Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers

Raviv Murciano-Goroff

November, 2017

— Draft —

Motivation

- Tech sector workforce is highly gender imbalanced.
- Supply-side and demand-side factors might contribute.
- What can be done to improve diversity requires answering two questions:

1. Do gender differences in the behaviors of job seekers exist?
2. Do recruiters adjust based on such gender differences in ways that could increase the diversity of the job applicant pool?

This paper

- Focus on recruiting and initial screening stage.
- Data from a digital recruiting platform for software engineers
  - Candidates have profiles. Self-reported information.
  - Recruiters contact candidates.
- In addition, merge in data on actual previous coding work demonstrating technical skills for each candidate.

Questions

- What do tech recruiters look for in candidates?
  Self-reported skills?

- Do men and women with the same previous experience in a programming language self-report knowing that technical skill?

- Do recruiters adjust to gender differences in the propensity of candidates to self-report their skills?

Literature Review and Contribution

- Previous literature has focused on one-side or the other:
  - Supply side: Gneezy, Niederle, and Rustichini (2003), Niederle and Vesterlund (2007), Baldiga (2014), Bursztyn, Fujiwara, and Pallais (2017).
  - Demand side: Goldin and Rouse (1997), Riach and Rich (2002), Bertrand and Mullainathan (2004).
- Unlike an audit study, I can see actual job seekers' and recruiters' behavior in a real and competitive labor market.
- I can see how the two sides of the market related to each other in equilibrium.

Table of Contents

Table of Contents

The Setting

- Many tech recruiters use online hiring platforms for initial screening of potential recruits.
- Platforms show "profiles" with information about "candidates."
- The platform where my data comes from combines information that candidates supply with assessments of the candidates' technical skills by the platform.
    - Individuals posted their information on a partner site.
    - Recruiters viewed the "profile" with both the candidate written and the platform's assessment on this platform.

Raviv Murciano-Goroff

Education

- Ph.D. Stanford University (May 2018)
- MS.c. University of Oxford (Oct 2010)
- B.A. Harvard University (June 2009)

Work History

- Research Assistant at NBER (2010-2012)

Verified Languages

- Python - High Experience
- Ruby - Low Experience

Self-Reported Skills

- Public Speaking, Stata, Python, JavaScript.

Raviv Murciano-Goroff

Education

- Ph.D. Stanford University (May 2018)
- MS.c. University of Oxford (Oct 2010)
- B.A. Harvard University (June 2009)

Work History

- Research Assistant at NBER (2010-2012)

Verified Languages

- Python - High Experience
- Ruby - Low Experience

Self-Reported Skills

- Public Speaking, Stata, Python, JavaScript.

Written by the candidate

Raviv Murciano-Goroff

Education

- Ph.D. Stanford University (May 2018)
- MS.c. University of Oxford (Oct 2010)
- B.A. Harvard University (June 2009)

Work History

- Research Assistant at NBER (2010-2012)

Verified Languages

- Python - High Experience
- Ruby - Low Experience

Self-Reported Skills

- Public Speaking, Stata, Python, JavaScript.

Added by the platform

- Open Source
- StackOverflow

Search and Save

- Recruiters can search the profiles in four ways:
  1. Verified languages.
  2. Schools attended.
  3. Previous employers.
  4. Geography.

- Recruiters cannot search based on self-reported languages or gender.

- The order of the results is not based on self-reported languages or gender.

- Profiles have a button to “save” a candidate.

“Profiles” Dataset

- Cross-section of profiles on the platform in December 2015.
  - Only those in the U.S. with either a BA in CS, a job involving coding, or self-reported skills related to software engineering (Git, “Computer programming”, “Machine Learning,” etc.)
- Includes 259 profile “attributes” that recruiters would see:
  - Educational credentials and work history.
  - Non-technical and social skills.
  - Technical skills and Programming languages.
- Observe if a recruiter “saved” a profile between 2014 and 2016.
  - Only use recruiters who are looking for software engineers.

