How Do You Say Your Name? Difficult-To-Pronounce Names and Labor Market Outcomes*

Qi Ge† Stephen Wu‡

August 1, 2022

Abstract: We test for labor market discrimination based on a previously unstudied characteristic: name fluency. Analysis on two recent cohorts of economics PhD job candidates shows that those with difficult-to-pronounce names are less likely to obtain an academic or tenure-track position and are placed at institutions with lower research productivity. Discrimination due to name fluency is also found using experimental data from two prior audit studies. Within a sample of African-American candidates (Bertrand and Mullainathan, 2004) and a sample of ethnic Indian, Pakistani, and Chinese candidates (Oreopoulos, 2011), job applicants with less fluent names have significantly lower callback rates.

Keywords: labor market discrimination; name pronunciation; job placement; economics PhD job market;
JEL Codes: A11, J44

*We thank Michèle Belot, Jeffrey Cross, Wayne Grove, Benjamin Ho, Rodrigo Schneider, seminar participants at Vassar College, and anonymous referees for helpful comments and suggestions. We also thank Matthew Bartok, Kathryn Biederman, Alex Eisert, Jacob Gliedman, Daniel Goldstein, Taicheng Jin, James Kaffenbarger, Erin Kuo, Dewayne Martin, Maroun Mezher, Griffin Perry, Josue Herrera Rivera, Matthew Surprenant, Kenneth Talarico, Gwendolyn Urbanczyk, Adam Valencia, Fiona Xiang, and Chenyu Zhou for outstanding research assistance.

†Department of Economics, Vassar College, Poughkeepsie, NY 12604, *qige@vassar.edu*.

‡Corresponding Author. Department of Economics, Hamilton College, Clinton, NY 13323, *swu@hamilton.edu*.

1 Introduction

There is strong evidence for the existence of labor market discrimination along a number of different dimensions, including but not limited to: age (Riach and Rich, 2010; Neumark et al., 2019; Carlsson and Eriksson, 2019), gender (Neumark et al., 1996; Goldin and Rouse, 2000; Riach and Rich, 2006), race and ethnicity (Reimers, 1983; Cross et al., 1990; Turner et al., 1991), national origin (Pierné, 2013), religion (Wright et al., 2013), sexual orientation (Elmslie and Tebaldi, 2007; Carpenter, 2007; Ahmed et al., 2013), and physical appearance (Hamermesh and Biddle, 1994; Mobius and Rosenblat, 2006; Belot et al., 2012). One strand of this literature focuses on discrimination related to the ethnic or racial origin of an individual's name. Much of this research is based on audit studies using fictitious resumes, where only the name is changed between otherwise similar applicants. Bertrand and Mullainathan (2004) show that applicants with White sounding names receive 50 percent more callbacks than otherwise similar applicants with African-American sounding names. Along the same vein, Jacquemet and Yannelis (2012) find that those with fabricated resumes with Anglo-Saxon names receive almost one third more callbacks than identical resumes with non Anglo-Saxon names, either African-American or foreign.

These types of experimental studies have been repeated in a number of different countries outside of the United States. Ahmad (2020) shows that employers in Finland have a strong preference for Finnish job applicants over other ethnic candidates, and Carlsson and Eriksson (2019) obtain similar results in Sweden, where those with Swedish sounding names received more interview requests than those with Middle Eastern sounding names. Some researchers have tested for discrimination against ethnic sounding names across multiple ethnic groups, including Booth et al. (2012), who analyzed the Australian labor market and found that relative to those with Anglo-Saxon names, all other ethnic groups had lower callback rates for interviews; and Oreopoulos (2011), who found substantial discrimination in Canada against job seekers with Indian, Pakistani, and Chinese names, relative to those with English names. And there is also research documenting name-based discrimination in other settings. For instance, Ahmed and Hammarstedt (2008) provide evidence for discrimination in Sweden's rental housing market, as landlords were much less likely to contact people with Arabic/Muslim names than those with Swedish names, and Sweeney (2013) shows that advertisements for public records on a person or ads containing the word "arrest" are more likely to appear in internet searches involving Black-associated names than in searches for White-associated names.

The existence of name-discrimination along racial or ethnic lines can have significant consequences. Some individuals who are cognizant of the potential drawbacks of having ethnic sounding names may have a desire to change them. Historically, this has been a common practice, as documented by Biavaschi et al. (2017), who show that European immigrants to the United States who Americanized their names experienced larger occupational upgrading and higher earnings than those who did not. Zhao and Biernat (2017) conduct an experiment where emails from Chinese students attending college in the United States were sent to several hundred White professors, with only the name of the sender being varied. Emails using a Chinese first name had significantly lower responses to requests for a meeting than those using an Americanized first name. This desire to hide one's race or ethnicity in a job application may also manifest itself in ways other than changing names. Kang et al. (2016) demonstrate that many minority job applicants engage in resume "whitening," where they strategically omit information that could signal their race, such as participation in race-based nonprofit organizations or affinity groups. Results from their audit study show that there are large benefits to resume whitening (both in terms of changing one's first name and altering activities/experience on one's resume) for Asian and Black job applicants, as measured by higher callback rates.

In this paper, we study the labor market effects of a related but distinct characteristic of names, name fluency. While there is now significant and growing evidence of discrimination against names that signal one's ethnic or racial background, there are no studies that we are aware of that have tested the hypothesis that having a difficult-to-pronounce name may lead to worse job market outcomes. Although many ethnic sounding names are also difficult to pronounce, particularly for those outside of that particular racial or ethnic group, there is still variation in the fluency of names within particular groups. For example, most non-Chinese speakers would consider *Chen* to be more familiar and easier to pronounce than *Xiang*; people without a Polish background will generally have much more trouble trying to pronounce the surname *Przybyłko* than they will with *Nowak*.

Even after controlling for the ethnic or racial origin of one's name, there are a few reasons that individuals with hard-to-pronounce names may experience worse outcomes in the labor market. There may be subconscious bias against those with difficult-to-pronounce names, leading potential employers to have more negative evaluations for these applicants and be more critical of their profiles. Recruiters will also have an easier time processing and remembering names that are more fluent and/or familiar sounding. Some individuals on hiring committees may undertake the mental effort to remember difficult sounding names, but others may not. Belot and Schröder (2022) study the setting of academic conferences at the University of California, San Diego (UCSD) and the University of Edinburgh and find that when there are significant numbers of people with minority attributes, they are often confused with one another and blended together. There is also evidence consistent with these theories in other settings. Alter and Oppenheimer (2006) demonstrate that short term stock market returns are affected by fluency of the names of stocks. Specifically, a basket of shares that are easy to pronounce yielded an 11 percent higher return than another comparable basket with less fluent stock names for the New York Stock Exchange. Laham et al. (2012) conduct a series of experiments in Australia and find that those with easier-to-pronounce names are judged more positively by others and also hold more prominent positions in law firms. In Newman et al. (2014), university students in New Zealand were more likely to believe claims from those with easy-to-pronounce names relative to those with more difficult names. And processing fluency can even affect motivation and behavior. For example, Song and Schwarz (2008) find that when instructions for a task were presented in hard to read fonts, there were negative downstream effects on their willingness to engage in the tasks.

In this study, we examine three sources of data and document significant evidence of discrimination based on name fluency. First, we utilize observational data from the academic labor market by assembling curriculum vitae (CV's) of over $1,500$ economics job market candidates from 96 top ranked economics PhD programs from the 2016-2017 and 2017-2018 job market cycles and find that name fluency is significantly related to job market outcomes. Specifically, candidates who have difficult-to-pronounce names are much less likely to be initially placed into an academic job or obtain a tenure track position, and they are placed in jobs at institutions with lower research productivity, as ranked by the Research Papers in Economics (RePEc) database. Our results are consistent and robust across three separate ways of measuring pronunciation difficulty: an algorithmic ranking based on commonality of letter and phoneme combinations, a survey-based measure that records the average time it takes individuals to pronounce a name, and a purely subjective measure based on individual ratings. These results also hold after controlling for a large number of covariates including PhD institution and home country.

We next augment these real world labor market results by analyzing experimental data from two prior audit studies discussed above. In analysis of data from Bertrand and Mullainathan (2004), we find that job applicants with less fluent names have lower callback rates, even after accounting for the implied race of the candidate. What is particularly striking is the fact that within the sample of resumes with distinctly African-American names, name fluency is still strongly correlated with callback rates. We also document similar results using data from another prior audit study by Oreopoulos (2011). Once again, job applicants are less likely to be called back when they have names that are difficult to pronounce, and even when restricting the sample to immigrants with ethnically Indian, Pakistani, and Chinese names, those whose names are less fluent are significantly less likely to be called for a job interview.

Despite data constraints, we also explore possible mechanisms for these results by testing for heterogeneity in the size and strength of name penalties according to the profiles of job applicants. Indeed, we find that for both the academic labor market data and the analysis of prior audit studies, those who are less qualified and have weaker resumes tend to encounter a larger degree of discrimination due to name pronunciation than their more qualified counterparts with stronger resumes. This suggestive evidence is consistent with the idea that there are mental costs to processing and remembering less fluent names, and that these mental costs are only worth spending on higher quality candidates.

There are several unique aspects of this paper. We test for the existence of labor market discrimination due to a previously unstudied characteristic: name fluency. In doing so, our empirical analysis also introduces key methodological innovations. Unlike most other studies of name-based discrimination, we employ observational data using actual names and resumes, in addition to our analysis of prior audit studies that randomly assign names to fictitious applicants. Our labor market data also incorporates hiring outcomes, which complements the analysis of initial callback rates that are the focus of most earlier research. Moreover, we contribute to the literature on word fluency by incorporating three separate ways of rating the difficulty of pronouncing names. Each of these measures is independently derived, but nonetheless highly correlated with one another. One of these measures is similar to what has been used in prior work on this topic: a collection of independent subjective assessments by three different individuals. The other two measures are novel in their approaches: an algorithmic ranking that provides a relatively objective way to rate the fluency of names; and the median time that it takes people to pronounce a first and last name based on survey data. These innovative name fluency measures might be used by others researchers working on the impacts of textual fluency.

2 Primary Data and Empirical Strategy

2.1 Data on Job Market Candidates and Outcomes

To study the potential impact of name pronunciation on job market outcomes, we employ data on economics PhD candidates who entered the job market in either 2016-17 or 2017-18.[1] We manually collected curriculum vitae from professional websites or departmental placement pages for PhD students from the 96 top-ranked economics doctoral programs, as ranked by the U.S. News & World Report (USNWR) in 2017.[2] Although this data represents only one specific labor market, there are advantages to studying the job market for economics PhDs. We have a fairly comprehensive pool of candidates searching for jobs requiring a doctoral degree in economics, given the large number of graduate programs included in our sample. In addition, curriculum vitae are publicly available and easily accessed, and they tend to follow similar and consistent formats, so they provide reliable and comparable information about candidates' research productivity and teaching experience. Furthermore, the ranking of departments according to research productivity indices offers one measure of job placement quality, albeit an incomplete and imperfect one.

Students' initial job market outcomes are primarily obtained from official departmental web pages that list placement histories of their graduates. When departments do not list specific names in their placement outcomes, we search for this information on personal websites or publicly available *LinkedIn* profiles. We categorize job placement into three categories: academia; government or research think tanks; and the private sector. Within academic positions, we distinguish between tenure track jobs and other academic positions such as post-doctoral fellowships or visiting (non-tenure track) positions.[3]

For academic, governmental, and think tank positions, placement outcomes are ranked

[1] This data contains a similar sample to one that has been used to study the impact of student-advisor matching on job placement outcomes in Ge et al. (2021).

