed to cycle until it stabilized (100 cycles were employed). This process was repeated for all of the 50 people represented in the model. Table 3 shows the final resting level of activation at the VRUs, PINs, and SIUs for two people whose activation levels are typical of the people represented in the model. Activation levels at the two most strongly activated SIUs (SIU1 & SIU2) and the two most weakly activated SIUs in the model (SIU3 & SIU4) are presented. The top line is a simulation of performance in the standard-face condition, and shows the final activation levels with FRU–PIN connection strengths the same as those in Burton et al. (1999).

It can be seen that although activation levels at the VRUs are similar to activation levels at the FRUs in the standard-face condition, activation at the PINs is lower in the voice than in the face condition. Nevertheless, the PIN threshold set by Burton et al. (1999) of 450 is consistently exceeded, which means that both Person 1 and Person 2 would be deemed familiar in the voice as well as in the face condition. The IAC model has never set a threshold level for SIUs, and instead of establishing an arbitrary level here, we believe

TABLE 3
The final activation levels of units in Burton et al.'s (1999) IAC model after 100 processing cycles in the standard face condition,[a] voice condition,[b] and blurred-face condition[c]

		FRU/VRU	*PIN*	*SIU1*	*SIU2*	*SIU3*	*SIU4*
Person 1	Standard face	747	508	359	319	129	97
	Voice	744	466	311	274	118	100
	Blurred face	516	526	301	256	150	132
Person 2	Standard face	612	494	370	318	127	97
	Voice	602	466	311	298	125	94
	Blurred face	442	532	243	203	148	143

[a] Standard face condition = normal associative connections between units.
[b] Voice condition = weakened FRU–PIN connections.
[c] Blurred face condition = weakened PCA–FRU connections.

that it is best to think of SIU activation levels as reflecting the probability that the information contained in the unit can be recalled. As the activation levels at SIU1 and SIU2 are lower in the voice condition despite threshold having been reached at the PINs, we interpret this to represent a successful simulation of our finding that semantic information is less likely to be recalled in the voice condition than in the face condition even when the voice has been found familiar. Activation levels at SIU3 and SIU4 are similar in the two conditions, but they are so low that the probability of recall of these units would be minimal.

An attempt was also made to simulate performance in the blurred-faces condition. A blurred face is likely to reduce the level of activation at the PIN not by a direct reduction of the connections between the FRU and PIN, but because of reduced activation of the FRU itself. Burton et al.'s (1999) new version of the IAC model includes an image-processing input layer, which operates according to the principles of principal components analysis (PCA). It is therefore possible to run a simulation in which the links between the PCA input layer and the FRUs have been weakened. In this simulation, the strength of these associative connections were reduced by 90%. The massive connectivity (each FRU receives input from all 50 units in the input layer) in the latest version of the IAC model means that there is automatically a great deal of redundancy in the model, and it is necessary to weaken severely the associative connections within the model in order to obtain a significant decrement in performance.

The bottom line of Table 3 presents the results of a simulation of performance in the blurred-face condition, which reveals that the FRU levels are much lower than in the standard-face condition. As a consequence of this, even though the PINs reached threshold in both conditions, the level of activation at SIU1 and SIU2 is much lower in the blurred-face condition. This provides a convincing simulation of the finding that recall of occupations in the blurred-face condition is more difficult than recall of occupations in the standard-face condition even when the blurred face has been found familiar. As in the previous simulation, the level of activation at SIU3 and SIU4 in both conditions is too low to allow recall of these items. The results of this simulation also have potential implications for explaining familiar-only responses in the voices condition. One might reasonably argue that the difficulty in recalling the occupation of a familiar voice relative to a familiar face is at least as likely to be the consequence of reduced activation of the VRU (relative to the FRU) as it is to be caused by weaker links between the VRU and PIN (than between the FRU and PIN). The results in Table 3 therefore demonstrate that the IAC model could successfully simulate the observed data regardless of whether weakened associative links between the VRU and the PIN or between the input layer and the VRU are the cause of poor recall of occupations in the voice condition.

