Matched Names Analysis Reveals No Evidence of Name Meaning Effects: A Collaborative Commentary on Silberzahn and Uhlmann (2013)

Raphael Silberzahn†
University of Cambridge

Uri Simonsohn†
The Wharton School
University of Pennsylvania

Eric Luis Uhlmann†
HEC Paris School of Management

† This collaborative commentary is the result of a productive debate between the authors of the original paper (Silberzahn and Uhlmann) and the second author (Simonsohn). Author order is alphabetical.

In an article recently published in this journal, Silberzahn & Uhlmann (2013) found that Germans with a last name with a noble meaning, such as *Kaiser* ("emperor") or *König* ("king"), were more likely to hold managerial positions than Germans with other last names. However, further data collection and new analyses reported in this collaborative commentary indicate that the described name-meaning effect is most likely attributable to name frequency. These findings suggest that the effects reported by S&U should not be interpreted as evidence of a causal effect of names on career outcomes.

S&U compared all noble-meaning surnames in the German language to the 100 most common German surnames. This was done to compare the effect of noble names to that of common German last names used throughout Germany. However, it also meant that noble names were on average less frequent than the other names in the sample. Noble last names are infrequent (few Germans have them), whereas the control last names they were compared to are frequent (many Germans have them).

Name-frequency could impact the estimated relationship between name-meaning and career outcomes if name-frequency were correlated with variables that are in turn correlated with career outcomes. Ex-ante plausible candidates include variables like socioeconomic status, urban/rural residency, and religion. More importantly, it turned out that an idiosyncratic feature of the social networking website from which the data originate, www.xing.com, created a mechanical channel by which the relative share of managers among low frequency last names is overestimated (see Supplement 1).

S&U sought to account for name frequency and other possible confounding variables using GEE regression (in their original submission) and hierarchical linear modeling (in the final published paper). Both approaches, however, assume the impact of name frequency on the probability of being a manager is linear. The skew in the distribution of frequency of German last names, and the nonlinear mechanism by which the database overestimates the share of low-frequency surnamed individuals who are managers, allowed the relationship between name-meanings and career outcomes to survive linear controls.

This makes a matched-names analysis, which compares each noble name to a set of similarly frequent names, a better analytic approach than controlling for name frequency using regression or hierarchical linear modeling. For each noble name we identified a comparison set of 50 names that were similar in frequency (see Supplement 2). For example, *Baron*– one of the noble names— is ranked as the 1016th most popular last name in Germany. *Faerber* and *Gerner*, ranked 1015th and 1017th respectively, were two of the 50 control names for *Baron*. When contrasting noble names to this more comparable set, the name-meaning effect was not observed.

For example, the data include 493 managers and 3,553 employees last named *Kaiser*, so 12.2% of *Kaisers* are managers. Among the 50 control names with the most similar overall frequency to *Kaiser*, there were 23,842 managers and 159,127 employees; so 13.0% of controls for *Kaisers* are managers. Dividing those percentages we arrive at 12.2/13.0=.94. A name-meaning effect implies a ratio greater than 1. The Figure reports the results from analogous calculations for all noble names. In aggregate there is no name-meaning effect. Noble names and their matched controls are similarly likely to be managers.

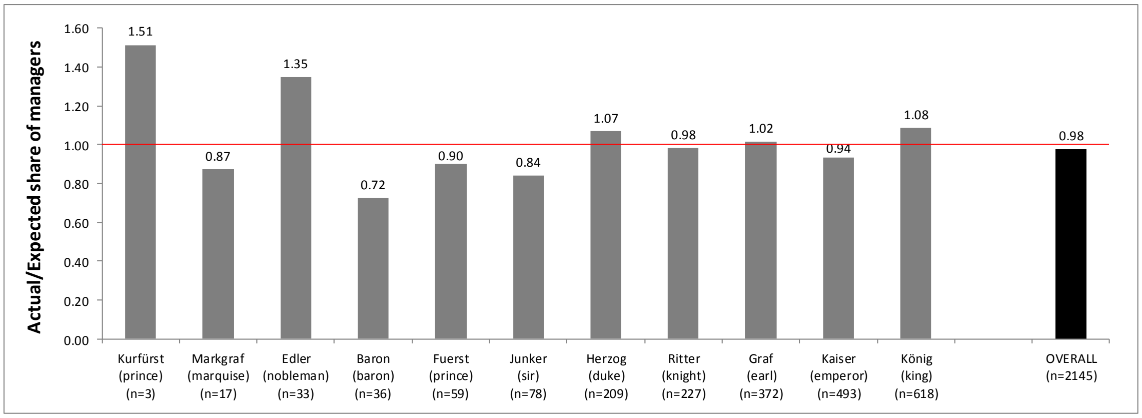

Figure 1. Germans with noble last names are as likely to be managers as controls with similarly frequent last names