[2] The 2017 USNWR listing of top graduate programs in economics rates programs all the way down to number 90, but because of ties, the list includes a total of 96 programs.

[3] If a candidate lists a postdoc as the initial placement, but also lists a tenure-track position that will start upon completion of the postdoc, we code such placement as the tenure track position.

based on institutional research productivity from the Research Papers in Economics database, as published in February 2019.[4] Compared to alternative indices on economics research productivity, such as those introduced in Coupé (2003) that focus exclusively on economics departments in colleges and universities, the RePEc productivity ranking covers a broader range of institutions and thus offers a more comprehensive indicator of placement quality. In particular, the RePEc index provides explicit productivity scores from 1 to 389 (in descending order of productivity) for the top 5 percent of all economics institutions around the world. Institutions outside of the top 5 percent are only ranked within their percentile in the ranking system and are not distinguishable from each other. For the purposes of this study, all institutions ranked in the top sixth percentile are coded as having a RePEc productivity score of 400. Similarly, the remainder of the institutions ranked in the top 10 percent are coded as having a RePEc score of 500 (top 7 percent), 600 (top 8 percent), 700 (top 9 percent), or 800 (top 10 percent). Institutions that are not listed among the top 10 percent are coded as having a RePEc score of $1,000$. For academic positions, we only use the rank of an institution in the top 10 percent when the individual has been placed in a tenure track job. Candidates that are placed in post-doctoral fellowships or non-tenure track academic positions are given a RePEc score of $1,000$, even if the institution is ranked among the top 10 percent. To maximize the number of observations, all individuals placed in the private sector and those with missing job information have been given an imputed RePEc score of $1,000$. However, we also show that our results are robust to excluding these imputed observations.

2.2 Name Fluency Ratings

To measure the difficulty in pronouncing individual names, we use several different methods: a computer-generated algorithm that assesses the difficulty of pronouncing various words, a rating based on the median time it takes for people to pronounce a particular name, and a

[4] These rankings tend to be fairly stable over time, so the specific choice of the ranking date does not affect results.

subjective measure based on three independent raters. We discuss each of these methods in more detail below.

In prior research on word fluency, a common way to rate difficulty of pronunciation is to have individuals use their own criteria to decide whether a word is easy or difficult to pronounce (Laham et al., 2012; Newman et al., 2014). While this type of rating scheme is straightforward, it is prone to potential bias, as an individual's subjective perception of the difficulty in pronouncing a name may not equate to how hard it actually is to pronounce a name in practice. Therefore, our primary measure of name fluency is based on a more objective method, where we created an algorithm to rate the difficulty of pronouncing words from the vantage point of a native English speaker. This algorithm developed two measures of difficulty in pronouncing words: one based on sequences of letters and the other based on sequences of phonemes, or distinct units of sound. Occurrence dictionaries of common letter and phoneme combinations were then used to develop one difficulty rating based on the frequency of letter combinations and another rating based on the frequency of phonemes embedded in various names, resulting in two relatively objective ratings of fluency/familiarity for first names and two ratings for last names, where each of these individual ratings were bounded between 0-100. We use an arithmetic average of these two ratings to measure the difficulty of pronouncing a particular first name or last name and summed these two ratings to provide a measure of the pronunciation difficulty for each individual's full name.[5]

As a second measure, we hired ten research assistants to independently go through a series of surveys of first and last names of the job market candidates in our data and attempt to pronounce each name.[6] Each screen on the survey contained one first name or one last name; we asked the research assistants to read through the name and click to advance to the next screen to see the next name. This procedure was then repeated until a particular survey of names was completed. Each survey contained timing functions that kept track of the length

[5]The two sub-rating schemes are highly correlated and results are robust to using only the letter based rating, only the phoneme based rating, or any weighted average of the two.

[6]We assembled a fairly diverse group of research assistants across race, ethnicity, gender, and country of origin.

of time that each person spent on a particular screen before moving onto the next name. We summed the total time it took each research assistant to pronounce a candidate's first name and last name to obtain the total time spent on the candidate's full name. We then used the median time (out of all research assistants) spent on a candidate's full name as another measure of pronunciation difficulty.[7] The full instructions on this method are included in Online Appendix A.

Lastly, as a third measure of rating name fluency, we implement the subjective method common in prior research, in spite of its shortcomings. In following this method, three individuals conducted independent ratings for each first name in a binary manner: the name was deemed either difficult to pronounce or easy to pronounce. Independent ratings of last names were done in the same fashion. This rating system provided a number of different possible ways to code the difficulty of pronouncing one's name, but we primarily used a dichotomous variable which was equal to one if at least two raters deemed a candidate's first name as being difficult to pronounce or if at least two people deemed a candidate's last name as being difficult to pronounce.[8]

We recognize that there could be several possible sources of measurement error in rating the fluency of names. The algorithm-based rating is based on the perspective of a native English speaker, and search committees may have individuals from a variety of native countries. Similarly, the team of survey-taking research assistants, while reflecting diversity across a number of demographic characteristics, may or may not reflect the typical makeup of a hiring committee. And subjective ratings are always potentially flawed for reasons discussed earlier. Nonetheless, we believe that the use of three independent ways of rating name fluency provides a comprehensive way to test for name-based labor market discrimination.

[7] Research assistants were instructed to advance to the next screen immediately after pronouncing each name.

[8] We also tried alternative schemes, such as having at least one person rate a first or last name as being difficult, having all three consider a name as difficult, or summing the total number of raters that deemed a candidate to have a difficult to pronounce first or last name. Results were consistent across different methods.

2.3 Empirical Strategy

As an initial test of the hypothesis that difficulty of name pronunciation affects a job candidate's initial placement, we estimate the following probit model:

$$Pr(Placement_i = 1) \quad = \Phi(\delta' Pronounce_i + \phi' X_i + \theta' Z_i) \tag{1}$$

where $Placement_i$ denotes dichotomous placement outcome variables for candidate i, including 1) an indicator that is equal to one if the candidate is placed in academia (including tenure track, visiting, or postdoc positions), and 2) an indicator that is equal to one if the candidate is placed into a tenure track position.

X_i is a vector of job candidate characteristics in line with prior literature (e.g., Krueger and Wu (2000), Athey et al. (2007), and Sullivan et al. (2018), among others) that includes gender, whether the undergraduate institution attended is an elite college,[9] prior graduate degree(s), number of publications, number of publications in top 5 economics journals,[10] number of papers currently under revise and resubmit status, number of revise and resubmit invitations at top 5 journals, and dichotomous variables for the receipt of a teaching award and for having independent instructor experience. In addition, Z_i captures the following characteristics of the candidate's primary dissertation advisor: the number of publications in top 5 journals, having experience of holding an editorial position, and having experience of holding an editorial position at a top 5 journal. Finally, we control for a range of fixed effects that may affect placement outcomes, including region of undergraduate degree (US and Canada, Latin America and the Caribbean, Eastern Europe, Western Europe, South Asia and the Middle East, East and Southeast Asia, Australia, and Africa), sub-field (theory, macro/finance, econometrics, and applied micro), job market cycle (2016-17 versus 2017-18),

[9]Similar to Athey et al. (2007), we define elite colleges as 1) Ivy League universities; 2) other top 15 national universities according to the 2019 U.S. News & World Report college ranking; and 3) top 5 liberal arts colleges, as rated by USNWR.

[10]We use the *Quarterly Journal of Economics*, *American Economic Review*, *Review of Economic Studies*, *Journal of Political Economy*, and *Econometrica* as our group of top 5 journals.

and PhD institution.

Our independent variable of interest, $Pronounce_i$, represents the difficulty of pronouncing a candidate's name. As discussed earlier, we incorporate three different measures of name difficulty: an algorithmic ranking based on the familiarity and frequency of letter and phoneme combinations within a name, the median time it took for a team of research assistants to pronounce a candidate's first and last names, and a purely subjective measure (akin to prior studies) based on individual ratings. One potential confounding issue in measuring name fluency is that longer names may be more difficult and take longer to pronounce, all else constant. To distinguish between the effects of name length and name difficulty, we also include variables controlling for the number of letters in a candidate's first and last names.

As a second measure of placement outcomes, we also investigate the determinants of job placement ranking based on the RePEc database's ranking of economics institutions according to their research productivity. For our main specifications, we impute the placement quality of candidates placed in the private sector or candidates obtaining non-tenure track jobs at academic institutions with a RePEc score of $1,000$ and run a tobit specification with an upper limit at $1,000$, but we also re-estimate our regressions by excluding these observations and obtain similar results. Similar to Equation 1, our tobit models include job candidate characteristics, advisor characteristics, rating of name fluency, name length, as well as region, sub-field, job market cycle, and PhD program fixed effects.[11]

2.4 Summary Statistics and Descriptive Evidence

Our raw data contains a total of $1,660$ candidates from the top 96 ranked economics PhD programs who were on the job market in either 2016-2017 or 2017-2018. After dropping observations with missing information, we obtain a sample of approximately $1,500$ candidates for our regression analyses.[12]

[11] We also implement a quantile (median) regression model on the RePEc rankings and find the estimates to be comparable to the tobit model. The quantile regression results are available upon request.

[12] Other than a handful of candidates for whom we could not locate any placement information, the majority of job candidates excluded from our regression sample were placed in the private sector and did

Table 1 presents the summary statistics for our sample. Most of the job market candidate characteristics are in line with prior studies. 70 percent of our sample is male, slightly over half completed a prior graduate degree, about 40 percent completed their undergraduate education in the U.S., and close to 10 percent attended an elite U.S. undergraduate institution. The majority of job market candidates have limited publication experience as revealed by the mean number of publications or papers receiving a revise and resubmit invitation, both of which are less than one. In terms of job market outcomes, 65 percent of candidates are initially placed into academic jobs, with the majority of these being tenure track positions. As explained earlier, we impute the RePEc ranking to be the highest (worst) ranking for those who are placed in private sector positions ($N = 611$). The mean imputed RePEc ranking is close to 800, equivalent to the top 10 percent of economics institutions around the world, but the median is $1,000$.[13]

Turning to our measures of name fluency, we find that the computer-generated algorithm produces a wide range of pronunciation ratings, with a low of 0.21 and a high of 185, and mean and median values of approximately 55. The median time it takes to pronounce a candidate's full name is approximately 2.5 seconds, though this also has a wide range from just over 1.5 seconds to nearly 6 seconds. Slightly less than 30 percent of names were subjectively rated as having a difficult to pronounce first or last name by at least two of three individual raters.

3 Results and Discussions

3.1 Baseline Findings

Table 2 shows the probit marginal effect estimates of the likelihood of being initially placed in an academic job (columns 1-2) or a tenure track job (columns 3-4), and tobit estimates for

not have publicly available resumes.

[13]The mean and median RePEc rankings without imputation ($N = 910$) are 635 and $1,000$, respectively.

the research ranking of initial placement (columns 5-6). In column 1 (3), we see that getting an academic (tenure track) job is positively related to the number of revise and resubmit invitations and publications, and to having completed a prior graduate degree or receiving a teaching award, though the last coefficient is not quite statistically significant for the tenure track regression. Column 5 exhibits negative and significant coefficients for prior graduate degrees, number of papers under revise and resubmit status, and teaching awards, implying that these are all correlated with placing in more productive research institutions, as a smaller RePEc score corresponds to an institution with greater research output.[14] Having an advisor who is a journal editor increases the probability of being placed in an academic job (column 1), but is not significantly related to the other outcome variables. We do not find significant gender differences for any of the job outcomes.