It is also important to point out that the present results are highly supportive of the kind of cascade architecture that is found in interactive activation and competition models of face processing (e.g., Bredart et al., 1995; Burton et al., 1990). Without the cascade architecture, the only available account of familiar-only responses in Burton et al.'s model would be that they arise as a consequence of a primary problem in activating SIUs from the PINs. A problem at this level cannot possibly be the cause of the significantly larger number of familiarity-only responses in the blurred-face than in the normal-face condition, however, because exactly the same faces were presented in the two conditions. Our findings and their simulation provide a clear example of a situation in which a processing problem at one level in a stage model reduces but does not prevent the activation of units at subsequent processing stages. The FRU-to-PIN weights and the PIN-to-SIU weights were equally strong in the normal-face and blurred-face conditions, yet a manipulation that directly reduced the level of activation at the FRU had knock-on effects in terms of lower levels of activation of units in the semantic pool even when a threshold for familiarity at the intervening level of the PIN had been reached.

Subjective ratings of familiarity

Hanley et al. (1998) also found that there were significantly more familiar-only responses in the voice condition even on trials where the familiarity rating of the voice was reported as being very strong. These findings were replicated in the present study. Hanley et al. argued that the only way in which Burton et al.'s model could accommodate this finding was if it were assumed that there is no direct relationship between the relative levels of familiarity that subjects report in the face and voice conditions and the level of activation at the PIN. Even when the feelings of familiarity in the voice condition were reported to be very strong, it could then still be assumed that levels of activation of the PIN were somewhat weaker than when a face was deemed "very familiar".

The results from the blurred-face condition provide strong support for this interpretation. Participants shown blurred faces recalled fewer occupations of faces that they deemed "very familiar" than did participants in the standard-face condition. Even though the same faces were presented in the two conditions, "very familiar" in the standard-face condition did not appear to mean the same thing to subjects as "very familiar" in the blurred-face conditions. What subjects mean by "very familiar" therefore appears to be different when overall levels of familiarity are much lower and the ability to distinguish target faces from distractors is more difficult. Under these circumstances, the activation level that the PIN must reach before the person is deemed to be "very familiar" might be lower than it is when overall levels of recognition are significantly higher.

However, one intriguing aspect of the simulation presented in the bottom line of Table 3 is that it suggests a more straightforward way in which the IAC model could accommodate the finding that semantic information can be inaccessible despite strong reported feelings of familiarity. Even though the activation levels of the FRU and SIUs are much lower when the PCA–FRU links have been weakened, the level of activation at the PIN is similar in both conditions. Let us assume that the strength of the feeling of familiarity is directly related to the extent to which the level of activation of the PIN exceeds threshold (see Young & Burton, 1999, for a similar proposal when two above-threshold activation levels are compared).[1] If so, then it is clear that the condition where the PCA–FRU links have been weakened could be associated with equally strong feelings of familiarity as those in the standard condition despite the inaccessibility of all SIUs. Future experimentation that examines familiarity ratings when the same subjects are asked to make recognition

[1] It is important to point out that Young and Burton's (1999) proposal does not allow for subjects to discriminate between two different *below*-threshold activations.

decisions about both faces and voices may help to determine which of these two accounts is more likely to be correct.[2]

GENERAL CONCLUSIONS

It remains true that under standard presentation procedures, there are many more familiar-only responses to voices than to normally presented faces. However, Hanley et al.'s (1998) suggestion that this finding comes about because it is much more difficult to associate biographical information with voices than with faces receives no support from the results of the experiment reported in this paper. When overall levels of familiarity were equated, then occupations proved equally difficult to recall in the face and voice conditions. Instead, the large number of familiar-only responses in the voice condition seems to occur because activation levels or associative connections are much weaker in the voice-processing system than in the face-processing system. The effect of this is that there will be relatively more occasions where a voice will activate the appropriate PIN sufficiently strongly to achieve a feeling of familiarity without there being sufficient levels of activation in the system for semantic information to be retrieved.