"Profiles" Summary Stats

- 3,927,150 profiles of which 20.63% are female.
- Male and female candidates are similar in some ways...
  - Mean candidate graduated in 2001.
  - Approximately 19% have a BA in CS.
  - 2.7% attend college at a top-10 school for CS.
  - 13% of men and 15% of women completed master's degrees.
- and different in others...
  - 54.40% of men and 64.16% of women list at least one self-reported skill.
  - 1.776% of male and 0.654% of the female candidates were saved by at least one recruiter.

Programming Languages on Resumes

- Languages can appear in one of five different ways on a profile.
  - Self-reported or not.
  - Verified - Low, Verified - High, not verified.
- In 2015, the programming language JavaScript was...
  - The most popular language to do open source.
  - The most popular skill listed on job postings for software engineering and web development jobs in 2015 (Burning Glass data).
  - 263,422 candidates have JavaScript on their profile (either self-reported or Verified).

Tabulation of JavaScript's Appearance on Profiles

	Verified - High	Verified - Low	Not Verified
Self-Reported	0.28%	0.62%	3.75%
Not Self-Reported	0.23%	1.80%	93.31%

Note: 263,422 profiles have JavaScript listed as either "Self-Reported" or "Verified" or both. Each cell shows the fraction of those profiles that have the programming language JavaScript listed in either the "Self-Reported Skills" list, the "Verified Languages" list, or both lists.

- % of Verified - Low who Self-Report: 25.60%
- % of Verified - High who Self-Report: 54.72%

Table of Contents

Question #1: What on profiles do recruiters pay attention to?

- We would like to know what recruiters searched for and which attributes recruiters responded to:
  - Educational credentials, work history, technical skills, or soft skills?
  - The searches were not logged, so cannot directly observe.
- By predicting which candidates got saved, I can infer which attributes are correlated with greater demand.

Inferring What Recruiters Sought in Candidates

- Predict if a profile got saved based on all their attributes in a linear probability model.
  - An observation in the regression is one of the 3,927,150 candidate profiles.
  - The dependent variable is whether or not the candidate is saved between 2014 and 2016.
  - The covariates are the 259 attributes of a candidate.
- Compare the magnitudes of the estimated coefficients.
- Mean is 0.015.

	Saved
BA in CS=1	0.008***
	(0.000)
Current Employer Ranked Top 10 for Tech	0.049***
	(0.002)
SR Map/Reduce	0.053***
	(0.006)
JavaScript - SR, No V	0.005***
	(0.001)
JavaScript - No SR, V-Low	0.049***
	(0.001)
JavaScript - SR, V-Low	0.115***
	(0.003)
JavaScript - No SR, V-High	0.147***
	(0.005)
JavaScript - SR, V-High	0.244***
	(0.005)
Controls	Yes
Dep. Mean	0.015
R^2	0.222

Note: Observation is a profile. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Results

- Specific technical skills, including programming languages, predict more recruiter interest than general educational credentials or aspects of work history.
- Languages capture 88% of the variation of all observables.
- Only 0.6% of candidates with no verified languages got saved.
  - Recruiters likely searched based on programming languages and technical skills.

Evidence Self-Reporting Matters

- Saw evidence that self-reporting a programming language matters even when experience in that language is "verified."

	Saved
BA in CS=1	0.008***
	(0.000)
Current Employer Ranked Top 10 for Tech	0.049***
	(0.002)
SR Map/Reduce	0.053***
	(0.006)
JavaScript - SR, No V	0.005***
	(0.001)
JavaScript - No SR, V-Low	0.049***
	(0.001)
JavaScript - SR, V-Low	0.115***
	(0.003)
JavaScript - No SR, V-High	0.147***
	(0.005)
JavaScript - SR, V-High	0.244***
	(0.005)
Controls	Yes
Dep. Mean	0.015
R^2	0.222

Note: Observation is a profile. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

	Saved
BA in CS=1	0.008***
	(0.000)
Current Employer Ranked Top 10 for Tech	0.049***
	(0.002)
SR Map/Reduce	0.053***
	(0.006)
JavaScript - SR, No V	0.005***
	(0.001)
JavaScript - No SR, V-Low	0.049***
	(0.001)
JavaScript - SR, V-Low	0.115***
	(0.003)
JavaScript - No SR, V-High	0.147***
	(0.005)
JavaScript - SR, V-High	0.244***
	(0.005)
Controls	Yes
Dep. Mean	0.015
R^2	0.222

Note: Observation is a profile. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Testing if Self-Reporting Matters

- Could be recruiters care about self-reported skills.
- Could also be caused by two other factors:
  1. The choice to self-report is correlated in a complex way with other profile attributes.
  2. Recruiters actually care about a complex interaction of candidate attributes.
- I'm going to show that the two factors above are not the reason why self-reporting predicts more recruiter attention.