We now turn to the focus of this study: the effect of name fluency on job outcomes. Our algorithmic measure of name fluency is normalized in all regressions for ease of interpretation. The coefficient in column 1 implies that for a one standard deviation increase in the difficulty of pronunciation, candidates are 4 percentage points less likely to land an academic job, and this coefficient has a p-value of 0.014. The analogous effect of a one standard deviation increase in name difficulty on the likelihood of obtaining a tenure track position (column 3) is 1.9 percentage points, though this coefficient is not precisely estimated. In terms of the ranking of an individual's initial placement, a one standard deviation increase in name difficulty is correlated with an initial job placement that is 83 spots worse in the RePEc ranking, which would translate to a difference between a job in the economics departments at MIT versus Georgia State, for example, or the difference between a job at Duke's Fuqua School of Business and a job at University College Dublin.

Next, we test to see if there are differences between the fluency effects of first names and last names. Column 2 suggests that hard-to-pronounce first names and last names are both

[14] As described earlier, our main specifications include a full set of PhD program fixed effects, which slightly reduces the sample size in the probit regressions from the original $1,510$ to $1,469$ (columns 1-2) and $1,499$ (columns 3-4) due to collinearity.

negatively correlated with placing in academia, though the coefficient for last names has a large standard error. The marginal effects for a one standard deviation increase in name difficulty of first and last names is −3.7 and −2.0 percentage points, respectively, though a t-test fails to reject the statistical equivalence of the two coefficients. In column 4, we observe that neither of the coefficients for first name or last name difficulty is statistically significant in predicting tenure track placements. With regards to the institutional ranking of initial placement, pronunciation difficulty of first and last names are both significant predictors, with statistically indistinguishable coefficients of 57 and 68, respectively.

3.2 Robustness Checks

We conduct additional analysis to assess the sensitivity of our baseline findings. First, we test for the presence of name fluency discrimination using two alternative ways to rate pronunciation difficulty and show these results in Table 3. The measure of pronunciation difficulty in columns 1, 3, and 5 is the normalized median time it took ten research assistants to pronounce a particular job market candidate's first and last names. Once again, we find that having a harder to pronounce name is correlated with a lower probability of being placed in an academic (marginal effect of −7.4 percentage points for a one standard deviation increase in difficulty) or tenure track position (marginal effect of −4.7 percentage points), and obtaining a less prestigious initial job placement (marginal effect of 83 ranking spots), and all coefficients are significant at the 5 percent level or better. Columns 2, 4, and 6 use information from three subjective ratings as the measure of name fluency, where the independent variable of interest is equal to one if at least two raters deemed the first name or the last name to be difficult to pronounce, and zero otherwise. Having a subjectively difficult name decreases the likelihood of obtaining an academic (tenure track) job by 11.6 (5.5) percentage points, and is linked with a job that is 82 spots worse according to the RePEc ranking, though the tobit coefficient in column 6 is not statistically significant.[15]

[15]In Table A1 in the Online Appendix, we also seek to distinguish the impact of name fluency by first and last names using alternative ways to rate name pronunciation difficulty. Similar to the results presented

Next, we perform a number of additional robustness checks on our empirical specifications, including adopting raw RePEc rankings without imputation as an alternative measure for placement quality, estimating multinomial logit regressions for different types of placements, and employing an ordered probit model based on categories of the RePEc ranking of job placements. We also account for the possibility that students who have advisors and committee members from the same country might be less likely to feel the need to Americanize/Anglicize their (first) names. To streamline the presentation in the main text, we present and discuss these robustness checks in Online Appendix B.

It is possible that difficult to pronounce names are concentrated in a few countries, and the lack of success that individuals from these countries have in finding prestigious academic jobs is not necessarily linked to their names but from more general discrimination due to national origin. All regressions shown in our tables have controlled for the region of one's undergraduate school, but we have also estimated specifications which include a full set of individual country effects, and the results are largely the same.[16] In addition, we have also run separate regressions for different regions, though statistical power is reduced in regions with few observations. In general, we observe that the effects of name fluency on placement types and quality are not driven by a particular region of undergraduate degree, as the magnitude of the effects are large and significant for several different regions.

3.3 Sources of Bias in the Effects of Name Fluency

One potential threat to the validity of our results is that name difficulty is related to unobserved measures of ability. As one check, we calculate correlations between our name difficulty measures and the control variables that we have available and present the results in Table A8 in the Online Appendix. Here, we observe that there are no significant correlations between candidates' names and their productivity-related characteristics, such as the

in Table 2, most specifications again do not reveal statistically significant differences between the fluency effects of first names and last names.

[16] These results are presented in Table A7 in the Online Appendix.

number of publications or papers under revise and resubmit status, teaching awards, and prior instructor experience. This suggests that a candidate's name is unlikely to correlate with unobserved ability as (imperfectly) proxied by these measures. Nonetheless, we will address the concern of omitted variable bias further when we analyse experimental data from prior audit studies later in the paper.

A related concern is that some individuals may choose to Americanize/Anglicize their names, and name changing may be endogenous. If individuals who choose to change their names are particularly savvy and would be more likely to succeed on the job market even if they did not switch their names, we will overestimate the size of name-based discrimination. In recent time periods, adoption of simplified or Anglicized names are more often confined to first names, and when testing for the separate effects of first names versus last names, we find the magnitudes of the effects roughly equivalent. Furthermore, as shown in Table A9 in the Online Appendix, our results also hold when we exclude all candidates with ethnically Chinese names, a group for which one is most likely to find individuals who adopt Americanized first names. The decision of whether or not to change one's last name after marriage may also be endogenous, though separate analysis by gender does not reveal any differences in the effects of name fluency. In particular, as shown in Table A10 in the Online Appendix, the fact that we find similarly sized effects for the sample of male job market candidates (where changing last names is much less common than for females) suggests that endogenous switching of surnames is not driving our results either.

Another possibility is that the uniqueness or commonality of names could also impact job outcomes. It is likely that those with very common names could be at a *disadvantage* because they do not stand out from other candidates. Because pronunciation difficulty is likely negatively correlated with commonality of names, our estimates of the name fluency effect might actually be underestimated. To address this, we conduct additional regressions by adding controls for having a common first name and common last name and discuss the findings in Online Appendix C. Overall, our results indicate that controlling for commonality of names does not affect the magnitude or significance of the name difficulty coefficients in our findings.

3.4 Potential Mechanisms

Overall, our analysis of recent cohorts of economics PhD job market candidates suggests that those with difficult-to-pronounce names are less likely to obtain an academic or tenure track position and tend to be placed at institutions with lower research productivity. While we are unable to identify the precise mechanisms through which this type of discrimination acts given the observational nature of the data, we can speculate about some possibilities. For job searches at academic, governmental, and research institutions, an initial screening generally involves committees getting together to discuss names of potential candidates, which may lead to some subconscious discrimination against names that are harder to pronounce and/or remember. Another possibility is that there exist mental costs to processing and remembering less fluent names, and that these mental costs are only worth spending on higher quality candidates. In our context, such explanation would imply that job market candidates with strong profiles may experience less name-based bias than those with weak profiles.

To test this hypothesis, we first follow Bertrand and Mullainathan (2004) and map the quality of the job market candidates according to their predicted RePEc rankings of placement institutions.[17] We then focus on subsamples of candidates with relatively strong and weak predicted profiles, i.e., those above and below the median predicted placement RePEc ranking, and separately estimate our placement type and quality specifications. Additionally, since the predicted RePEc ratings are highly correlated with PhD program ranking published by the USNWR (correlation: 0.927), we also consider an analogous subsample analysis based on students coming from top 20 versus non-top 20 programs.

Table 4 reports the estimated coefficients for the algorithm-based name difficulty measure. When comparing candidates with high (column 1) versus low (column 4) predicted RePEc

[17] Predicted RePEc rankings of placement institutions are estimated using the set of candidate and advisor characteristics described in Section 3.2.

rankings of initial placement, those with weak predicted profiles experience a significantly larger name penalty in the likelihood of being placed into an academic position, where the p-value on the difference between coefficients in columns 1 and 4 is 0.045. However, when comparing tenure track placement (columns 2 vs. 5) and placement quality (columns 3 vs. 6), we do not observe statistically significant differences between the two groups. Likewise, we find that students with difficult-to-pronounce names from non-top 20 programs are significantly less likely to land an academic placement compared to their counterparts from top 20 programs (p-value for the difference in coefficients is 0.077), but there are no significant differences between the two groups for their chances of obtaining a tenure track position or for the RePEc ranking of their placement.[18]

Conditional on candidate profiles, it is also possible that employers may be more likely to overcome mental costs associated with processing and remembering less fluent names when the recruitment involves higher stakes. For example, recruiting for a tenure track position represents a greater long-term investment for the department relative to a temporary position. When focusing on candidates with a tenure track placement and examining the role of name fluency on their placement quality, we once again find that name fluency is negatively correlated with RePEc placement ranking (shown in column 1 of Table A12 in the Online Appendix). However, when we focus on those who were placed in the most research-intensive tenure track positions (RePEc<400), the name penalty almost completely disappears (point estimate = 0.591 with a p-value of 0.95). In other words, the comparison here suggests that the name penalty is minimal in the recruitment of the most research-intensive and, presumably, most selective positions, which further reinforces the potential role of mental costs in facilitating name bias.

Our results here offer suggestive evidence that candidates with relatively weaker profiles and those coming from lower ranked PhD programs are more likely to suffer from name bias than those with stronger profiles and those from top ranked doctoral programs, where

[18]We also explore similar specifications for students from top 30 and non-top 30 program and observe analogous patterns. These results are available upon request.

the name penalty is minimal. Take together, these findings hint at the possibility that the observed name-based bias may work through the channel of increased mental costs necessary to process and remember candidates with difficult names, and hiring committees are only willing to expend this energy on candidates with strong credentials. In the next section, we will leverage experimental data from prior audit studies to further corroborate our baseline findings and gain more insights regarding the mechanisms at play.

4 Experimental Data from Prior Audit Studies

Our analysis of two recent cohorts of economics PhD candidates shows that name fluency is strongly correlated with labor market outcomes in one specific setting. In particular, having a name that is difficult to pronounce decreases the chances of landing an academic or tenure track position and leads to placement in a less prestigious institution as ranked by research output. Although we have controlled for a large set of characteristics related to the job candidates and to their graduate schools and advisors, as discussed in Section 3.3, it is possible that name fluency may be related to some unobserved variables that could also impact job placement. As an additional way to test the hypothesis that name fluency may penalize an individual's job prospects, we incorporate experimental data from two important and well-cited audit studies that look at how callback rates for job applicants are impacted by having either a distinctively African-American name Bertrand and Mullainathan (2004) or being a foreigner with an ethnic sounding name Oreopoulos (2011).[19]

4.1 Job Applicants with Distinctively Black Sounding Names

The audit study by Bertrand and Mullainathan (2004) involved sending fictitious resumes to postings for jobs in the areas of sales, administrative support, clerical, and customer services in the cities of Boston and Chicago, where applicants were randomly assigned either distinc-

[19] The publicly available replication data for Bertrand and Mullainathan (2004) and Oreopoulos (2011) can be obtained from openICPSR at `https://www.openicpsr.org/openicpsr/`.

tively White-sounding names or distinctively Black-sounding names.[20] In the first column of Table 5, we replicate and show the main finding of the paper: those with Black-sounding names have significantly lower callback rates than those with White-sounding names. Specifically, the marginal effect of the probit regression shows that the average callback rate is 3.2 percentage points lower for Black job applicants than for White applicants, representing a very large effect, as the average callback rates for Blacks and Whites are only 6.45 and 9.65 percent, respectively.

