

THE ROLE OF THE ASK GAP IN GENDER PAY INEQUALITY^{*†}

Nina Roussille

[Click here for the most recent version](#)

August, 2023

Abstract

The gender *ask gap* measures the extent to which women ask for lower salaries than comparable men. This paper studies its role in generating wage inequality, using novel data from an online recruitment platform for full-time engineering jobs: Hired.com. To use the platform, job candidates must post an *ask salary*, stating how much they want to make in their next job. Firms then apply to candidates by offering them a *bid salary*, solely based on the candidate's resume and ask salary. If the candidate is hired, a *final salary* is recorded. After adjusting for resume characteristics, the ask gap is 2.9%, the bid gap is 2.2%, and the final offer gap is 1.4%. Further controlling for the ask salary explains the entirety of the residual gender gaps in bid and final salaries. To further provide evidence of the causal effect of the ask salary on the bid salary, I exploit an unanticipated change in how candidates were prompted to provide their ask. For some candidates in mid-2018, the answer box used to solicit the ask salary was changed from an empty field to an entry pre-filled with the median bid salary for similar candidates. I find that this change drove the ask, bid, and final offer gaps to zero. In addition, women did not receive fewer bids or final offers than men did due to the change, suggesting they faced little penalty for demanding comparable wages.

JEL codes: J31; J16; J49

Keywords: Gender wage gap, gender ask gap, job search, online recruitment

^{*}Email: nroussil@mit.edu. I am indebted to Patrick Kline, Hilary Hoynes and Gabriel Zucman for their guidance and encouragement on this project. For very helpful comments, I would like to thank David Autor, Anne Boring, David Card, Stefano DellaVigna, Larry Katz, Supreet Kaur, Juliana Londoño-Vélez, Claire Montialoux, Jesse Rothstein, Ben Scuderi and Chris Walters as well as numerous participants at the Summer Institute Labor Studies session and several other seminars. I would also like to thank Hired.com for letting me access their data, and in particular I am grateful to Nate West, Andrew Collins, Michael Mitchell, John Sapienza, Saba Sedighi and Erica Yamamoto for helping me navigate the platform and for providing me with insightful comments. I would finally like to thank Christian Höhne and Emiliano Sandri for excellent research assistance. I gratefully acknowledge financial support from the UC Berkeley Economics department, the Burch Center, the UC Berkeley Opportunity Lab, the Strandberg Family Graduate Fellowship Fund and the Hub for Equal Representation.

[†]The views expressed here are my own and do not necessarily reflect those of Hired.com. This manuscript was not subject to prior review by any party, except to verify that it does not contain any confidential data or information, as per the research contract signed at the outset of this project.

“We cannot change what we are not aware of, and once we are aware, we cannot help but change.”

— Sheryl Sandberg *Lean In: Women, Work, and the Will to Lead*

1 Introduction

Over the past several decades, the raw gender pay gap in the U.S. has declined significantly, falling from about 40% in the 1960s to 20% today. While the raw gap has narrowed, the residual pay gap - the portion of the pay gap that cannot be accounted for by gender differences in measured qualifications - has stagnated at around 10% for the past 30 years (Blau and Kahn (2017)). In parallel, there is mounting evidence that women still have lower salary expectations than comparable men, especially at the top of the income distribution (Reuben, Wiswall, and Zafar (2017), Bergerhoff et al. (2021)). Taken together, these facts raise concerns that women’s lower salary expectations contribute to the persistence of the residual pay gap (Babcock et al. (2003), Leibbrandt and List (2015), and Biasi and Sarsons (2022)).

This paper investigates how gender differences in salary demands influence the wage gap in a high-skilled online labor market. Recent survey evidence indicates that the majority of high-wage workers in the U.S. are asked to state their desired salary during the recruitment process (Agan, Cowgill, and Gee (2020)). Yet, quantifying the role of the candidates’ desired salary in the determination of salary offers in traditional labor markets has proven challenging. Data on workers’ salary demands is typically collected via surveys or laboratory experiments that may not capture the salary negotiations that actually arise in high-stakes recruitments. In addition, available wage data usually provides information on only one side of the market: either the candidate’s side (e.g., survey evidence on salary expectations) or the firm’s side (e.g., administrative data on firm salary offers). No dataset simultaneously combines information on candidate salary demands and on how these demands influence the salary offers they receive from firms.