Checking If Recruiters Care about Self-Reported Languages

- Take the profiles with JavaScript verified.
- Find groups of profiles with a large number of the exact same attributes, but some self-report JavaScript and some do not.
- Do the self-reporters get saved more?
- Normally matching on high dimensions is hard, but lots of data and limited number of salient profile attributes.

Match Groups

Match all the profiles based on the following:

- "Verified" level in JavaScript
- Bachelors degree year, major, school.
- Masters degree, Ph.D.
- Currently in a coding job, previous employers.
- Platform computed "Expertise" and "Market Value" scores
- Self-reported agile methods, Git/SVN, machine learning, REST, databases.
- Numbers of "years of equivalent experience" in the following programming languages as computed by the platform and displayed on the profile JavaScript, Java, Python, and C#,
- Way that languages Java, Python, C# are reported
- Gender

Predicting Recruiter Interest with Self-Reports

$$saved_i = \alpha_g + \beta sr_i + \epsilon_i$$

- where profile i is a member of profile group g, a set of profiles that share the exact same displayed values for many salient features.
- sr_i is an indicator for if candidate i self-reported JavaScript.
- Estimate separately for profiles "Verified - Low" and "Verified - High" in JavaScript.

Estimates – Self-Reporting JavaScript

	Saved	
	Verified - Low	Verified - High
SR JavaScript	0.063***	0.086*
	(0.010)	(0.047)
Group Fixed Effect	Yes	Yes
N	12,336	1,169
N Groups	1,624	591
Dep. Mean	0.08	0.34
R^2	0.007	0.008

Note: Observation is a profile where JavaScript is "Verified." Within-R^2 shown. Standard errors clustered a profile group level in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

- "Verified Low" self-reporters are 30.7% more likely to be saved.
- "Verified High" self-reporters are 16.7% more likely to be saved.

Why might recruiters focus on self-reported skills?

- Recruiters are incentivized by the "recruiting funnel":
  - The fraction of the candidates that they contact who apply for open positions.
  - The fraction of those applicants who score well on technical interviews.
  - The fraction of those extended job offers who accept and join their firm.
- Self-reporting conveys interest.
  - Tambe, Ye, and Cappelli (2017) workers pay to work with the tech they want to work with.

Table of Contents

Question #2: Are there gender differences in the propensity to self-report programming languages?

- Find candidates who have previous experience in a programming language, and check if they self-reported the language.
- Dataset is 86,365 candidate-language pairs derived from 53,713 distinct candidates.
  - Only candidate-language pairs where candidate uploaded at least one line of open source code.
  - Candidate must have at least one self-reported skill.
  - 2.2% of the candidates from "Profiles" dataset, but 37.6% of the candidates saved on the platform.
  - Only 7% female.
  - 51.9% of men and 46.1% of women have job titles associated with software engineering.

“OS Contributors” Dataset

- For each candidate–language observation, I see:
  - Lines of code uploaded to open source over lifetime.
  - StackOverflow questions and answers.
  - Number of days with uploads.
  - Number of “watchers” of open source projects.
- Mean number of lines of open source code in JavaScript is 48,341.68.
- Mean number of “watchers” of JavaScript code: 20.76.