Notes: the bars indicate the percentage of people with a given last name that are managers, divided by the percentage of people with the control names that are. A ratio of 1 indicates absence of a name-meaning effect. *n* below noble name is the number of managers total with that name. The "Overall" bar is based on the simple sum of total managers and employees across all names.

The new data and matched names analysis reported here indicate that at present, no significant relationship between the meaning of a person's name and his or her career outcomes can be confirmed. Potential name meaning effects remain an interesting avenue for future research, but currently lack empirical support. We hope this collaboration, where disagreements gave rise to a joint data-driven effort, can be emulated by others in the future.

References

Silberzahn, R., & Uhlmann, E. L. (2013). It Pays to Be Herr Kaiser Germans with Noble-Sounding Surnames More Often Work as Managers Than as Employees. *Psychological Science, 24*(12), 2437-2444. doi: 10.1177/0956797613494851

Author Contributions

US conducted the analyses and wrote the first draft. RS confirmed the analyses. RS, EU, and US worked together to shape the final version of the paper. Author order is alphabetical.

SUPPLEMENTARY MATERIALS

Matched Names Analysis Reveals No Evidence of Name Meaning Effects: A Collaborative Commentary on Silberzahn and Uhlmann (2013)

Raphael Silberzahn
University of Cambridge

Uri Simonsohn
The Wharton School
University of Pennsylvania

Eric Luis Uhlmann
HEC Paris School of Management

Outline

Supplement 1. Name-frequency and xing.com searches

Silberzahn and Uhlmann (2013) searched the German social networking site www.xing.com for members with a given last name, living in Germany, and identifying themselves as managers or employees (the English version of the site uses the word *executive* rather than manager). S&U restricted their data collection to German professionals for whom industry information was available to avoid including German government employees, for whom promotions are often based on seniority rather than personal characteristics. An additional reason was that in www.xing.com searches counts above N=10,000 are automatically rounded down to 10,000. For some control names, thus, accurate numbers would have not been available. Limiting data collection to German professionals for whom industry information was available allowed staying below the 10,000 limit and was intended to result in higher data quality. However, as detailed below, the S&U approach inadvertently overestimates the number of managers (relative to employees) with low frequency names.

The results page for xing.com queries include the first ten people with the requested last name and position. The information of interest is on the side of the page, indicating the aggregate membership information for the query (how many people in total have that last name). This additional information is what’s actually used to generate the analyses. Figure S1 shows print-screens from the results for “Baron” as employee and executive.