Of particular benefit for our study is the fact that there are many different White-sounding and Black-sounding first names (18 for each race, or 36 in total) assigned to the various job applicants, so we use our previously described algorithm to generate ratings of name fluency/familiarity.[21] In column 2, we estimate a probit regression that only includes the algorithmic rating of name fluency (once again, normalized for ease of interpretation) as an independent variable, and we find a statistically large and significant effect. A one standard deviation increase in name difficulty is correlated with a 1.7 percentage point decrease in the callback rate. Given that name fluency is correlated with race, we include both of these factors in the regression shown in column 3. The magnitude of the coefficient on the indicator variable for Blacks is now reduced to -0.018 (p-value$= 0.1$), and the coefficient on the standardized measure of name difficulty is -0.011 (p-value$= 0.07$). This suggests that one (though certainly not the only) mechanism through which name-based racial discrimination works is the fluency of the name. In column 4, we include a set of additional resume characteristics employed in Bertrand and Mullainathan (2004),[22] and we find that name difficulty continues to be an important and significant factor in callback rates. The indicator variable for Black names is still negative, though it is no longer statistically significant.

[20] Note that the publicly available dataset for Bertrand and Mullainathan (2004) only contains first names without last names. Our analysis on name difficulty will therefore be limited to first names, which are also the focus for their paper.

[21] Bertrand and Mullainathan (2004) also randomized last names, but they do not discuss any of these results and the publicly posted data does not include last names.

[22] These resume characteristics include gender, education, number of jobs listed on resume, years of work experience, honors received, volunteering and military experience, and dummies for working in school, listing an email address on resume, computer skills and other special skills.

To more clearly distinguish discrimination due to name fluency from discrimination based on perceived race, we restrict our analysis to the sample of African-American job applicants and show marginal effects of probit regressions in Table 6. In column 1, we see that even within the sample of African-American candidates, those with less fluent/familiar sounding names are less likely to receive callbacks. The coefficient of -0.011 (p-value$= 0.03$) implies that a one standard deviation increase in name difficulty leads to a decrease in the rate of callbacks by 1.1 percentage points, which represents a 20 percent change in the overall rate for Blacks. After including a set of additional resume characteristics in column 2, the impact of name fluency actually increases (coefficient of -0.018) and remains statistically significant (p-value< 0.01).

In documenting the relationship between name difficulty and callback rates within a particular minority group, an important question is the manner in which this discrimination occurs. One possibility is that there is some general bias (conscious and/or subconscious) towards those with names that are hard to pronounce. But when we conduct separate analyses for the group of high quality resumes and low quality resumes (as constructed by Bertrand and Mullainathan (2004)),[23] we find interesting differences in the results. In particular, we find no evidence of name-based discrimination towards Black applicants with high quality resumes (columns 3-4), but we observe large and statistically significant effects of name fluency on callback rates for Black applicants with low quality resumes (columns 5-6). This difference in effects complements analogous findings from our analysis of recent cohorts of economics PhD job market candidates (as discussed in Section 3.4) and is consistent with the idea that there are mental costs to processing and remembering names that are less familiar or difficult to pronounce, and that these mental costs are only worth spending on higher quality candidates. The implication is that discrimination based on name fluency will mostly occur for those with lower quality resumes, as employers may develop quick heuristics to make decisions about these candidates and for certain types of positions, which they do

[23]Specifically, Bertrand and Mullainathan (2004) consider a resume of "high-quality" if it ranks above the median predicted callback rate and "low-quality" if it ranks below the median predicted callback rate.

not necessarily do for those with better resumes.

As a further way to measure the relationship between name difficulty and callback rates, we calculate a simple Pearson correlation between callback rates and algorithmic ratings of name difficulty for the 18 different Black-sounding names in the data. Table 7 shows that four of the five most familiar/fluent names (Ebony, Kenya, Leroy, Jermaine) have relatively high callback rates of 9.62, 8.67, 9.38, and 9.62 percent, while four of the five least familiar/fluent names (Aisha, Keisha, Tanisha, Lakisha) have relatively low callback rates of 2.22, 3.83, 5.8, and 5.5 percent. While there are some exceptions to this general relationship, the overall correlation coefficient is -0.49, which is statistically significant with a p-value of 0.04.

An alternative interpretation discussed by Bertrand and Mullainathan (2004) is that names may convey information about a candidate other than just their race, particularly their socioeconomic background. This hypothesis would imply that those with names that convey a worse social background would have lower callback rates. But these authors show the opposite finding: across the set of Black-sounding names, mother's education and callback rates are actually negatively correlated.[24] And in terms of our study of name pronunciation, we find no significant relationship (correlation coefficient of -0.06) between name difficulty and mother's education across the set of African-American names, so there is not strong support for the social background theory in our context.

4.2 Job Applicants from India, Pakistan, and China

While Bertrand and Mullainathan's study concentrated on White and African-American applicants to blue-collar jobs in the United States, Oreopoulos (2011) conducts an audit study in Canada on skilled immigrant workers, with a particular focus on those with Indian, Pakistani, and Chinese backgrounds. In his data, all jobs required a bachelor's degree or higher and all job applicants possessed at least this level of education, so this data set acts

[24] As defined in Bertrand and Mullainathan (2004), mother's education for a given first name is based on the percent of babies born under that first name in the state of Massachusetts between 1970 and 1986 whose mother had at least completed a high school degree.

as a nice complement to the one from Bertrand and Mullainathan (2004).

We begin by estimating a probit regression of callback rates, where the only independent variables are a set of indicators for various backgrounds based on ethnicity and national origin. In column 1 of Table 8, we see that relative to Canadians with English sounding names, all other groups have lower callback rates, ranging from a decrease of 2.4 percentage points for applicants from Britain, to a decrease of 5.7 percentage points for Pakistani applicants. These are very large effects, given that the average callback rate for all job seekers is roughly 10 percent. In column 2, we see that name fluency matters, with a one standard deviation increase in difficulty corresponding to a decrease in callback rate of 1.4 percentage points. When including controls for both ethnicity and name difficulty together, we see in column 3 that both remain significant, though the magnitudes of the coefficients on the ethnicity indicators are decreased by between 10 to 25 percent. Once again, there is suggestive evidence that racial discrimination based on one's name may be partly working through the difficulty of pronouncing (and potentially processing and remembering) that name. After controlling for other resume characteristics in column 4,[25] the coefficient on name difficulty now has a p-value of just over 0.1. Column 5 shows that difficulty of first and last names are both negatively related to getting called for a job, but these coefficients are not statistically significant after the inclusion of the other control variables.

Next, we restrict the sample to foreign applicants with ethnically Indian, Pakistani, or Chinese names, to see if name-based discrimination holds within these particular groups of job seekers. In Table A13 in the Online Appendix, we first show results for samples of each of these ethnic groups separately. Name difficulty is negatively related to callback rates for each separate group, with a one standard deviation increase in difficulty associated with a decrease in callback rates of 0.6, 1.2, and 3.1 percentage points for Indian, Pakistani, and Chinese applicants, respectively, though none of the coefficients is precisely estimated.

[25] These resume characteristics closely follow those employed in Oreopoulos (2011) that include gender, quality of undergraduate programs, extracurricular activities, fluency in foreign languages, master's degrees, additional credentials, prior work experience, list of references, and work authorization in Canada.

When we combine these three groups together and include indicators to control for ethnic background (shown in column 1 of Table 9), the coefficient on name difficulty is negative and significant at the 10 percent level.

Similar to our analysis for economics doctoral candidates and for job applicants with African-American sounding names, we also explore the idea that discrimination based on name fluency may depend on the quality of an applicant. Here, the results share similar patterns. Analogous to Bertrand and Mullainathan (2004), we use the predicted callback rate as a proxy for resume quality and consider the sample of Indian, Pakistani, or Chinese candidates who are in the top quarter of predicted callback rates (column 2) and their counterparts within the remainder of the sample. The coefficient on name difficulty is actually positive in column 2 and negative and significant in column 3, and the difference in coefficients is statistically significant at the 5% level, suggesting that Indian, Pakistani, and Chinese candidates with low quality resumes suffer from much greater name-based discrimination relative to their counterparts with high quality resumes. This result echoes a similar finding using African-American names from Bertrand and Mullainathan (2004). We subsequently also consider partitioning the sample based on whether the candidate holds a master's degree or not. As shown in columns 4 and 5, there is greater discrimination due to name fluency among Indian, Pakistani, and Chinese applicants without a Master's degree, relative to those with a Master's degree, though the difference in regression coefficients is not statistically significant in this case.

Finally, analogous to our analysis of African-American names, we calculate the correlation between callback rates and the standardized algorithmic ratings of name difficulty for the 24 different ethnic names associated with applicants from India, Pakistan, and China. Table 10 shows that within each of these groups, the rating of name difficulty is negatively related to callback rates, with the correlations between name difficulty and callback rates of -0.680, -0.588, and -0.338, for the sample of Indian, Pakistani, and Chinese names, respectively. The correlation coefficient is statistically significant at the 10 percent level for the eight

Indian names, though not for the other two groups. When we pool these three ethnic groups together, the correlation coefficient is -0.594, which has a p-value of less than 0.01 for a sample of 24 observations. There are significantly higher callback rates for easier to pronounce names such as Tara Singh (10.3 percent) and Min Liu (11.3 percent), compared to more difficult names such as Rabab Saeed (4.3 percent) and Xiuying Zhang (7.4 percent).

5 Conclusions

In this paper, we analyze two recent cohorts of economics PhD candidates and find strong evidence for labor market discrimination against individuals with names that are hard to pronounce. Job candidates with difficult-to-pronounce name are less likely to be placed into an academic job or to land a tenure track position, and also are placed in jobs at lower ranked institutions, as measured by research productivity. These results are statistically significant and economically large in magnitude. A one standard deviation increase in the difficulty of a candidate's full name lowers the likelihood of obtaining an academic (tenure track) job by somewhere between 4 to 7 (2 to 5) percentage points, depending on the measure used, and results in placing in an institution that is ranked more than 80 spots lower on RePEc's global ranking. We obtain strong and consistent results across three ways of measuring name fluency (an algorithmic score based on frequency of letter and phoneme combinations, a measure based on median time for survey takers to pronounce a name, and a purely subjective rating), and the results persist after the inclusion of a comprehensive set of control variables, including fixed effects for PhD institution and country of undergraduate degree. The effect of name fluency on job outcomes does not seem to be driven by the endogeneity of name changes, as there is not a significant difference in the magnitude of effects for first and last names, and the effects are similar for both male and female candidates.

As a test of the validity of our results in other settings, we also employ experimental data from two prior audit studies and find that labor market discrimination based on the perceived race or ethnicity of a candidate can be partially explained by the difficulty of a job applicant's name. Using Bertrand and Mullainathan's (2004) data with randomly assigned White-sounding and Black-sounding names, we find that even after accounting for perceived race, people with more difficult-to-pronounce names receive fewer callbacks for job listings in Boston and Chicago. And within the set of Black applicants, those with more difficult names are less likely to get called. We obtain similar results using data on applicants to jobs in Canada (Oreopoulos, 2011). Once again, controlling for ethnic background, having a less fluent name corresponds to a lower probability of being called back for a job. And within the set of ethnically Indian, Pakistani, and Chinese applicants, name difficulty is negatively related to callback rates.