Gender Differences in Self-Reporting of JavaScript – Code

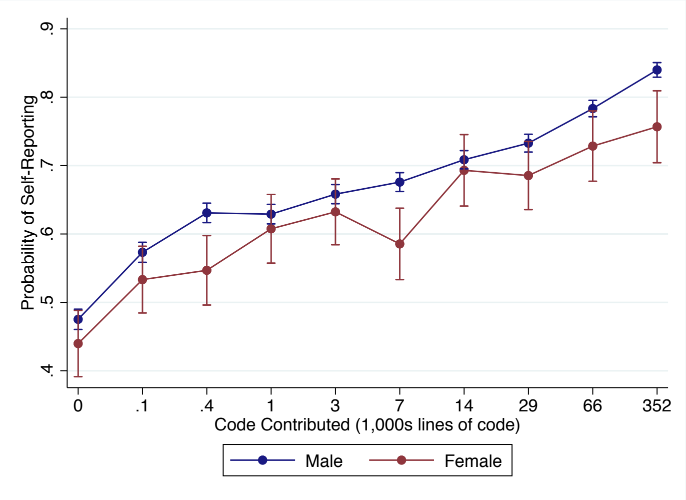

Note: An observation is a profile that is "Verified" in JavaScript, has at least one line of code uploaded to open source, and at least one self-reported skill. Profiles are put into deciles by the lines of code uploaded. The average lines of code per decile is shown on the horizontal axis. 95% confidence intervals around the mean

Gender Differences in Self-Reporting of JavaScript – Watchers

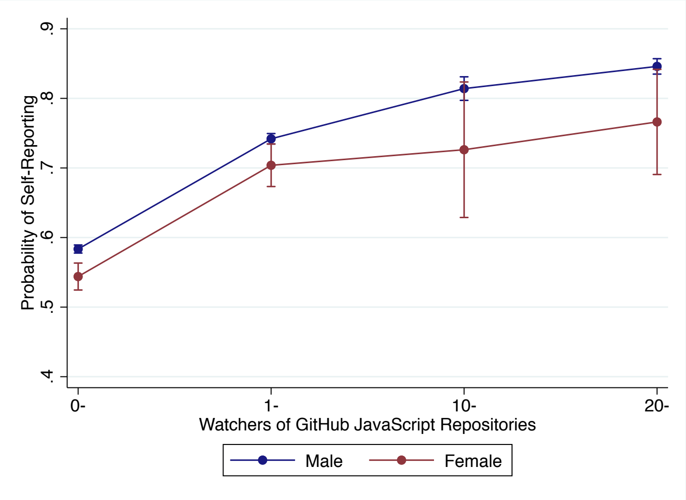

Note: An observation is a profile that is “Verified” in JavaScript, has at least one line of code uploaded to open source, and at least one self-reported skill. A “watcher” is a subscriber to updates about the code.

Estimating Gender Differences in Self-Reporting

$$
\begin{aligned}
sr_{i,l} =& \alpha \\
&+ \beta female_i + \gamma_1 experience_{i,l} + \gamma_2 (female_i \times experience_{i,l}) \\
&+ \gamma_3 X_i + \epsilon_{i,l}
\end{aligned}
$$

- where $sr_{i,l}$ is an indicator for if candidate i lists programming language l in the self-reported skills section.
- Run this regression on candidate-language pairs where the candidate has uploaded at least one line of open source code and has at least one self-reported skill on their resume.

Estimated Gender Difference in Self-Reporting

	Self-Reported		
	Code (10k)	Year	SO Answers
Female	-0.060***	-0.063***	-0.061***
	(0.008)	(0.008)	(0.007)
Experience	0.003***	0.010***	0.001***
	(0.000)	(0.002)	(0.000)
Female x Experience	0.001	0.030***	0.004***
	(0.001)	(0.007)	(0.001)
Edu. Controls	Yes	Yes	Yes
N	82,779	82,779	82,779
N Programmers	52,811	52,811	52,811
Dependent Mean	0.61	0.61	0.61
R^2	0.011	0.009	0.004

Note: An observation is a profile-language pair in which the candidate uploaded at least one line of code. Languages are either JavaScript, Java, C#, or Python. Standard errors clustered at the candidate level in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Possible Explanations

1. Preferences: male and female coders have on average different preferences over occupations.

2. Confidence gap: male and female coders believe the experience cutoff for self-reporting is different.

Difference in preferences over occupations?

- Examine two samples with similar average preferences:
  1. Among those with at least one self-reported skill associated with programming jobs, 53.9% of men and 52.5% of women are currently employed as software engineers.
  2. Among recent bachelors degrees in Computer Science, NSF SESTAT survey says equally likely to do programming.

Prediction:

- If the self-reporting gap is entirely because of preferences, men and women in these two subsamples will self-report at the same rate.