To fill this gap, I analyze data from Hired.com, a leading online recruitment platform for full-time, high-wage engineering jobs. The key novelty of this platform is that it records previously unexplored components of the salary negotiation process. First, every candidate has to provide the salary they are looking for in their next job. This *ask salary* is visible to firms recruiting on the

platform, along with the candidate’s resume information. Second, companies signal their interest to candidates with a *bid salary*, indicating how much they are willing to pay the candidate before interviewing them. Last, the platform records a *final salary* if the candidate is hired. Given that the average annual salary on the platform is \$120,000, the candidates on Hired.com are a highly relevant population for studying high-stakes wage bargaining.

Using data on more than 110,000 candidates over several years, I first document a 6.6% raw ask gap on the platform. After controlling for all the candidates’ resume characteristics, the ask gap is still 2.9%. In other words, women ask for 2.9% less than men with comparable resumes. This gap is both statistically significant and economically meaningful: it represents \$3,830 every year, on average. I also find significant heterogeneity in the ask gap. Using the Sorted Partial Effects method of Chernozhukov, Fernández, and Luo (2018), I find ask gaps ranging from 8.5% to -2.1%, with the largest gap arising among candidates who are not currently employed, have more experience, and fewer credentials.

Second, I document the relationship between the ask salary and firms’ bid and final offer gaps. Using data on more than 460,000 bids, I find a raw bid gap on the platform of 3.3%. Adjusting for candidates’ resume characteristics but excluding their ask salary leaves a 2.2% residual bid gap. When candidates’ ask salaries are included as a control, and even when candidates’ resume characteristics are not, this residual bid gap disappears. In other words, while accounting for resume characteristics can only reduce the raw bid gap by 33%, gender differences in ask salaries can explain 100% of it. Similarly, for a given job, resume characteristics account for 3 ppts of the 4.8% unadjusted bid gap, while further controlling for the ask salary brings the bid gap to zero, indicating that the bid gap doesn’t arise from the composition of jobs for which women interview. These results are qualitatively the same when restricting the sample to firms that make a final offer or when adding firm fixed effects. A linear model conditioning solely on candidates’ resume characteristics explains 82% of the variation in bid salaries, while adding the ask salary to the controls raises the R^2 to 0.95, leaving little room for omitted variable bias. For the sub-sample of 7,582 hired candidates, gender differences in ask salaries explain nearly all of the gap in final offers. In particular, while conditioning on resume characteristics only narrows the final offer gap to 1.4%, adding the ask salary to the controls reduces the final offer gap to -0.9% and further controlling for firm fixed effects brings it to zero.

To further provide evidence of the causal effect of ask salaries on bid salaries, and thus final offers, I take advantage of an unanticipated feature change that affected a subset of candidates on the platform and induced women to ask for more. In mid-2018, Hired.com unexpectedly changed the way that some candidates were prompted to provide their ask salary. Until mid-2018, candidates stated their ask salary by filling out an empty text box. Starting in mid-2018, the answer box for San Francisco software engineers was pre-filled with the median bid salary over the past 12 months for the candidate’s combination of desired location, job title, and experience. This change gave candidates information on the typical offers received by similar candidates on the platform and provided them with an anchor to benchmark their own ask salary. Using an interrupted time series design, I show that the new framing of the ask salary elicitation eliminated both the ask and bid gap. These results are driven by women asking for higher salaries after the reform. Further, I find no discernible impact on the number of bids that women received or their likelihood of receiving a final offer, suggesting that there was no downside, for women, to asking for more. Finally, I leverage the reform effects to discuss plausible mechanisms behind women’s initial lower ask. The evidence I gather is most consistent with an information channel: women had downward beliefs about the market wage for their resumes and the reform corrected them.

This paper contributes to several lines of research. First, it integrates the ask gap into the prominent literature on gender wage gaps. The most common concept measured in this literature is the gender gap in realized wages (Blau and Kahn (2017), Olivetti and Petrongolo (2016)), but a more recent strand of the literature has turned to investigate gender gaps in salary expectations (Reuben, Wiswall, and Zafar (2017), Bergerhoff et al. (2021)). Unlike traditional expectation measures, the ask salary plays a direct role in the salary negotiation, as it is one of the few signals voluntarily transmitted by the candidates to potential employers. Relative to survey measures, Hired data have several strengths: a large sample size, no missing values due to non-response, and real labor market relevance. Finally, the recruitment process on the platform allows for the direct measurement of the impact of candidates’ ask gap on the firms’ offer gap, while most studies only observe either the candidate or the firm side of the market. Some exceptions can be found in the literature on reservation wages (e.g. Le Barbanchon, Rathelot, and Roulet (2021)), but in contrast with the ask salary, reservation wages are not observable by firms.