Despite data limitations, our study also explores the potential mechanisms through which discrimination based on name pronunciation may occur. For job searches at academic, governmental, and research institutions, an initial screening generally involves committees getting together to discuss names of potential candidates, which may lead to some subconscious discrimination against names that are harder to pronounce and/or remember. This may also occur in the settings of prior audit studies, where recruiters must decide which applicants to call for potential interviews. Another possibility is that there are mental costs to processing and remembering less fluent names, and that these mental costs are only worth spending on higher quality candidates or when the jobs carry significant stakes. Additional analysis from both observational and experimental data seems to support this explanation. In particular, we find that PhD job market candidates with relatively weak profiles or from non-top PhD programs are more likely to suffer from name penalty in their search for academic positions. Likewise, we document similar patterns from separate analyses of Black Bertrand and Mullainathan (2004) and Indian, Pakistani, and Chinese Oreopoulos (2011) applicants: those who are less qualified and have weaker resumes tend to encounter much greater discrimination due to name pronunciation than those who are more qualified and have stronger resumes.

For our analysis of actual labor market data, we have only investigated this type of discrimination in the context of one specific labor market, but one might expect that the level of discrimination due to name fluency would be greater in other labor markets for several reasons. For jobs that require a doctoral degree in economics, search committee members will generally consist of other PhD economists, so the level of education for individuals reading job applications will be much higher than in many other labor markets. The selection process also extends far beyond just resume screenings, which are typical of the empirical contexts in prior audit studies. In our setting, committees have many different ways to evaluate candidates, including recommendation letters, interviews, and job talks. In addition, faculty members at colleges and universities have great exposure to people from many different countries and cultural backgrounds, which would increase the likelihood of encountering less popular names. Finally, many universities, research institutes, and governmental agencies have now incorporated anti-bias training that is not as typical in markets for other workers. Nonetheless, we still find a substantial penalty for candidates with names that are difficult to pronounce in the market for doctoral candidates in economics, providing a potential lower bound for the size of this phenomenon in other settings.

One implication of our academic labor market findings is that candidates with difficult-to-pronounce names may initially be under-placed, and then subsequently outperform (relative to their initial placement) in terms of research output. Future work could test to see whether there is greater likelihood of job switching for these individuals. A practical implication for search committees and recruiters is to be cognizant of another potential source of bias that exists in the hiring process. While it is impossible to entirely eliminate all sources of bias and labor market discrimination, our analysis suggests that with sufficient attention and effort, the bias against difficult-to-pronounce names can be minimized. The fact that this type of bias is much weaker for those with the highest quality resumes suggests that this bias can at least be partially overcome by those who choose to spend the mental energy and cost of paying attention to and remembering less familiar/fluent names.

References

Ahmad, A. (2020). "When the name matters: An experimental investigation of ethnic discrimination in the Finnish labor market." *Sociological Inquiry*, *90*(3), 468–496.

Ahmed, A. M., Andersson, L., and Hammarstedt, M. (2013). "Are gay men and lesbians discriminated against in the hiring process?" *Southern Economic Journal*, *79*(3), 565–585.

Ahmed, A. M., and Hammarstedt, M. (2008). "Discrimination in the rental housing market: A field experiment on the internet." *Journal of Urban Economics*, *64*(2), 362–372.

Alter, A. L., and Oppenheimer, D. M. (2006). "Predicting short-term stock fluctuations by using processing fluency." *Proceedings of the National Academy of Sciences*, *103*(24), 9369–9372.

Athey, S., Katz, L. F., Krueger, A. B., Levitt, S., and Poterba, J. (2007). "What does performance in graduate school predict? Graduate economics education and student outcomes." *American Economic Review*, *97*(2), 512–528.

Belot, M., Bhaskar, V., and Van De Ven, J. (2012). "Beauty and the sources of discrimination." *Journal of Human Resources*, *47*(3), 851–872.

Belot, M., and Schröder, M. (2022). "Remember me? The role of gender and racial attributes in memory." Working paper.

Bertrand, M., and Mullainathan, S. (2004). "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination." *American Economic Review*, *94*(4), 991–1013.

Biavaschi, C., Giulietti, C., and Siddique, Z. (2017). "The economic payoff of name Americanization." *Journal of Labor Economics*, *35*(4), 1089–1116.

Booth, A. L., Leigh, A., and Varganova, E. (2012). "Does ethnic discrimination vary across minority groups? Evidence from a field experiment." *Oxford Bulletin of Economics and Statistics*, *74*(4), 547–573.

Carlsson, M., and Eriksson, S. (2019). "Age discrimination in hiring decisions: Evidence from a field experiment in the labor market." *Labour Economics*, *59*, 173–183.

Carpenter, C. S. (2007). "Revisiting the income penalty for behaviorally gay men: Evidence from NHANES III." *Labour Economics*, *14*(1), 25–34.

Coupé, T. (2003). "Revealed performances: Worldwide rankings of economists and economics departments, 1990-2000." *Journal of the European Economic Association*, *1*(6), 1309–1345.

Cross, H., Kenney, G., Mell, J., and Zimmermann, W. (1990). *Employer Hiring Practices: Differential Treatment of Hispanic and Anglo Job Seekers*. Urban Institute Press.

Elmslie, B., and Tebaldi, E. (2007). "Sexual orientation and labor market discrimination." *Journal of Labor Research*, *28*(3), 436–453.

Ge, Q., Wu, S., and Zhou, C. (2021). "Sharing common roots: Student-graduate committee matching and job market outcomes." *Southern Economic Journal*, *88*(2), 828–856.

Goldin, C., and Rouse, C. (2000). "Orchestrating impartiality: The impact of 'blind' auditions on female musicians." *American Economic Review*, *90*(4), 715–741.

Hamermesh, D. S., and Biddle, J. E. (1994). "Beauty and the labor market." *American Economic Review*, *84*(5), 1174–1194.

Jacquemet, N., and Yannelis, C. (2012). “Indiscriminate discrimination: A correspondence test for ethnic homophily in the Chicago labor market.” *Labour Economics*, *19*(6), 824–832.

Kang, S. K., DeCelles, K. A., Tilcsik, A., and Jun, S. (2016). “Whitened résumés: Race and self-presentation in the labor market.” *Administrative Science Quarterly*, *61*(3), 469–502.

Krueger, A. B., and Wu, S. (2000). “Forecasting job placements of economics graduate students.” *Journal of Economic Education*, *31*(1), 81–94.

Laham, S. M., Koval, P., and Alter, A. L. (2012). “The name-pronunciation effect: Why people like Mr. Smith more than Mr. Colquhoun.” *Journal of Experimental Social Psychology*, *48*(3), 752–756.

Mobius, M. M., and Rosenblat, T. S. (2006). “Why beauty matters.” *American Economic Review*, *96*(1), 222–235.

Neumark, D., Bank, R. J., and Van Nort, K. D. (1996). “Sex discrimination in restaurant hiring: An audit study.” *Quarterly Journal of Economics*, *111*(3), 915–941.

Neumark, D., Burn, I., and Button, P. (2019). “Is it harder for older workers to find jobs? New and improved evidence from a field experiment.” *Journal of Political Economy*, *127*(2), 922–970.

Newman, E. J., Sanson, M., Miller, E. K., Quigley-McBride, A., Foster, J. L., Bernstein, D. M., and Garry, M. (2014). “People with easier to pronounce names promote truthiness of claims.” *PloS One*, *9*(2), e88671.

Oreopoulos, P. (2011). “Why do skilled immigrants struggle in the labor market? A field experiment with thirteen thousand resumes.” *American Economic Journal: Economic Policy*, *3*(4), 148–71.

Pierné, G. (2013). “Hiring discrimination based on national origin and religious closeness: Results from a field experiment in the Paris area.” *IZA Journal of Labor Economics*, *2*(1), 1–15.

Reimers, C. W. (1983). “Labor market discrimination against Hispanic and black men.” *Review of Economics and Statistics*, 570–579.

Riach, P. A., and Rich, J. (2006). “An experimental investigation of sexual discrimination in hiring in the English labor market.” *Advances in Economic Analysis & Policy*, *5*(2).

Riach, P. A., and Rich, J. (2010). “An experimental investigation of age discrimination in the English labor market.” *Annals of Economics and Statistics/Annales d’économie et de Statistique*, 169–185.

Song, H., and Schwarz, N. (2008). “If it’s hard to read, it’s hard to do: Processing fluency affects effort prediction and motivation.” *Psychological Science*, *19*(10), 986–988.

Sullivan, R. S., Dubnicki, A., and Dutkowsky, D. H. (2018). “Research, teaching, and ‘other’: What determines job placement of economics Ph.Ds?.” *Applied Economics*, *50*(32), 3477–3492.

Sweeney, L. (2013). “Discrimination in online ad delivery: Google ads, black names and white names, racial discrimination, and click advertising.” *Queue*, *11*(3), 10–29.

Turner, M. A., Fix, M., and Struyk, R. J. (1991). *Opportunities Denied, Opportunities Diminished: Racial Discrimination in Hiring*. Urban Institute Press.

Wright, B. R., Wallace, M., Bailey, J., and Hyde, A. (2013). “Religious affiliation and hiring discrimination in New England: A field experiment.” *Research in Social Stratification and Mobility*, *34*, 111–126.

Zhao, X., and Biernat, M. (2017). ""Welcome to the US" but "change your name"? Adopting Anglo names and discrimination." *Journal of Experimental Social Psychology*, *70*, 59–68.

Table 1: Summary Statistics ($N = 1,510$)

	(1) Mean	(2) Median	(3) Std. Dev.	(4) Min	(5) Max
CANDIDATE CHARACTERISTICS					
Male	0.697	1	0.460	0	1
US Undergrad	0.415	0	0.493	0	1
Elite US Undergrad	0.089	0	0.284	0	1
Prior Grad Degree	0.552	1	0.497	0	1
PhD Rank	32.934	27	25.812	1	90
No. of RRs	0.160	0	0.465	0	5
No. of Top 5 RRs	0.022	0	0.151	0	2
No. of Pubs	0.674	0	1.303	0	17
No. of Top 5 Pubs	0.011	0	0.102	0	1
Teaching Award	0.206	0	0.405	0	1
Instructor Exp	0.521	1	0.500	0	1
ADVISOR CHARACTERISTICS					
No. of Top 5 Pubs	5.584	3	7.826	0	61
Top 5 Journal Editor Exp	0.125	0	0.330	0	1
Editor Exp	0.784	1	0.412	0	1
PLACEMENT OUTCOMES					
Academia	0.647	1	0.478	0	1
Tenure Track	0.460	0	0.499	0	1
Visiting	0.078	0	0.269	0	1
Postdoc	0.186	0	0.311	0	1
Government/Think Tank	0.152	0	0.359	0	1
Industry	0.201	0	0.401	0	1
US Job	0.701	1	0.458	0	1
Imputed RePEc Ranking	779.852	1,000	361.874	1	1,000
NAME DIFFICULTY					
Algorithm Rating	55.727	55.670	45.194	0.210	185.770
Pronunciation Time	2.525	2.429	0.537	1.539	5.818
Subjectively Difficult	0.286	0	0.452	0	1

Notes: The sample is based on PhD candidates from the 96 top-ranked economics doctoral programs who entered the job market in either 2016-17 and 2017-18 job market cycles. The raw data contain information about $1,660$ job market candidates.