Estimated for Candidates with Programming Skills

	Self-Reported		
	Code (10k)	Year	SO Answers
Female	-0.031***	-0.033***	-0.031***
	(0.008)	(0.008)	(0.008)
Experience	0.003***	0.009***	0.001***
	(0.000)	(0.002)	(0.000)
Female x Experience	0.001	0.023***	0.003***
	(0.001)	(0.006)	(0.001)
Edu. Controls	Yes	Yes	Yes
N	72,287	72,287	72,287
N Programmers	45,611	45,611	45,611
Dependent Mean	0.66	0.66	0.66
R^2	0.012	0.011	0.006

Note: An observation is a profile-language pair in which the candidate uploaded at least one line of code. Only candidates with programming skills self-reported. Languages are either JavaScript, Java, C#, or Python. Standard errors clustered at the candidate level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Estimated for Recent BA in CS Grads

	Self-Reported		
	Code (10k)	Year	SO Answers
Female	-0.024	-0.025	-0.031
	(0.020)	(0.019)	(0.019)
Experience	0.004***	0.028***	0.003**
	(0.001)	(0.004)	(0.001)
Female x Experience	-0.003	0.003	0.001
	(0.003)	(0.016)	(0.002)
Edu. Controls	Yes	Yes	Yes
N	11,273	11,273	11,273
N Programmers	7,184	7,184	7,184
Dependent Mean	0.65	0.65	0.65
R^2	0.011	0.014	0.004

Note: An observation is a profile-language pair in which the candidate uploaded at least one line of code. Only candidates with Computer Science bachelors degrees who graduated between 2010 and 2015 are included. Languages are either JavaScript, Java, C#, or Python. Standard errors clustered at the candidate level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Difference in Preferences

- A small gap in self-reporting remains even after using subgroups with very similar occupational trajectories.
- Differences in preferences likely account for a large portion of the self-reporting gap, but cannot account for the entire gap.

Confidence gap?

Say different experience cutoffs before self-report:

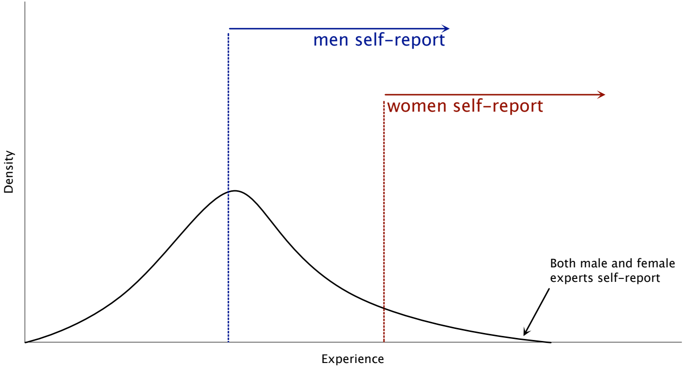

Predictions:

1. $E[\text{experience}|SR, Male] < E[\text{experience}|SR, Female]$

2. Men and women with the highest experience should self-report equally.

Gender Differences in Self-Reporting of JavaScript – Code

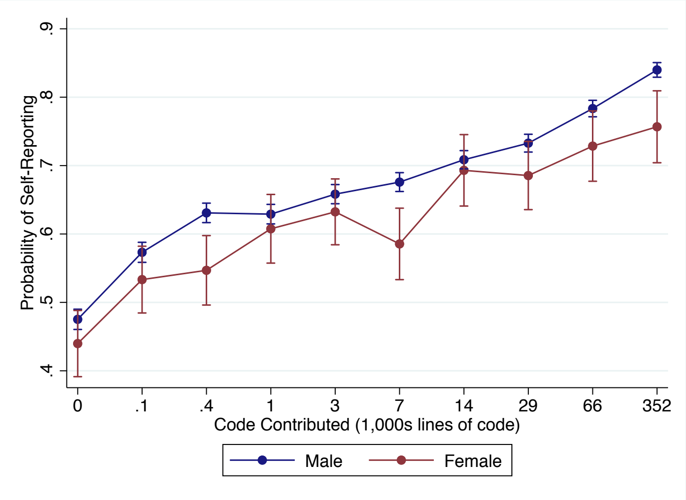

Note: An observation is a profile that is "Verified" in JavaScript, has at least one line of code uploaded to open source, and at least one self-reported skill. Profiles are put into deciles by the lines of code uploaded. The average lines of code per decile is shown on the horizontal axis. 95% confidence intervals around the mean