Second, my research relates to the literature on gender differences in negotiation, especially at

the top of the income distribution (Bertrand (2017), Goldin (2014), and Garbinti, Goupille-Lebret, and Piketty (2018)). Most of the evidence in this literature comes from laboratory experiments (Babcock et al. (2003), Bowles, Babcock, and McGinn (2005), Small et al. (2007), Exley and Kessler (2022)) or surveys (Babcock and Laschever (2006)). These papers find that, in the lab or in self-reported survey data, women have lower salary expectations, negotiate less and receive lower salary offers. I contribute to this literature first by showing that women indeed ask for significantly less in high-stakes environments and second by providing direct evidence that this gap is consequential for resulting salary offers.

Finally, my research contributes to a strand of literature in behavioral labor economics that examines the role of information in the job search process and salary decisions. Some recent papers (Bennedsen et al. (2022), Baker et al. (2023), Cullen and Pakzad-Hurson (2023), Cortés et al. (forthcoming) and Jäger et al. (forthcoming)) illustrate, in the field, how accurate information and pay transparency can correct workers’ misperceptions about wages and reduce the gender wage gap. In the lab, Rigdon (2012) shows that, in a “Demand-Ultimatum” game where participants have to share \$20, women initially request less than men but after they are informed about the amounts demanded by other participants, they start requesting the same as men. In contrast, recent lab-based evidence finds that nudging women to “lean in” can result in worse outcomes for them. For instance, Exley, Niederle, and Vesterlund (2020) show that, when workers and firms have to ex-post split the sum of their respective contributions in a series of (modified) ultimatum games, negotiations are not helpful and may actually harm women. I see my paper as complementary to these lab experiments and argue that better understanding the contexts and conditions under which asking for higher pay benefits, rather than harms, women is an important avenue for research.

Section 2 provides details on the empirical setting. Section 3 presents a detailed description of the data. Section 4 describes the empirical strategy to estimate the ask gap and documents its existence and magnitude. Section 5 provides evidence of the impact of the ask gap on the bid gap and final salary gap. Section 6 details the reform on elicitation of candidates’ ask salaries and reports estimates of the effects of the reform and Section 7 provides a framework to interpret the results of the reform. Section 8 concludes.

2 Institutional setting

2.1 Market description

Several previous papers have studied online labor markets, such as Amazon MTurk, to explore the causes of the gender pay gap (Litman et al. (2020), Gomez-Herrera and Mueller-Langer (2019)). These markets allow researchers to run experiments and to precisely record the impacts of experimentally assigned treatments on labor market outcomes. However, most of these markets offer task-based, remote, and low-wage jobs. Hence, even experimental evidence on bargaining on those platforms may not reflect behaviors in more traditional labor markets. In contrast, Hired.com mostly features full-time, onsite, high-wage engineering jobs based in the U.S.: 96.9% of the candidates on the platform state that they are looking for a full-time job and the average salary offered by firms on the platform is high (\$119,548). In short, Hired.com should be thought of as a job board for highly-educated candidates, with a focus on the tech industry. The candidates and jobs on Hired.com are comparable to those listed on other recruitment platforms for similar careers. For instance, the most common profile on Hired.com is a software engineer in San Francisco. As of April 2020, Glassdoor’s average salary for this profile was \$119,488 and Paysa’s was \$132,000.¹ Hired’s salary for such profiles is \$130,349, which is in the bracket between Glassdoor’s (lower bound) and Paysa’s (upper bound) salaries. The Hired.com sample also features profiles with different levels of seniority; for instance, the years of experience of San Francisco software engineers² are distributed similarly to their equivalent found on Payscale. Additionally, the 6,532 firms in the Hired sample are also representative of the digital economy ecosystem: they are a mix of early-stage firms, more mature start-ups (e.g. Front, Agolia), and larger, more established firms (e.g. Zillow, Toyota). Finally, the gender ratio on Hired.com (20.8% female) is similar to the general population of computer science and engineering graduates.³

¹ Paysa is a personalized career service offering salary compensation and job matching for corporate employees. It is a useful reference for comparing employee salaries in the tech industry.