Table 2: Name Fluency and Placement Outcomes – Algorithm Rating

	(1) Academia	(2) Academia	(3) TT	(4) TT	(5) RePEc_Imputed	(6) RePEc_Imputed
<u>CANDIDATE CHARACTERISTICS</u>						
Male	-0.003	-0.002	-0.021	-0.020	-16.274	-16.122
	(0.029)	(0.029)	(0.031)	(0.031)	(57.935)	(57.935)
Elite US Undergrad	0.052	0.051	0.044	0.043	-48.135	-44.690
	(0.045)	(0.046)	(0.050)	(0.050)	(87.957)	(87.725)
Prior Grad Degree	0.097***	0.097***	0.082**	0.082**	-135.797**	-135.187**
	(0.031)	(0.031)	(0.033)	(0.033)	(62.335)	(62.546)
No. of RRs	0.067*	0.068*	0.110***	0.112***	-142.150**	-142.740**
	(0.035)	(0.035)	(0.034)	(0.034)	(57.315)	(57.570)
No. of Top 5 RRs	0.147	0.151	0.129	0.132	-287.441**	-290.038**
	(0.108)	(0.108)	(0.105)	(0.106)	(133.806)	(134.798)
No. of Pubs	0.028**	0.029**	0.025**	0.026**	-26.166	-26.682
	(0.013)	(0.012)	(0.012)	(0.012)	(19.829)	(19.901)
No. of Top 5 Pubs	0.271***	0.274***	0.147	0.148	-159.884	-156.518
	(0.069)	(0.067)	(0.129)	(0.130)	(164.138)	(165.978)
Teaching Award	0.076**	0.077**	0.061	0.063*	-116.441*	-116.980*
	(0.034)	(0.034)	(0.037)	(0.038)	(63.973)	(64.017)
Instructor Exp	0.010	0.010	-0.010	-0.010	65.549	64.713
	(0.032)	(0.032)	(0.033)	(0.033)	(60.724)	(60.814)
<u>ADVISOR CHARACTERISTICS</u>						
No. of Top 5 Pubs	0.003	0.004	-0.000	-0.000	-1.774	-1.877
	(0.003)	(0.003)	(0.002)	(0.002)	(3.950)	(3.956)
Top 5 Journal Editor Exp	-0.048	-0.051	0.019	0.015	-9.411	-8.115
	(0.053)	(0.053)	(0.052)	(0.052)	(78.483)	(78.447)
Editor Exp	0.095***	0.094***	0.047	0.046	11.593	11.954
	(0.036)	(0.036)	(0.036)	(0.036)	(74.406)	(74.419)
<u>NAME DIFFICULTY</u>						
Algorithm Rating: Full Name	-0.040**		-0.019		82.771***	
	(0.016)		(0.017)		(31.479)	
Algorithm Rating: First Name		-0.037**		-0.018		56.701*
		(0.017)		(0.018)		(32.596)
Algorithm Rating: Last Name		-0.020		-0.009		68.384*
		(0.019)		(0.019)		(36.238)
Observations	1,469	1,469	1,499	1,499	1,510	1,510
Control for Name Length	Yes	Yes	Yes	Yes	Yes	Yes
Sub-Field/Program FE	Yes	Yes	Yes	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The coefficients in columns 1-4 are marginal effects of probit regressions. The dependent variable in columns 1-2 (3-4) is a dichotomous variable for being placed in an academic (tenure track) position. Columns 5-6 are estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of $1,000$, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Name Fluency and Placement Outcomes – Pronunciation Time and Subjective Rating

	(1) Academia	(2) Academia	(3) TT	(4) TT	(5) RePEc_Imputed	(6) RePEc_Imputed
Pronunciation Time: Full Name	-0.074***		-0.047**		82.734**	
	(0.018)		(0.018)		(33.930)	
Subjectively Difficult Name		-0.116***		-0.055*		82.110
		(0.033)		(0.033)		(60.897)
Observations	1,469	1,469	1,499	1,499	1,510	1,510
Control for Name Length	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Sub-Field/Program FE	Yes	Yes	Yes	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The coefficients in columns 1-4 are marginal effects of probit regressions. The dependent variable in columns 1-2 (3-4) is a dichotomous variable for being placed in an academic (tenure track) position. Columns 5-6 are estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of $1,000$, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: Name Fluency and Placement Outcomes by Candidate Profiles

	High Predicted RePEc			Low Predicted RePEc		
	(1)	(2)	(3)	(4)	(5)	(6)
	Academia	TT	RePEc_Imputed	Academia	TT	RePEc_Imputed
Algorithm Rating: Full Name	-0.025	-0.032	95.528***	-0.064***	-0.014	81.147
	(0.023)	(0.025)	(36.592)	(0.024)	(0.023)	(64.481)
Observations	729	742	756	702	732	754
Control for Name Length	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Sub-Field/Program FE	Yes	Yes	Yes	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes	Yes	Yes	Yes

	Top 20 Program			Non-Top 20 Program		
	(7)	(8)	(9)	(10)	(11)	(12)
	Academia	TT	RePEc_Imputed	Academia	TT	RePEc_Imputed
Algorithm Rating: Full Name	-0.020	-0.025	96.866**	-0.055***	-0.017	58.089
	(0.026)	(0.028)	(41.790)	(0.021)	(0.022)	(48.562)
Observations	594	594	594	869	905	916
Control for Name Length	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Sub-Field/Program FE	Yes	Yes	Yes	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The coefficients in columns 1-2, 4-5, 7-8, and 10-11 are marginal effects of probit regressions. The dependent variable in columns 1, 4, 7, and 10 (2, 5, 8, and 11) is a dichotomous variable for being placed in an academic (tenure track) position. Columns 3, 6, 9, and 12 are estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of 1,000, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. In the top panel, high and low predicted RePEc rankings correspond to being above and below the median predicted RePEc ranking of placement institutions, respectively. In the bottom panel, program rankings are based on the U.S. News & World Report (USNWR) in 2017. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Name Fluency and Callback Rates: Experimental Data from Bertrand and Mullainathan (2004)

	(1) Callback	(2) Callback	(3) Callback	(4) Callback
Black	-0.032***		-0.018*	-0.015
	(0.008)		(0.011)	(0.010)
Female				0.007
				(0.010)
Education				0.006
				(0.007)
Number of Jobs on Resume				-0.001
				(0.005)
Years of Experience				0.003***
				(0.001)
Honors				0.042
				(0.028)
Volunteering Experience				-0.006
				(0.012)
Military Experience				-0.002
				(0.017)
Working in School				-0.007
				(0.010)
Listing Email				0.019
				(0.013)
Computer Skills				-0.015
				(0.016)
Special Skills				0.047***
				(0.011)
First Name Length				0.004
				(0.003)
Algorithm Rating: First Name		-0.017***	-0.011*	-0.012**
		(0.004)	(0.006)	(0.006)
Observations	4,870	4,870	4,870	4,794
Occupation Dummies	No	No	No	Yes

Notes: The sample is derived from publicly available replication data for Bertrand and Mullainathan (2004). The reported coefficients are marginal effects of probit regressions. The dependent variable is a dichotomous variable for receiving a callback. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 6: Name Fluency and Callback Rates: Experimental Data from Bertrand and Mullainathan (2004) – Sample of Black Applicants

	Black Applicants		High Quality Resume		Low Quality Resume	
	(1)	(2)	(3)	(4)	(5)	(6)
	Callback	Callback	Callback	Callback	Callback	Callback
Algorithm Rating: First Name	-0.011**	-0.018***	-0.004	-0.009	-0.017***	-0.025***
	(0.005)	(0.006)	(0.007)	(0.009)	(0.007)	(0.008)
Observations	2,435	2,181	1,223	1,058	1,212	1,071
Occupation Dummies	No	Yes	No	Yes	No	Yes
Control for Resume Characteristics	No	Yes	No	Yes	No	Yes
Control for First Name Length	No	Yes	No	Yes	No	Yes

Notes: The sample is derived from publicly available replication data for Bertrand and Mullainathan (2004). All specifications in this table focus on the subsample of black job applicants. The reported coefficients are marginal effects of probit regressions. The dependent variable is a dichotomous variable for receiving a callback. High quality and low quality resumes correspond to those ranked above and below the median predicted callback rate, respectively. Robust standard errors are in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 7: Name Fluency, Callback Rates, and Socioeconomic Background: Experimental Data from Bertrand and Mullainathan (2004) – Correlations Within Black Applicants

Name	Name Difficulty	Percent Callback	Mother's Education
Ebony	-0.973	9.62	65.6
Kenya	-0.973	8.67	70.2
Leroy	-0.523	9.38	53.3
Tyrone	-0.361	5.33	64.0
Jermaine	0.004	9.62	57.5
Jamal	0.153	6.56	73.9
Tremayne	0.200	4.35	–
Tamika	0.297	5.47	61.5
Darnell	0.675	4.76	66.1
Rasheed	0.770	2.99	77.3
Latonya	0.826	9.13	31.3
Hakim	0.970	5.45	73.7
Kareem	1.038	4.69	67.4
Aisha	1.148	2.22	77.2
Keisha	1.547	3.83	68.8
Latoya	1.549	8.41	55.5
Tanisha	1.839	5.80	64.0
Lakisha	2.161	5.50	55.6
Average	0.575	6.21	63.7
Correlation:	Name Difficulty & Callback	-0.488 [0.040]	
	Mother's Education & Callback	-0.609 [0.007]	
	Name Difficulty & Mother's Education	-0.055 [0.828]	

Notes: The sample is derived from publicly available replication data for Bertrand and Mullainathan (2004). The correlations reported in this table focus on the subsample of black job applicants. P-values for correlations are in brackets.

Table 8: Name Fluency and Callback Rates: Experimental Data from Oreopoulos (2011)

	(1) Callback	(2) Callback	(3) Callback	(4) Callback	(5) Callback
Female				0.018***	0.018***
				(0.005)	(0.005)
Top 200 World Ranking University				-0.003	-0.003
				(0.005)	(0.005)
Listing Extracurricular Activities				-0.002	-0.002
				(0.005)	(0.005)
Fluent in French & Other Languages				0.019***	0.018***
				(0.006)	(0.006)
Master's Degree				0.006	0.006
				(0.007)	(0.007)
High Quality Work Experience				0.009	0.009
				(0.006)	(0.006)
Additional Required Credentials				0.041***	0.041***
				(0.010)	(0.010)
Listing Canadian References				-0.029*	-0.029*
				(0.015)	(0.015)
Accreditation of Foreign Education				-0.012	-0.012
				(0.014)	(0.014)
Permanent Resident				-0.007	-0.007
				(0.015)	(0.015)
Indian	-0.046***		-0.036***	-0.035***	-0.034***
	(0.006)		(0.008)	(0.009)	(0.009)
Pakistani	-0.057***		-0.049***	-0.049***	-0.049***
	(0.007)		(0.009)	(0.009)	(0.009)
Chinese	-0.041***		-0.038***	-0.035***	-0.033***
	(0.006)		(0.006)	(0.010)	(0.011)
Chinese Canadian	-0.053***		-0.053***	-0.050***	-0.049***
	(0.006)		(0.006)	(0.007)	(0.007)
Greek	-0.031**		-0.018	-0.018	-0.022
	(0.012)		(0.015)	(0.016)	(0.018)
British	-0.024***		-0.024***	-0.023**	-0.023**
	(0.009)		(0.009)	(0.009)	(0.009)
Name Length				0.000	0.000
				(0.002)	(0.002)
Algorithm Rating: Full Name		-0.014***	-0.008**	-0.007	
		(0.003)	(0.004)	(0.004)	
Algorithm Rating: First Name					-0.006
					(0.004)
Algorithm Rating: Last Name					-0.002
					(0.004)
Observations	12,910	12,910	12,910	12,910	12,910

Notes: The sample is derived from publicly available replication data for Oreopoulos (2011). The reported coefficients are marginal effects of probit regressions. The dependent variable is a dichotomous variable for receiving a callback. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 9: Name Fluency and Callback Rates: Experimental Data from Oreopoulos (2011) – Sample of Foreign Ethnic Applicants