Gender Differences in Self-Reporting of Ruby – Code

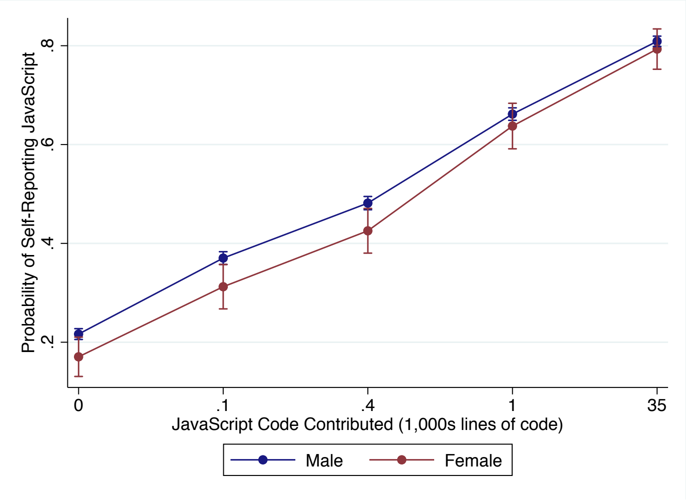

Note: An observation is a profile that is "Verified" in Ruby, has at least one line of code uploaded to open source, and at least one self-reported skill. Profiles are put into quintile by the lines of code uploaded. The average lines of code per decile is shown on the horizontal axis. 95% confidence intervals around the mean are also shown.

Takeaway

- Differences in preferences can account for a large amount of the gender gap in self-reporting, but not all.
- A confidence gap contributes to the gender differences in self-reporting.
- Therefore, among the open source coders who go into coding jobs at similar rates, self-reporters are on average more experienced.

Table of Contents

Question #3: Do profile attributes predict different recruiter responses by gender?

- Given similar profile attributes, do female candidates have a higher probability of being saved?
- If a recruiter is looking for someone with a programming language skill, the women who self-report knowing that language have on average higher previous experience.
- Do recruiters adjust based on this gender difference in self-reporting?

Difference in Predict Probability

- Predict which profiles get saved using all attributes and interactions with all attributes.
- At the mean value of all profile attributes, the predicted probability of being saved:
  - 0.0157 for men and 0.0137 for women = 12.74% lower
- On average, not advantageous to be a female candidate.
- But it might still be that for the skilled candidates, those who know a programming language, recruiters are leaning toward female candidates.

Coefficients on JavaScript

- Plot the coefficients from the regression for how JavaScript appeared on the profile.

- See if female candidates who know JavaScript are more likely to be saved.

Average Predicted Probability Saved by Gender and JavaScript

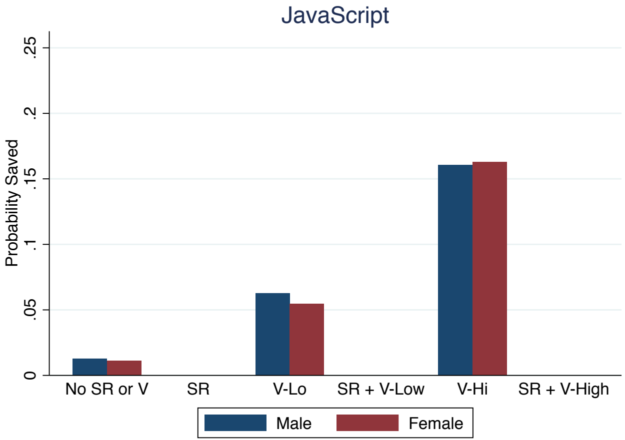

Note: The average predicted probability is computed by holding fixed the way that JavaScript appeared on the profile as well as the gender of the candidate.