(A) Employee results *(B) Executive results*

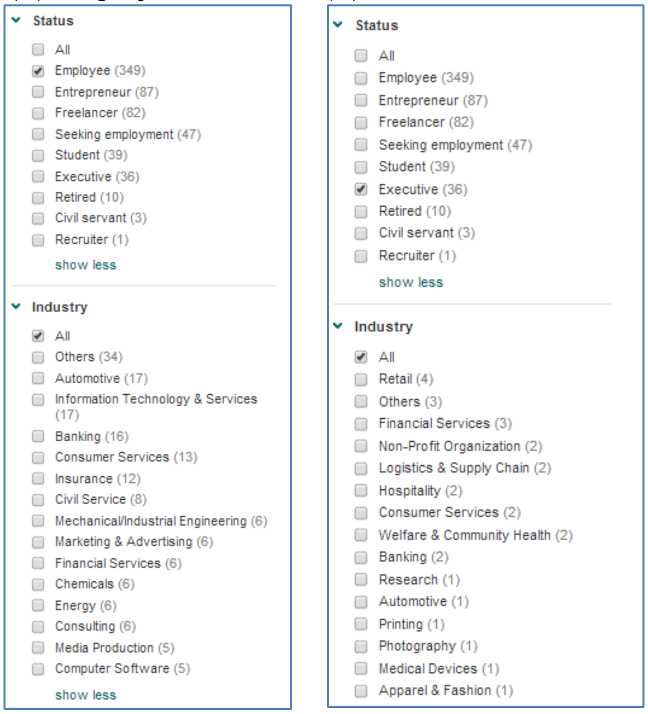

Figure S1. Print screen of search results for "Baron" in www.xing.com

Both panels show the same results under "Status," those correspond to the actual total number of people last named *Baron* that report being employees, executives (i.e., managers), students, etc. at xing.com.

As noted above, S&U added up the numbers appearing under *Industry*. Those differ between (A) and (B) because they are the subsets of employees and executives respectively. So among *Baron* employees, 34 selected "Other" as their industry, 17 "Automotive," etc. Xing.com only displays the 15 most frequent industries for the particular query so the sums will almost always be a subset of the total. For instance, if we add up the numbers of *Baron* employees with a specified industry, we get 163, less than half the actual total of 349.

Whereas all names were processed using the same procedure, this overestimated the number of managers among low-frequency names. Panel B in Figure S1 helps illustrate this bias. Adding up all the *Baron* executives listing an industry we get to 28, more than half the actual number of 36. The sample contains 47% of employees but 78% of the managers; the procedure is biased towards a higher share of managers.

This problem is less severe for more common last names because for them both manager and employee estimates are similarly biased down, and hence the ratio of managers is less biased. Take *Becker*, the 7th most frequent last name in Germany. Figure S2 shows the print-screen results for that name. There are actually 8174 employees and 1220 managers with that name on the xing.com network. Using the subset that indicated a top-15 industry (e.g., for employees, 801+409+358..), we get only 3740 employees and 636 managers, still dramatic underestimates, but they are much more even: we observe 46% of *Becker* employees and 52% of *Becker* managers. For the low-frequency *Baron*, recall, these were 47% and 78%.

(A) Employees results *(B) Executives results*

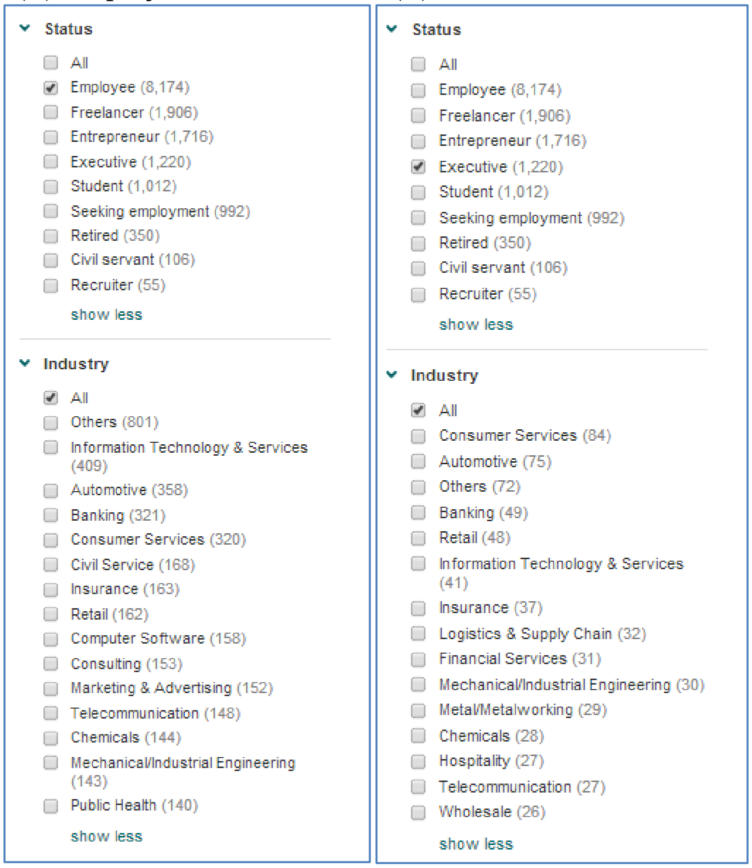

Figure S2. Print screen of search results for "Becker" in www.xing.com

Those two examples are not outliers. The median name in the top-100 most frequent names used by S&U as controls, estimates 46% as many employees as there really are and 53.5% as many managers as there really are. The median noble name, also estimates only 47% as many employees as there really are, but estimates 61% as many managers as there really are. The bias is nearly twice as pronounced for low frequency names.