² Among San Francisco software engineers, 6% have 0-2 years of experience in software engineering, 21% have 2-4 years of experience, 23% have 4-6 years of experience, 35% have 6-10 years of experience, 9% have 10-15 years of experience, and 6% have more than 15 years of experience.

³ Chamberlain and Jayaraman (2017) showed that among science and engineering graduates, only 26% are female, and a disproportionate number of these female graduates end up working in fields other than computer science. This gender imbalance in a high-wage sector makes the tech industry a particularly interesting case study of the gender pay gap among top earners.

2.2 Recruitment process

The hiring process on Hired.com differs from a traditional job board in two main ways. First, on a traditional job board, firms post a job description (that may contain a posted wage), and then candidates apply to each posted job separately. Afterwards, the company interviews a selection of applicants and decides whether and who to hire. In contrast, on Hired.com, companies apply to candidates based on their profiles, then candidates decide whether or not to interview with the company based on the job description and bid salary they receive. Second, in a wage posting context, candidates' demands do not directly influence firms' posted wages. In contrast, on Hired.com firms make salary offers only after observing the candidates' resumes and asks. Formally, the recruitment process can be divided into the following three sequential steps, also described in Figure I:

Supply side: Candidates create a profile that contains standardized resume entries and the salary that the candidate wants to make in their next job: their *ask salary*.⁴ Figure A.1 is a screenshot of a typical candidate's profile, and Table B.1 further provides the listing of all fields on a profile. In short, a profile includes the current and desired location of the candidate, their job title (e.g. software engineering or web design), their experience in this position, their top skills (e.g. coding languages such as R or Python), their education (degree and institution), the firms they worked at, their contract preferences (remote or on-site, contract work or full-time), as well as their search status, which describes whether the candidate is actively searching or simply exploring new opportunities. Importantly, the ask salary is a required field prominently featured on all profiles.

Demand side: Firms get access to candidate profiles that match standard requirements for the job they want to fill (job title, experience, and location). To apply for an interview with a candidate, the company sends them a message - the *interview request* - that contains a basic description of the job as well as the salary at which they would be willing to hire the candidate: their *bid salary*. Figure A.2 is a screenshot of a typical message sent to a candidate by a company. The bid salary is prominently featured in the subject line of the message and is required to be able to send the message. The equity field also exists but is optional.

Demand meets supply: Hired.com records whether the candidate accepts or rejects the interview

⁴ Specifically, the ask salary is the answer that candidates give to the question: "What base salary are you looking for in your next role?". It then appears on a candidate's profile (see Figure A.1) as a bullet point saying: "Prefers base salary of X per year." (where X is the answer of the candidate to the ask salary question.)

request. While interviews are conducted outside of the platform, Hired.com gathers information on whether the company makes a job offer to the candidate and at what *final salary*. It is important to note that the bid salary is non-binding, so the final salary can differ from it. Finally, we observe whether the candidate accepts the final salary offer, in which case the candidate is hired.⁵

2.3 Relevance of the recruitment process to other wage bargaining settings

While the ability to record granular steps of the negotiation is unique, some of these steps are similar in the broader labor market, especially for high-wage candidates. For instance, using a 2019 survey of 504 Americans in the labor force, Agan, Cowgill, and Gee (2020) found that 55% of workers making above \$68,000 a year were asked for their desired salary during the recruitment process (compared to 42% of the full sample). Therefore, Hired.com makes explicit what effectively occurs during the majority of high-wage interviews: candidates are asked to disclose their desired salary. There is also evidence that, in a non-trivial share of wage negotiations, candidates are asked for their desired salary before the company makes them an offer. For instance, in a Google survey of approximately 400 subjects, Barach and Horton (2021) found that, among candidates who negotiated their wages, 39.2% proposed a wage before the firm did. It is therefore not uncommon for the candidate to state their ask first, although, in more traditional settings it might occur later in the recruitment process (e.g. after, rather than before, the interview).