Average Predicted Probability Saved by Gender and JavaScript

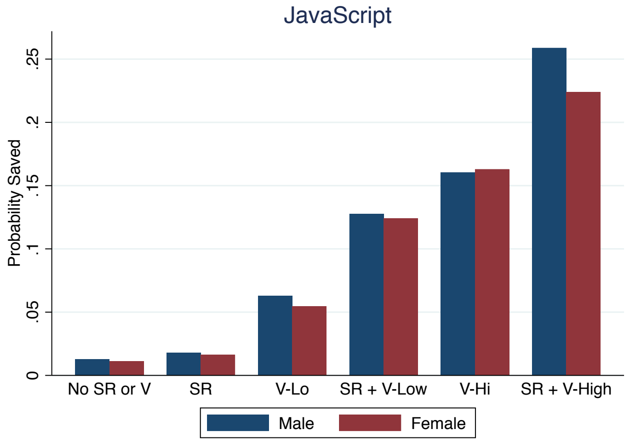

Note: The average predicted probability is computed by holding fixed the way that JavaScript appeared on the profile as well as the gender of the candidate.

Only Interact Self-Reported Skills and Languages

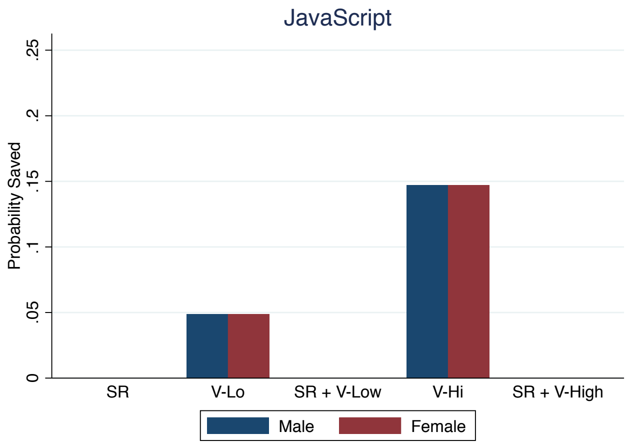

Note: The average predicted probability is computed by holding fixed the way that JavaScript appeared on the profile as well as the gender of the candidate.

Only Interact Self-Reported Skills and Languages

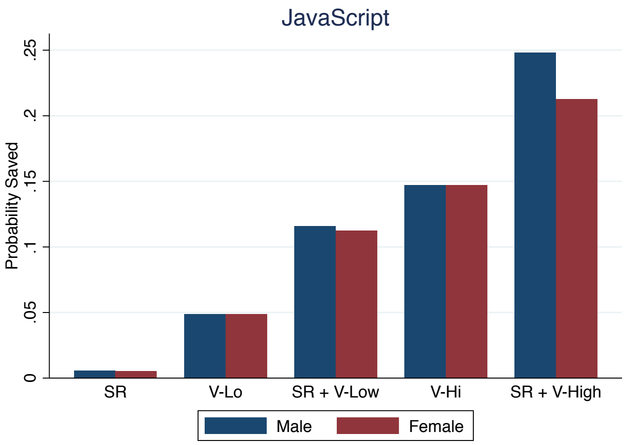

Note: The average predicted probability is computed by holding fixed the way that JavaScript appeared on the profile as well as the gender of the candidate.

Findings

- Overall, recruiters don't appear to be leaning towards female candidates.
- If recruiters are trying to find the most experienced programmers, they should be more inclined towards female candidates who self-report.
  - They are not doing so significantly.
- We do see significant differences in the response to non-technical skills.
  - Recruiters can adjust.

Why might recruiters not use gender and adjust?

- If recruiters are not biased...
  - They might know be aware of gender differences or how to use this information (none that I have talked to knew).
  - They don't want to appear to be using gender in their decisions.
  - Might have a different objective (attrition/retention?).
  - Might have multiple channels.
- If some recruiters are biased...
  - Gender blind recruiting could be beneficial.
  - Algorithmic recruiting might want to incorporate gender information.
  - Change the questions candidates are asked.

Recruiters vs the Crowd

- On GitHub, two measures of reputation:
  - Users can "follow" other users (subscribe to updates about code they contribute to any open source project).
  - Users can "watch" for changes in a particular open source project (subscribe to bug fixes and new features in a particular project).
- Do recruiters save the male and female candidates with similar numbers of "followers" at the same rate?
- Do recruiters save the male and female candidates with similar numbers of "watchers" to their JavaScript projects at the same rate?