	Ind/Pak/Chin	High Quality	Low Quality	Master's	No Master's
	(1)	(2)	(3)	(4)	(5)
	Callback	Callback	Callback	Callback	Callback
Algorithm Rating: Full Name	-0.007*	0.036*	-0.009**	-0.003	-0.008*
	(0.004)	(0.020)	(0.004)	(0.010)	(0.005)
Observations	7,158	467	6,705	1,119	6,029
Control for Name Length	Yes	Yes	Yes	Yes	Yes
Control for Ethnicity	Yes	Yes	Yes	Yes	Yes
Control for Resume Characteristics	Yes	Yes	Yes	Yes	Yes

Notes: The sample is derived from publicly available replication data for Oreopoulos (2011). All specifications in this table focus on the subsample of job applicants with ethnically Indian, Pakistani, and Chinese names. The reported coefficients are marginal effects of probit regressions. The dependent variable is a dichotomous variable for receiving a callback. High quality and low quality resumes correspond to those ranked above and below the median predicted callback rate, respectively. Robust standard errors are in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 10: Name Fluency and Callback Rates: Experimental Data from Oreopoulos (2011) - Correlations Within Foreign Ethnic Applicants

Name	Name Difficulty	Percent Callback
Indian		
Tara Singh	-0.603	10.29
Maya Kumar	-0.538	8.66
Shreya Sharma	0.348	9.54
Arjun Kumar	0.742	7.82
Samir Sharma	0.985	8.59
Panav Singh	1.264	8.25
Rahul Kaur	1.913	8.14
Priyanka Kaur	2.557	7.61
Average:	0.834	8.613
Correlation:		-0.680 [0.064]
Pakistani		
Hassan Khan	-0.304	6.30
Fatima Sheikh	0.245	8.11
Sana Khan	0.392	8.82
Ali Saeed	0.705	8.33
Chaudhry Mohammad	1.102	6.12
Asif Sheikh	1.296	3.85
Hina Chaudhry	1.348	7.80
Rabab Saeed	3.142	4.26
Average:	0.991	6.699
Correlation:		-0.588 [0.125]
Chinese		
Na Li	-0.802	7.65
Min Liu	-0.671	11.34
Lei Li	-0.644	9.32
Tao Wang	-0.557	10.98
Dong Liu	-0.534	7.88
Fang Wang	-0.283	12.57
Yong Zhang	-0.279	8.60
Xiuying Zhang	1.511	7.42
Average:	-0.283	9.470
Correlation:		-0.338 [0.412]
Indian/Pakistani/Chinese		
Average:	0.514	8.260
Correlation:		-0.594 [0.002]

Notes: The sample is derived from publicly available replication data for Oreopoulos (2011). The correlations reported in this table focus on the subsample of job applicants with ethnically Indian, Pakistani, and Chinese names. P-values for correlations are in brackets.

Online Appendix

A. Instructions for Name Fluency Surveys

Thank you for agreeing to assist with research projects related to the pronunciation of names. I have designed a set of Qualtrics surveys which have a series of names for you to pronounce.

1. Before you start a particular survey, start an audio recording of yourself. Then, you will see a series of names for you to pronounce, with one name per screen. Read through the name, and then click the arrow to advance to the next screen to see the next name. Continue to repeat this until you have finished the survey. You may then stop the recording and save it. You will repeat this process for all of the different groups of names, though you may wish to do break up your work across several different times in the day or the week to complete the work.

2. Please complete a particular group in one sitting without taking any breaks in between. Once you complete that group, then feel free to take as long of a break as you need until you start the next survey, but again, please do not take breaks once you have started a new survey until you complete that one. Names will be separated in groups of approximately 50 (with some groups listed as first names and some groups listed as last names), so perhaps you may want to do a bunch at one time, with short breaks in between each of the individual surveys. Then, you can come back and do another chunk of them at another day/time when you are free.

3. If you are unsure of how to pronounce a particular name, simply do your best to make a guess or sound it out before you click the arrow to advance to the next screen. You should not search the internet to hear a recording of the name, but simply make an attempt at pronouncing it.

4. It is possible that you may see some names that are duplicates or are very similar to other names in one of the surveys, but please pronounce each of the names you see on

the screen even if you think you have seen that name before.

5. Please complete each survey only one time. To make sure that you do every survey only once, take careful notes about which ones you have completed and which ones you still need to complete. The most logical way would be to complete the surveys in numerical order (perhaps starting with the first names and then the last names).

6. Before you begin the surveys that contain the actual names in our database, please do the initial test survey, which is a short survey that will not be part of the actual database, but will allow you to see how the interface and the survey works.

7. Sharing your audio recordings: If possible, you may save these onto a google drive and share with me.

8. Calculating your work hours: You may include the short breaks you take in between individual surveys when you calculate your hours worked. For example, if you complete 10 surveys in 1.5 hours for one particular day that you are working, and you are taking a few minutes of breaks in between each of the ten, it is fine to include those short breaks as part of your time worked.

B. Additional Robustness Checks

In this section, we present a number of robustness checks on our baseline findings. We first consider an alternative placement quality measure based on the raw RePEc ranking that excludes private sector and non-tenure track academic placements. Tobit estimates based on this alternative RePEc measure are reported in Table A2 in the Online Appendix and are similar in direction and statistical significance compared to the full sample results with imputed RePEc values as shown in Tables 2 and 3, though the relevant coefficients on name fluency measures are smaller in magnitude, ranging from 70 for algorithmic ratings to 28 for subjective ratings.

Next, to assess the robustness of our specifications on placement quality, we estimate multinomial logit regressions for different types of placements and present the estimates in Table A3 in the Online Appendix.[26] We observe that relative to the reference group placement type of government or think tank jobs, the coefficient on name difficulty is significantly negative for being placed into academia, and this result is consistent across different name fluency measures.[27] On the other hand, in separate specifications reported in Table A4, when we further decompose academic job types and set the baseline category as visiting/postdoc, the coefficient on name difficulty for the tenure track category is not significant relative to the baseline. Taken together, this suggests that name fluency impacts the likelihood of being placed into academia relative to industry or government jobs, but does not affect the probability of obtaining a tenure track job, conditional on being placed in academia.

Analogously, to check the robustness of the results on placement quality, we also adopt an ordered probit model using categories of the imputed RePEc ranking of job placements as the outcome of interest. Given the ordinal nature of RePEc rankings, we categorize the ranking of imputed RePEC productivity index into the following five categories and estimate an ordered probit model: 1) $RePEc \leq 50$; 2) $50 < RePEc \leq 200$; 3) $200 < RePEc \leq 400$; 4) $400 < RePEc \leq 800$; and 5) $RePEc = 1,000$. The estimates on name fluency measures, as presented in Table A5 in the Online Appendix, are qualitatively similar to our main findings and again confirm that candidates with harder-to-pronounce names tend to be placed in institutions with lower research productivity.

A concern discussed in the main text is the possibility that name changes are endogenous. One possibility is that students who have advisors and committee members from the same country might be less likely to feel the need to Americanize/Anglicize their (first) names. Ge et al. (2021) document a beneficial impact of student-graduate committee matching, in

[26]Note that the multinomial logit specifications do not include program fixed effects due to non-convergence.

[27]The difference between the coefficients for academic and industry positions is statistically significant at the 10%, 1%, and 1% levels for name fluency measures based on algorithmic ratings, pronunciation time, and subjective ratings, respectively.

the form of country of origin and native language, on students' initial placement outcomes in the economics PhD job market, which could lead to a downward bias in the estimate of the magnitude of the name fluency effect. To address this issue, we re-estimate our baseline specifications and add controls for student-graduate committee matching based on country (U.S. vs. non-U.S.) or language (English vs. foreign language),[28] and the resulting estimates, as reported in Table A6 in the Online Appendix, remain identical to those in Tables 2 and 3.

C. Common Names

We also explore the possibility that the uniqueness or commonality of names may affect job outcomes. It is likely that those with very common names could be at a disadvantage because they do not stand out from other candidates. Because pronunciation difficulty is likely negatively correlated with commonality of names, our estimates of the name fluency effect might be underestimated. To address this concern, we augment our baseline specifications by controlling for having a common first name or common last name. Due to data constraints, we focus on common names in the U.S. Specifically, we code someone as having a very common name if their first name is among the 50 most common female first names or the 50 most common male first names according to the 1990 U.S. Census, and having a very common last name if their last name is among the 50 most common surnames according to the 2010 U.S. Census.[29]

We present the resulting estimates in Table A11 in the Online Appendix. As shown in columns 1-3 that focus on the full sample of job market candidates, none of the variables on name commonality (i.e., indicator for common first name, indicator for common last name, and their interaction) are statistically significant, and their inclusion does not impact

[28] Following Ge et al. (2021), we code "country match" as being equal to one when at least one of the student's committee members went to an undergraduate institution in the same country as the student's undergraduate institution. Similarly, we code "language match" as being equal to one when a student's country of origin has the same official language as that of at least one of the committee members.

[29] The 1990 and 2010 U.S. Census data respectively represent the most recent data sources for tabulations on common first and last names. The lists of common first and last names are available upon request.

the magnitude or significance of the name difficulty coefficient in any of our regressions. In addition, since the data sources for our common name analysis are based on the U.S. Census, we also conduct a subsample analysis by limiting the set of common names to job market candidates who are from U.S. and Canada. As shown in columns 4-6, the results on placement types and quality as well as the coefficients on common name indicators remain unaffected.

Table A1: Name Fluency and Placement Outcomes – Pronunciation Time and Subjective Rating by First and Last Names

	(1) Academia	(2) Academia	(3) TT	(4) TT	(5) RePEc_Imputed	(6) RePEc_Imputed
Pronunciation Time: First Name	-0.045***		-0.027*		58.527**	
	(0.016)		(0.016)		(29.752)	
Pronunciation Time: Last Name	-0.055***		-0.031		52.474	
	(0.020)		(0.020)		(38.977)	
Subjective Difficult: First Name		-0.084**		-0.022		133.809*
		(0.041)		(0.040)		(76.149)
Subjective Difficult: Last Name		-0.068*		-0.034		21.741
		(0.039)		(0.040)		(76.552)
Observations	1,469	1,469	1,499	1,499	1,510	1,510
Control for Name Length	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Sub-Field/Program FE	Yes	Yes	Yes	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The coefficients in columns 1-4 are marginal effects of probit regressions. The dependent variable in columns 1-2 (3-4) is a dichotomous variable for being placed in an academic (tenure track) position. Columns 5-6 are estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of 1,000, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A2: Name Fluency and Placement Quality (Raw RePEc Ranking) – Tobit Estimates

	(1) RePEc	(2) RePEc	(3) RePEc
Algorithm Rating: Full Name	70.426***		
	(25.006)		
Pronunciation Time: Full Name		51.069*	
		(26.546)	
Subjectively Difficult: Full Name			28.278
			(49.645)
Observations	910	910	910
Control for Name Length	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Sub-Field/Program FE	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes

Notes: The dependent variable across all specifications is the RePEc ranking of the institution of initial job placement, where individuals obtaining private sector jobs are excluded from the sample. All specifications are estimated using a tobit model censored with an upper limit of $1,000$. Robust standard errors are in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Table A3: Name Fluency and Placement Types – Multinomial Logit Estimates

	(1) Academia	(2) Industry
Algorithm Rating: Full Name	-0.225**	-0.087
	(0.088)	(0.108)
Observations	1,510	1,510
Control for Name Length	Yes	Yes
Other Controls	Yes	Yes
Sub-Field FE	Yes	Yes
Region/JM Cycle FE	Yes	Yes

