Appendix A. Identifying Frequent Ethnic Group Specific Names

In this appendix we describe how we calculated the scores to identify those names which predominantly occur with a single ethnic category, as discussed in the Experimental Evaluation section and shown in Tables 3 and 4. We also discuss the uniformity score calculation we used to rank the names shown in Table 5.

We used the matrix $\boldsymbol{E}$, as we defined in the Methodology section, where for each name in the data set we have the number of its occurrences associated with each ethnic category, which indicate how strongly a name is associated with a specific ethnic category. By using this information, we can identify the most frequent names across ethnicities. However, rating names based on how exclusively a name belongs to a specific ethnic category alone does not capture how frequent a name is within that category. For example, the first name *Nam Woong* might only occur with the *Asian* ethnic category, however, it might only occur in two records. As a result, *Nam Woong* might be ranked at the top of the *Asian* ethnic category (with many other uniquely Asian names) even though it is a rare name (in the NCVR data set which contains several million records).

To address this issue, we define a ranking score which jointly considers both the relative frequency of a name within its predominant ethnic category, and its exclusivity to that category. For each name $n_x \in \boldsymbol{n}$ and its predominant ethnic category, $e_d \in \boldsymbol{e}$, we calculate its ethnic-specific frequency score, *efs*, as follows:

$$efs(n_x, e_d) = \left(\frac{\boldsymbol{E}[x,d]}{\sum_x \boldsymbol{E}[x,d]}\right) \times \left(\frac{\boldsymbol{E}[x,d]}{\sum_y \boldsymbol{E}[x,y]}\right),$$

where $\boldsymbol{E}[x,d]$ is the count of name n_y within its predominant ethnic category e_d, $\sum_x \boldsymbol{E}[x,d]$ is the total count of the occurrence all names within the category e_d, and $\sum_y \boldsymbol{E}[x,y]$ is the total count of occurrence of name n_x across all ethnic categories. The first term in this equation provides the relative frequency of a specific name within its predominant ethnic category, while the second term captures how strongly the name is associated with that category.

The results obtained by calculating this score for all first names and last names are presented in Tables 3 and 4, respectively, together with their *efs* scores.

To identify names that occur (close to) uniformly across ethnic categories, we calculate a uniformity score, u for each name $n_x \in \boldsymbol{n}$ based on the name frequencies across the four ethnic categories. For each name, we first calculate the mean $f'(n_x)$ and standard deviation $\sigma(n_x)$ of its frequency counts across all ethnic categories. The uniformity score of a name is then calculated as the ratio of its mean frequency divided by its standard deviation:

$$u(n_x) = \frac{f'(n_x)}{\sigma(n_x)}.$$

This uniformity score is equivalent to the inverse of the coefficient of variation (cv) [47], $u = \frac{1}{cv}$. A higher value of u indicates lower relative variability in a name's frequency across the ethnic groups. As such, names with higher uniformity scores have a more balanced distribution across the ethnic categories than names with lower uniformity scores. In Table 5 we show the top ten first and last names ranked by their uniformity score to provide examples of names that occur consistently across ethnic groups.