Probability Saved by GitHub Followers

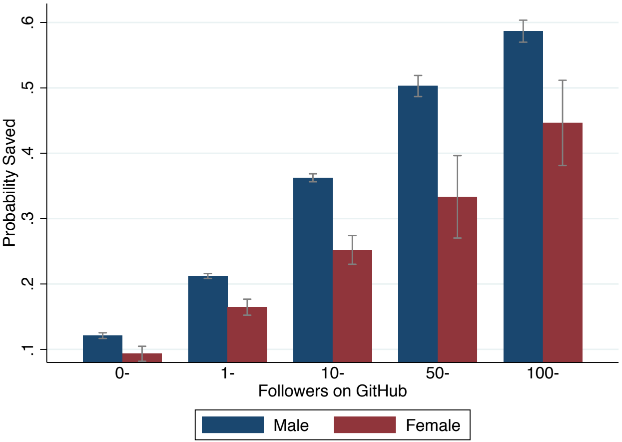

Note: The probability that candidates are saved is computed for each group based on the number of GitHub followers they have.

Probability Saved by “Watchers” of JavaScript Projects

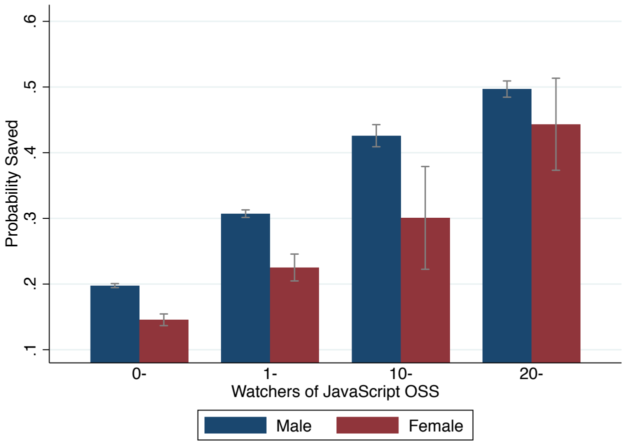

Note: The probability that candidates are saved is computed for each group based on the number of GitHub watchers of JavaScript repositories they have.

Findings

- Conditional on the number of "followers" or "watchers" to open source work, there is a large gender gap in recruiter interest for male and female candidates.

Table of Contents

Conclusion

- Recruiters pay attention to self-reported skills on profiles.
- Female coders are less likely than male coders to list programming languages they write open source code in.
- Recruiters do not adjust for gender differences in self-reporting when deciding which candidates to contact.
- The market does not provide corrective pressure on either the labor demand or the labor supply side that would increase the representation of women among software engineers.

Table of Contents

Appendix Slides

Distribution of Lines of Code Contributed to Open Source in JavaScript

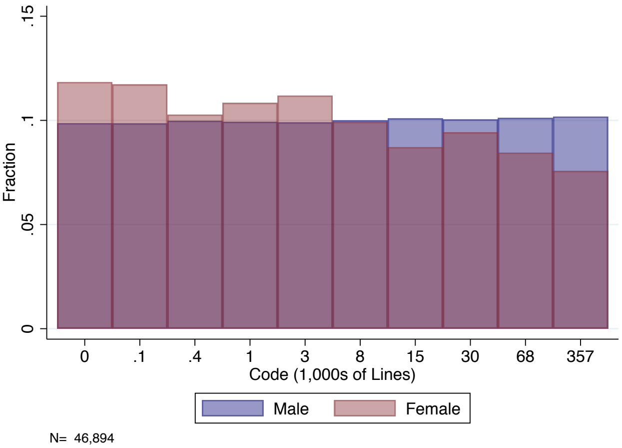

Note: The above figure shows distribution of the number of lines of JavaScript code uploaded to open source amongst candidates with at least one line of code and at least one self-reported skill on their resume. The blue distribution represents the distribution of male candidates, while the red distribution is that of the female candidates. The number of lines of code is adjusted from the raw numbers in order to exclude copied or "forked" code from other open source contributors.