3 Data

3.1 Sample size

Table I reports the sample sizes for the main units of observation on the candidate side (first row of Panel (a)) and company side (first row of Panel (b)). The final dataset has 113,777 candidates, 39,839 jobs, and 6,532 firms located in 20 different cities. Each job is sent out on average to 11.6 candidates so there are a total of 463,860 interview requests ($\approx 39,839 \times 11.6$) sent out by firms,

⁵ While I can't ensure that all final offers are recorded correctly, there are a number of features that guarantee high-quality data all the way to the final offer. First, in the time period of this study, Hired.com was paid by most firms only if the firm made a final hire. Therefore, the platform had strong incentives to ensure that firms report these final hires. Second, it is quite easy for Hired to detect fraud (i.e. a match made on the platform that results in a hire outside of it). Indeed, Hired records all the profiles interviewed by the firm, and most firms have a career page with their current employees. Therefore, checking interview records against hires is quite straightforward. Finally, a one-time fraud could result in the high cost of being kicked out indefinitely from Hired.

resulting in 7,582 final offers. The data spans several recent years but, per the research contract signed with the company, the exact start and end dates of the period cannot be disclosed.

3.2 Gender

Gender is an optional field on the profile and only 50% of the candidates self-declared their gender. In order to obtain gender data for the other 50%, I use a standard prediction algorithm based on first names.⁶ Reassuringly, for the sub-sample that self-declared their gender (i.e. 50% of the full sample), I verified that the algorithm guessed incorrectly only 0.6% of the time. Firms are informed of the gender of candidates since most profiles contain pictures and first names. Combining explicit declarations and imputation, I can classify 84.6% of the profiles. Women represent 20.8% of the classified sample, while men represent the remaining 80.2%.

3.3 Candidate summary statistics

Table I Panel (a) provides information on the resume characteristics of the candidates. They have, on average, 11.3 years of experience, which corresponds to the industry average in this sector (Visier and Insights (2017)). They are highly educated: 97.6% of the candidates have at least a bachelor’s degree and 41.4% have at least a master’s degree. Given that the platform targets engineers, it is not surprising that 55.2% of the candidates have a degree in Computer science and that 61.7% of them are looking for software engineering positions. The platform’s focus on the tech industry is also reflected in the location of its candidates: 31.6% of them are looking for a job in San Francisco. About 3 out of 4 candidates are looking for job-to-job transitions.

Men and women differ in experience, occupation, and location. On average, women have 1.6 fewer years of experience than men. However, mirroring the overall U.S. population, women appear to be more educated (45.2% of them have a master’s vs 40.3% of the men). With respect to occupation, 66.6% of the men are looking for software engineering positions, while only 43.2% of the women are. The other women are mainly looking for either a web design (16.6%) or a product management position (11.4%). Accordingly, the share of men with a computer science (CS) degree is higher (57.2% vs. 47.7%). Finally, women are more likely to be looking for a job in SF (37.5%

⁶ The prediction can take 5 values: “male”, “mostly male”, “ambiguous”, “mostly female”, and “female”. When available, I used the self-declared gender of the candidate; otherwise, I impute gender using the algorithm, assigning a gender only to candidates for whom the algorithm predicted “male” or “female”.

vs 30.0%).

Candidates can also express preferences about the size and industry of their ideal firm, as well as some preferred future job features. Around 75% of the candidates express at least one preference. Table B.2 presents gender differences in these preferences controlling for candidates' resume characteristics. The main takeaway from this table is that, while men and women differ in their preferences in the expected direction (e.g. women are more likely than men to prefer firms that are socially conscious, more likely to seek a mentorship role, and less likely to seek a leadership role), the differences are quite small in magnitude (e.g. 18.9% of men express a preference for leadership, that share is only 0.5ppt lower for women with the same resume characteristics).

3.4 Firm summary statistics

Table I Panel (b) provides information on firm characteristics such as revenue, age, size, or industry. Around a third of companies are early-stage firms that were founded within 5 years of the end of the sample period, half of them report less than 25 million in revenue, and almost half the firms enlist between 1 and 50 employees. Medium-sized companies or matured start-ups with 51 to 500 employees make up around 40% of the sample and the remaining 11% consist of established companies with more than 500 workers. The overall distribution of revenue is strongly right skewed with a median just above 25 million USD, but with almost a quarter of the sample reporting a revenue higher than 500 million USD. Consistent with candidates' current and preferred location, the most common location among firms is San Francisco (40%), followed by New York (24%) and Los Angeles (7%). The three most frequent industries in which companies operate are Enterprise Software (15%), Banking & Finance (10%), as well as Analytics (8%).