Supplement 2. Selecting control names for matched names analysis

From the website http://nachname.gofeminin.de/w/nachnamen/haeufigste-nachnamen-in-deutschland.html we obtained the frequency of nearly 30,000 German last names. For each of the 11 noble last names we identified 50 similarly frequent last names.

The procedure consisted of selecting the 25 names ranked just above and just below each noble name as controls. For example, the noble name Herzog ("prince") is ranked 176th most frequent. As controls for Herzog, therefore, we used last names ranked 151-175 and 177-201. Practical considerations led to a few deviations from this general approach, see details below. These decisions were all made before collecting the data, and were not revised after conducting the analyses.

First names. S&U eliminated first names so we did the same here. Given the large number of names we proceeded in a two-step automatized process. First, when creating the set of 50 controls for each noble name, if a last name jumped out as evidently also a first name we replaced it with a contiguous name. For example, *Karl* is ranked near *Herzog* in frequency but it is a first name so we did not include it and moved to the next most similarly frequent name.

We obtained a list of the most frequent ~1000 German first names and automatically dropped all control last name that appeared in that list. After this second step, each noble name has 48 control names on average. The posted spreadsheet has the entire list and the results at the individual name level (if after 4 attempts the scraper did not obtain the data for a control name, the control name was dropped, so some noble names have closer to 40 control names).

Non-ranked noble names. One of the noble names, *Kurfurst,* does not appear on the list of the nearly 30,000 most frequent last names. As controls for this name we chose 50 random names from the least frequent last names in the list.

Low-ranked noble names. Some of the noble names, e.g., *Edler,* are sufficiently infrequent that there were more than 50 other last names with the exact same number of individuals in Germany. We selected 50 last names at random among similarly frequent last names.

Koenig and Kaiser. These two noble names are very similarly ranked: 45th and 39th. Moreover, 2 of the 6 names between then are first names (Peter & Frank), so we used the same set of control names for both noble names.

Supplement 3. Explanations for posted Excel sheet with all data

The Excel spreadsheet contains four sheets, brief explanations for them are below.

(1) Fig 1
The calculations in these sheets and hence Figure 1 use the total numbers of employees and managers reported by xing.com.

Explanation of cells:

- Cells E6:H17 contain the total frequencies of employees and managers for the control names for each noble name. They are copy-pasted from the STATA output reprinted in cells T21:Z42. These STATA analyses, in turn, use the raw data from the sheet ***(3) Fig1 Data***, see below.
- Cells J7:M17 contain the frequencies of employees and managers for the noble names. They are copy pasted from the third sheet (F1 data) as well.
- Cells P7:Q17 compute the % of managers for noble names and their controls
- Cells S7:S17 divides the two ratios and generates the values plotted in Figure 1.

(2) Fig 1 Data

- Column A has the last name of the person, column B is relevant for the control names, as it indicates which name the control is for. For example, row 19 shows that *Anderson* is a control for *Baron.*
- Columns C and D have the frequencies of employees and managers with each name. Again, these are the actual totals, not the ones obtained by industry.

(3) Name Frequency
Columns A-C were scraped from the website indicated in Row 1.
A: last name
B: ranking of the last name in terms of frequency (all ties given the same rank, next last name adds 1, so if there are three last names tied in third place, the next last name gets the 4^{th} rank).
C: # of people with that last name in the database

Columns L-N are used to compute the median rank of noble names.

(4) Xing Bias
Used to compute the degree of bias discussed in Supplement 1

Columns B&C have the frequencies adding up the industry subtotals.
Columns F&G have the total actual frequencies.
Columns I&J indicate the % share covered by the industry total from the actual total
Some rows are missing. Those correspond for last names with more than 10,000 entries in xing.com. As noted earlier, the website does not give numbers above such a threshold (reporting instead “>10,000” and hence the calculations of interested cannot be performed.

Note that for Markgraf the total from the industries is HIGHER than the actual total, this is likely caused by people indicating more than one industry and being hence double-counted.