	(3) Academia	(4) Industry
Pronunciation Time: Full Name	-0.251***	0.116
	(0.090)	(0.103)
Observations	1,510	1,510
Control for Name Length	Yes	Yes
Other Controls	Yes	Yes
Sub-Field FE	Yes	Yes
Region/JM Cycle FE	Yes	Yes

	(5) Academia	(6) Industry
Subjectively Difficult: Full Name	-0.340*	0.161
	(0.177)	(0.210)
Observations	1,510	1,510
Control for Name Length	Yes	Yes
Other Controls	Yes	Yes
Sub-Field FE	Yes	Yes
Region/JM Cycle FE	Yes	Yes

Notes: Each panel is estimated using a separate multinomial logit model with the dependent variable being a categorical variable capturing placement types, including academia, government/think tank, and industry (private sector). Government/think tank positions are the baseline category across all specifications. The reported coefficients are in log-odds. Program fixed effects are not included due to non-convergence. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A4: Name Fluency and Placement Types – Multinomial Logit Estimates: Alternative Placement Categories

	(1) TT	(2) Govt/Think Tank	(3) Industry
Algorithm Rating: Full Name	0.045	0.253**	0.166
	(0.089)	(0.109)	(0.105)
Observations	1,510	1,510	1,510
Control for Name Length	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Sub-Field FE	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes
	(4) TT	(5) Govt/Think Tank	(6) Industry
Pronunciation Time: Full Name	0.028	0.271**	0.387***
	(0.099)	(0.114)	(0.112)
Observations	1,510	1,510	1,510
Control for Name Length	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Sub-Field FE	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes
	(7) TT	(8) Govt/Think Tank	(9) Industry
Subjectively Difficult: Full Name	0.248	0.524**	0.683***
	(0.179)	(0.219)	(0.207)
Observations	1,510	1,510	1,510
Control for Name Length	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Sub-Field FE	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes

Notes: Each panel is estimated using a separate multinomial logit model with the dependent variable being a categorical variable capturing placement types, including tenure track, visiting/postdoc, government/think tank, and industry (private sector). Visiting/postdoc positions are the baseline category across all specifications. The reported coefficients are in log-odds. Program fixed effects are not included due to non-convergence. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A5: Name Fluency and Placement Quality – Ordered Probit Estimates

	(1) RePEc_Imputed	(2) RePEc_Imputed	(3) RePEc_Imputed
Algorithm Rating: Full Name	0.100**		
	(0.046)		
Pronunciation Time: Full Name		0.100**	
		(0.048)	
Subjectively Difficult: Full Name			0.106
			(0.087)
Observations	1,510	1,510	1,510
Control for Name Length	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Sub-Field/Program FE	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes

Notes: All specifications are estimated using an ordered probit model, where the dependent variable is based on the following ordered categories of the imputed RePEc research productivity index: 1) $RePEc \leq 50$; 2) $50 < RePEc \leq 200$; 3) $200 < RePEc \leq 400$; 4) $400 < RePEc \leq 800$; and 5) $RePEc = 1,000$. Robust standard errors are in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Table A6: Name Fluency and Placement Outcomes – Controlling for Advisor Match

	(1) Academia	(2) Academia	(3) TT	(4) TT	(5) RePEc_Imputed	(6) RePEc_Imputed
Algorithm Rating: Full Name	-0.041**	-0.040**	-0.020	-0.019	83.740***	82.891***
	(0.016)	(0.016)	(0.017)	(0.017)	(31.313)	(31.415)
	(7) Academia	(8) Academia	(9) TT	(10) TT	(11) RePEc_Imputed	(12) RePEc_Imputed
Pronunciation Time: Full Name	-0.074***	-0.074***	-0.047***	-0.046**	80.610**	80.334**
	(0.018)	(0.018)	(0.018)	(0.018)	(33.735)	(33.952)
	(13) Academia	(14) Academia	(15) TT	(16) TT	(17) RePEc_Imputed	(18) RePEc_Imputed
Subjectively Difficult: Full Name	-0.117***	-0.116***	-0.060*	-0.057*	82.380	83.393
	(0.033)	(0.033)	(0.033)	(0.033)	(61.258)	(60.958)
Observations	1,469	1,469	1,499	1,499	1,510	1,510
Control for Country Match	Yes	No	Yes	No	Yes	No
Control for Language Match	No	Yes	No	Yes	No	Yes
Control for Name Length	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Subfield/Program FE	Yes	Yes	Yes	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The coefficients in columns 1-4, 7-10, and 13-16 are marginal effects of probit regressions. The dependent variable in columns 1-2, 7-8, and 13-14 (3-4, 9-10 and 15-16) is a dichotomous variable for being placed in an academic (tenure track) position. Columns 5-6, 11-12, and 17-18 are estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of 1,000, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. The country/language match variables are indicator variables based on matching with at least one of the committee members. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A7: Name Fluency and Placement Outcomes – Country Fixed Effects

	(1) Academia	(2) TT	(3) RePEc_Imputed
Algorithm Rating: Full Name	-0.039**	-0.005	82.746***
	(0.017)	(0.018)	(31.692)
Observations	1,199	1,200	1,236
Control for Name Length	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Subfield/Program FE	Yes	Yes	Yes
Country/JM Cycle FE	Yes	Yes	Yes

Notes: The coefficients in columns 1 and 2 are marginal effects of probit regressions. The dependent variable in column 1 (2) is a dichotomous variable for being placed in an academic (tenure track) position. Column 3 is estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of 1,000, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A8: Correlations with Name Fluency Measures

	Algorithm Rating		Pronunciation Time		Subjectively Difficult	
	Corr.	P-Value	Corr.	P-Value	Corr.	P-Value
CANDIDATE CHARACTERISTICS						
Male	-0.111	[0.000]	-0.044	[0.091]	-0.074	[0.004]
US Undergrad	-0.284	[0.000]	-0.187	[0.000]	-0.263	[0.000]
Elite US Undergrad	-0.042	[0.099]	-0.019	[0.464]	-0.069	[0.008]
Prior Grad Exp	0.234	[0.000]	0.160	[0.000]	0.157	[0.000]
PhD Rank	0.044	[0.090]	-0.012	[0.631]	0.073	[0.004]
No. of RRs	-0.042	[0.105]	-0.058	[0.024]	-0.022	[0.395]
No. of Top 5 RRs	-0.041	[0.107]	-0.009	[0.738]	-0.043	[0.093]
No. of Pubs	-0.007	[0.785]	-0.025	[0.340]	-0.062	[0.016]
No. of Top 5 Pubs	-0.015	[0.567]	-0.004	[0.879]	-0.051	[0.047]
Teaching Award	-0.002	[0.941]	-0.065	[0.012]	-0.007	[0.781]
Instructor Exp	0.011	[0.672]	0.035	[0.170]	0.017	[0.506]

Notes: The sample is based on PhD candidates from the 96 top-ranked economics doctoral programs who entered the job market in either 2016-17 and 2017-18 job market cycles. The raw data contain information about $1,660$ job market candidates.

Table A9: Name Fluency and Placement Outcomes – Excluding Candidates With Ethnically Chinese Names

	(1) Academia	(2) TT	(3) RePEc_Imputed
Algorithm Rating: Full Name	-0.049***	-0.027	89.133**
	(0.019)	(0.019)	(35.818)
Observations	1,184	1,185	1,221
Control for Name Length	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Subfield/Program FE	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes

Notes: The sample excludes job market candidates with ethnically Chinese names, including candidates from China and Taiwan. The coefficients in columns 1 and 2 are marginal effects of probit regressions. The dependent variable in column 1 (2) is a dichotomous variable for being placed in an academic (tenure track) position. Column 3 is estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of 1,000, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A10: Name Fluency and Placement Outcomes by Gender

	(1) Academia	(2) TT	(3) RePEc_Imputed
Algorithm Rating × Male	-0.040	-0.021	45.668
	(0.028)	(0.029)	(56.840)
Observations	1,469	1,499	1,510
Control for Name Length	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Subfield/Program FE	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes

Notes: The coefficients in columns 1 and 2 are marginal effects of probit regressions. The dependent variable in column 1 (2) is a dichotomous variable for being placed in an academic (tenure track) position. Column 3 is estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of 1,000, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A11: Name Fluency and Placement Outcomes – Accounting for Common Names

	All Candidates			Candidates from U.S. and Canada		
	(1)	(2)	(3)	(4)	(5)	(6)
	Academia	TT	RePEc_Imputed	Academia	TT	RePEc_Imputed
Algorithm Rating: Full Name	-0.042**	-0.026	85.257***	-0.080***	-0.040	117.797**
	(0.016)	(0.017)	(32.017)	(0.029)	(0.028)	(54.321)
Common First Name	-0.005	-0.043	-51.067	-0.015	-0.030	-7.563
	(0.042)	(0.044)	(87.130)	(0.056)	(0.053)	(111.094)
Common Last Name	0.005	-0.079	64.360	0.047	-0.035	76.325
	(0.069)	(0.068)	(123.867)	(0.104)	(0.098)	(170.516)
Common First Name ×	-0.235	-0.181	386.031	-0.127	-0.167	340.241
Common Last Name	(0.168)	(0.137)	(307.099)	(0.213)	(0.155)	(349.644)
Observations	1,469	1,499	1,510	586	600	648
Control for Name Length	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Sub-Field/Program FE	Yes	Yes	Yes	Yes	Yes	Yes
Region/JM Cycle FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The coefficients in columns 1-2 and 4-5 are marginal effects of probit regressions. The dependent variable in columns 1 and 3 (2 and 5) is a dichotomous variable for being placed in an academic (tenure track) position. Columns 3 and 6 are estimated using a tobit model, with the dependent variable being the imputed RePEc ranking of the institution of initial job placement, where private sector jobs are given an imputed ranking of 1,000, the highest (worst) ranking. All tobit regressions are censored with an upper limit of $1,000$. Common first and last names are derived from the 1990 and 2010 U.S. Census, respectively. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A12: Name Fluency and Placement Quality within Tenure Track Placements – Tobit Estimates

	TT	TT & RePEc < 400
	(1)	(2)
	RePEc	RePEc
Algorithm Rating: Full Name	72.019***	0.591
	(26.678)	(9.393)
Observations	695	253
Control for Name Length	Yes	Yes
Other Controls	Yes	Yes
Sub-Field/Program FE	Yes	Yes
Region/JM Cycle FE	Yes	Yes

Notes: The sample is restricted to job market candidates placed in tenure track positions. The dependent variable across all specifications is the RePEc ranking of the institution of initial job placement, where individuals obtaining private sector jobs are excluded from the sample. Column 1 is estimated using a tobit model censored with an upper limit of $1,000$. Robust standard errors are in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Table A13: Name Fluency and Callback Rates: Experimental Data from Oreopoulos (2011) – Sample of Foreign Ethnic Applicants

	Indian	Pakistani	Chinese
	(1)	(2)	(3)
	Callback	Callback	Callback
Algorithm Rating: Full Name	-0.006	-0.012	-0.031
	(0.006)	(0.009)	(0.020)
Observations	3,312	952	2,848
Control for Name Length	Yes	Yes	Yes
Control for Resume Characteristics	Yes	Yes	Yes

Notes: The sample is derived from publicly available replication data for Oreopoulos (2011). All specifications in this table focus on the subsample of job applicants with ethnically Indian, Pakistani, and Chinese names. The reported coefficients are marginal effects of probit regressions. The dependent variable is a dichotomous variable for receiving a callback. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.