3.5 Candidate - Firm interactions

For a given job, firms contact on average 11.6 candidates. Importantly, for the same job, there can be as many bid salaries as there are candidates contacted. In fact, only 2.4% of jobs offer the same bid salary to all candidates. The within-job variation in bids is also quite large: the average standard deviation of bids for a given job is \$16,575. On the candidate side, the average number of interview requests, conditional on receiving at least one, is 4.5, and candidates agree to interview 62% of the time.

Once a candidate profile is reviewed and approved by Hired.com, it becomes visible to firms. The default length of a spell on the platform is two weeks.⁷ On the company side, a separate identifier is created for each job that the company wants to fill. The company may be looking to hire several candidates for the same job. If we restrict the sample to jobs that make hires, 77.3% of them hire a single person and 14.3% hire two, the remaining 8.4% hire three or more. Only a subset of jobs find a suitable candidate on the platform, and similarly, only some of the candidates are hired. Firms that hire a candidate for the job exert additional search efforts on the platform: on average, they send almost three times as many interview requests to candidates than the average (30.2 vs 11.6). Similarly, candidates who get hired receive about 1.5 times as many interview requests as the average candidate (6.6 vs 4.5) and they are more likely to accept an interview request.

3.6 How do the ask and bid salaries relate to more traditional salary measures?

This paper measures two previously unobserved components of salary negotiation: the ask and bid salaries. Therefore, it is important to understand how these relate to more traditional measures. For instance, how does the ask salary compare to salary expectations or the reservation wage? Further, given that the bid is non-binding, how does it relate to final offers?

The ask salary is defined as the answer that candidates give to the question: “What base salary are you looking for in your next role?”. Candidates record this ask knowing that it will be visible to firms hiring on the platform. The closest concept previously measured in workers’ and job seekers’ survey data is salary expectations, i.e., how much people expect to make in their next job (e.g., Reuben, Wiswall, and Zafar (2017)). The key conceptual difference with the ask is that salary expectations are not observable by firms. This difference has important implications: the ask is disclosed in the salary negotiation while salary expectations can be measured outside of a recruitment context. Given the strategic game at play in salary negotiations, candidates may reveal an ask that is different from their “true” salary expectations in order to maximize their final offer.

Candidates can adopt different strategies for the choice of the ask salary. Some candidates may choose to record their reservation wage, i.e. the lowest wage at which they would accept a job. Others may provide an estimation of their market value while some may put the highest salary at

⁷ Candidates can request to remain visible for two to four additional weeks. 55% of the candidates are live for two weeks, 22% remain visible for four, and the remaining 23% for six.

which they think they can be hired. These possible interpretations are, to some extent, testable since they give rise to different responses to the bids received. For instance, if the ask is interpreted as a reservation wage, then we should observe that very few candidates accept interviews with firms that make bids below their ask. We test this prediction in Figure II Panel (a), plotting the probability of acceptance of an interview request against the ratio of the bid to ask salary. We first observe that, even when a bid is below the ask, candidates still accept the interview request on average 49% of the time. Therefore, the ask salary is not strictly conveying a reservation wage. Second, candidates do react to higher bids: the probability of acceptance is an increasing function of $\frac{bid}{ask}$, especially in the neighborhood of $\frac{bid}{ask} = 1$. There is however no detectable difference between men and women in their acceptance behavior.

When declining an interview request, candidates are given the option to provide a reason for their decision, and 55% of them choose to do so. The candidates can choose from justifications such as “company culture,” “company size,” and “insufficient compensation.” The latter is the justification I label as “bid too low.” Figure A.3 relates the share of candidates listing “bid too low” as the reason for turning down the interview request to $\frac{bid}{ask}$. As expected, candidates are much more likely to list “bid too low” as a reason for their decision when $\frac{bid}{ask} < 1$. In particular, while this reason is virtually never brought up when the ask is equal to or above the bid, it explains more than 31% of the rejections when the bid is less than 0.8 times the ask, and it is still mentioned in 12.5% of cases when the bid is between 0.8 and 1 times the ask.

The bid salary is what firms declare they are willing to pay the candidate solely based on their profile, before any interaction with them. The final salary is offered to a candidate at the hiring stage. Given that companies are by no means contractually bound by their bids, final salaries may differ from bids. Figure II Panels (b) and (c) show that the relationship between the two is linear, except at the very top, and the slope is close to one. Additionally, 36% of all final offers are identical to the bid, and 78% of all final offers are within \$10,000 of the bid.