Seen in the context of the simulations and the empirical findings of the present study, Burton et al.'s (1999) IAC model provides an entirely satisfactory account of the finding that occupations are sometimes more difficult to recall from voices found familiar than from faces found familiar. Furthermore, Burton et al.'s account of familiar-only experiences is uniquely successful in accommodating neuropsychological evidence from a patient who is unable to recall biographical information about people she finds familiar (Burton, Young, Bruce, Johnston, & Ellis, 1991; De Haan, Young, & Newcombe, 1991; Young & Burton, 1999). As a consequence, there is strong converging evidence to support Burton et al.'s claim that familiarity decisions are made at the level of the modality-free PINs rather than at the level of the modality-specific recognition units.

Finally, the present study has revealed that information about voices and faces is retrieved in a very similar way when the presentation format of the faces ensures that initial familiarity level is matched. The findings reported here therefore join a growing body of research (e.g., Ellis, Jones, & Mosdell, 1997; Schweinberger, Herholtz, & Sommer, 1997) that demonstrates the existence of strong parallels between the ways in which familiar faces and voices are processed.

[2] It can also be seen from Table 3 that the activation levels of the PINs are strikingly high in the simulation of the blurred-face condition and actually exceed the activation level in the standard-face condition. Why is this? In the new model, as we pointed out in the text, the PCA and FRU units are fully connected; there is a negative or positive link between all 50 PCA units and every FRU. As a consequence, there is a different pattern of activation in the blurred-face condition from the pattern observed in the standard condition. In both cases, there is a clear "winner" in the FRUs, but the residual activation in the other FRUs is different in the two cases. This residual activation also gets passed on to PINs, which means that activation of the PIN can rise as well as fall when the PCA–FRU link is weakened.

REFERENCES

Bredart, S., Valentine, T., Calder, A., & Gassi, L. (1995). An interactive activation model of face naming. *Quarterly Journal of Experimental Psychology*, *48A*, 466–486.

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

Burton, A.M., Bruce, V., & Hancock, P.J.B. (1999). From pixels to people: A model of familiar face recognition. *Cognitive Science*, *23*, 1–31.

Burton, A.M., Bruce, V., & Johnston, R.A. (1990). Understanding face recognition with an interactive activation model. *British Journal of Psychology*, *81*, 361–380.

Burton, A.M., Young, A.W., Bruce, V., Johnston, R.A., & Ellis, A.W. (1991). Understanding covert recognition. *Cognition*, *39*, 129–166.

De Haan, E.H.F., Young, A.W., & Newcombe, F. (1991). A dissociation between the sense of familiarity and access to semantic information concerning familiar people. *European Journal of Cognitive Psychology*, *3*, 51–67.

Ellis, H.D., Jones, D.M., & Mosdell, N. (1997). Intra and inter-modal repitition priming of familiar faces and voices. *British Journal of Psychology*, *88*, 143–156.

Hanley, J.R., & Cowell, E.S. (1988). The effects of different types of retrieval cues on the recall of names of famous faces. *Memory and Cognition*, *16*, 545–555.

Hanley, J.R., Smith, S.T., & Hadfield, J. (1998). I recognise you but I can't place you: An investigation of *familiar-only* experiences during tests of voice and face recognition. *Quarterly Journal of Experimental Psychology*, *51A*, 179–195.

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

Maylor, E.A. (1997). Proper name retrieval in old age: Converging evidence against disproportionate impairment. *Aging and Cognition*, *4*, 211–226.

Schweinberger, S.R., Herholz, A., & Sommer, W. (1997). Recognizing famous voices: Influence of stimulus duration and different types of retrieval cues. *Journal of Speech and Hearing Research*, *40*, 453–463.

Yarmey, A.D., Yarmey, A.L., & Yarmey, M.J. (1994). Face and voice identifications in showups and lineups. *Applied Cognitive Psychology*, *8*, 453–464.

Young, A.W., & Burton, A.M. (1999). Simulating face recognition: Implications for modelling cognition. *Cognitive Neuropsychology*, *16*, 1–48.

Young, A.W., Hay, D.C., & Ellis, A.W. (1985). The faces that launched a thousand slips: Everyday difficulties and errors in recognising people. *British Journal of Psychology*, *76*, 495–523.

Original manuscript received 8 February 1999
Accepted revision received 12 October 1999