4 Documenting the gender ask gap

4.1 The gender ask gap: Methodology

Following the literature, we define the raw gender ask gap as the coefficient β_0 in the regression:

$$\text{Log}(\text{Ask}_i) = \alpha + \beta_0 \text{Female}_i + \gamma_t + \epsilon_i \quad (1)$$

where Ask_i is the ask salary of candidate i , Female_i is a dummy equal to one if the candidate is female, γ_t is the Month \times Year fixed effect, and ϵ_i is the error term.⁸

The adjusted gender ask gap is given by the coefficient β_0 in the regression:

$$\text{Log}(\text{Ask}_i) = \alpha + \beta_0 \text{Female}_i + \beta_1 X_i + \gamma_t + \epsilon_i \quad (2)$$

where the controls X_i are the candidates' resume characteristics, as described in detail in Table B.1. These controls include the variables we typically find in the gender pay gap literature using CPS or PSID data (e.g., education level and job title category), as well as more granular resume characteristics capturing, for instance, education quality and work history.

An alternative take on the ask gap is to consider each interview request a candidate receives as a separate observation. Column (7) of Table II therefore implements the following strategy:

$$\text{Log}(\text{Ask}_{ib}) = \alpha + \beta_0 \text{Female}_i + \beta_1 X_{ib} + \gamma_t + \epsilon_{ib} \quad (3)$$

where Ask_{ib} is the ask salary of candidate i when he or she receives her b 'th bid⁹, Female_i is a dummy equal to one if the candidate is female, γ_t is a Month \times Year FE, ϵ_{ib} is an error term, and t is a function of i and b , $t(i, b)$, the time at which candidate i received bid b . The advantage of this specification is that the units of analysis are the same as those in Table III, which investigates the relationship between the ask and the bid gap.

⁸ When collapsing the data to the candidate level, I select the first listed ask of each candidate. The results are qualitatively the same if we opt for the last ask salary (Table B.3).

⁹ Note that candidates may update their ask salary in a given spell and a small share (7.4%) of them do so. Appendix F discusses the behavior of these candidates.

4.2 Results

Graphical evidence Figure A.4 Panel (a) plots kernel density estimates of the distributions of ask salaries, separately by gender. The figure shows that men’s and women’s distributions have a similar shape, except that women’s distributions are comparatively shifted to the left: On average, women ask for \$6,826 less than men (\$115,116 vs \$121,942).¹⁰

Regression results Estimates of β_0 in Equation 1, reported in Table II Column (1), indicate that there is a 6.6% raw ask gap between men and women. Once we have linearly controlled for all the resume characteristics from the candidate’s profile in Column (5), the adjusted ask gap from Equation 2 is 2.9%. This gap is both statistically significant and economically meaningful: it represents \$3,830 in annual salary, on average. Columns (2) to (5) progressively add the resume characteristics detailed in Table B.1. This exercise identifies which resume controls reduce the gender ask gap, from a raw 6.6% to an adjusted 2.9%. Column (6) includes fixed effects on candidates’ most recent company and the adjusted gender ask gap goes to 3.2%. Further, I implement a selection exercise on observed and unobserved variables following Altonji, Elder, and Taber (2005). I obtain $[-0.028; -0.011]$ as a bounding set for β^{11} (see Table B.4). Since 0 does not belong to this set, I can reject the null of a zero gender ask gap. Last, to account for potential complex interaction effects among control variables, I ran a Double-lasso procedure ala Belloni, Chernozhukov, and Hansen (2014), with 2 and 3 way interactions between explanatory variables, which resulted in a 3.1% ask gap. Adding controls for experience, location, and job title first narrows the gap down to 4.3% (Column (2)). This is mostly due to women having on average less experience or opting for lower-paid occupations. Conversely, adding education controls (Column (3)) increases the ask gap by 0.3 ppts. This is in line with recent studies showing that women have surpassed men in educational attainment. Since the effect of the choice of major is likely already captured by the job title variable added in Column (2), adding the education controls mostly captures the level and quality of education. As evidenced in Table B.2 and described in Section 3.3 women and men have similar work preferences, so adding these controls in Column (4) does not affect the ask gap. Adding employment history in Column (5) takes the gender gap further down to 2.9%. This is mostly driven by the coding skills listed on candidates’ profiles, not by differences in exposure to an

¹⁰ These asks are here weighted by the number of offers received; the unweighted ask gap is larger, at \$8,853.

¹¹ I use the standard assumption that δ and R_{max} are 1.

“elite” tech company in the past. In particular, women are less likely than men to list high-demand coding skills such as JavaScript or Python.¹² Appendix D.1 discusses how the magnitude of the ask gap compares to other related salary measures such as salary expectations or reservation wages. Table IV Column (1) provides information on the coefficients of variables other than the female dummy. These coefficients affect the ask salary in the expected way: more experience and more education are associated with higher asks. For instance, keeping other variables constant, an individual with 2 to 4 years of experience in their current occupation tends to ask for 11.2% more than a candidate with 0 to 2 years of experience in that occupation. In a similar fashion, the coefficient on the employment dummy is positive and significant: all else equal, job-to-job switchers ask for 7.1% higher salaries than candidates who are not currently employed. Finally, more education also leads to higher ask salaries: all else constant, candidates whose highest degree is a PhD ask for 6.7% more than candidates whose highest degree is a master’s.

In Appendix C, a classification analysis using Chernozhukov, Fernández, and Luo (2018) sorted partial effect method highlights that experience is the resume characteristic that captures the greatest share of heterogeneity in ask salaries. Hence, I explore the effects of experience on the ask gap in Figure III Panel (a), which plots the coefficient on the female dummy in Equation 2, controlling for all resume characteristics but estimated separately for different experience groups. The ask gap increases considerably with experience: it is insignificant for the 0-4 years and 4-6 years of experience groups and is only 1.5% for the 6-8 years of experience group. It then jumps to 4% for the 8-15 years of experience group. The largest gap, for candidates with more 15-20 years of experience, reaches 5.4%.¹³

¹² Murciano-Goroff (2022) found that female programmers with previous experience in a programming language were 9.10% less likely than their male counterparts to self-report knowledge of that programming language on their resume. Therefore, it could be that the listed skill gap on Hired.com reflects a gender gap in the propensity to list a programming language, rather than a gap in the actual experience in this language.

¹³ While it is beyond the scope of this paper to explain this gradient, my analysis of the reform described in Section 6 demonstrates that a simple change in the way the website prompts candidates to provide their ask salary narrows the ask gap down to zero, even for candidates with more experience.

5 Descriptive evidence on the role of the ask gap in gender pay inequality

5.1 The gender bid gap: Methodology

Whether the 2.9% residual ask gap relates to the gender pay gap on the platform is an empirical question. Indeed, firms could value skill and experience regardless of what the candidates ask for and we would observe no gender differences in the bids sent by firms to candidates.

To empirically test the relationship between the bid gap and the candidates' resume characteristics and ask salary, I proceed in three steps. First, I estimate the raw gender bid gap. Then, I estimate how much of the bid gap can be explained by the candidates' resume characteristics. Finally, I estimate the effect of the ask salary on the bid gap, with and without the resume characteristics controls. Formally, these three models can be written as:

Model 1:

$$\text{Log}(\text{Bid}_{ib}) = \alpha + \beta_1 \text{Female}_i + \gamma_t + \epsilon_{ib} \quad (4)$$

Model 2:

$$\text{Log}(\text{Bid}_{ib}) = \alpha + \beta_1 \text{Female}_i + \beta_2 X_{ib} + \gamma_t + \epsilon_{ib} \quad (5)$$

Model 3a:

$$\text{Log}(\text{Bid}_{ib}) = \alpha + \beta_1 \text{Female}_i + \beta_3 \text{Log}(\text{Ask}_{ib}) + \epsilon_{ib} \quad (6)$$

Model 3b:

$$\text{Log}(\text{Bid}_{ib}) = \alpha + \beta_1 \text{Female}_i + \beta_2 X_{ib} + \gamma_t + \beta_3 \text{Log}(\text{Ask}_{ib}) + \epsilon_{ib} \quad (7)$$

where $\text{Log}(\text{Bid}_{ib})$ is the b 'th log bid salary received by candidate i . X_{ib} and $\text{Log}(\text{Ask}_{ib})$ are respectively candidate i 's resume characteristics and log ask salary, when he or she receives her b 'th log bid salary. X_{ib} contains the same controls as in Table II Column (5), and γ_t is a Month \times Year FE, where $t = t(i, b)$, the time at which candidate i received bid b .

5.2 The gender bid gap: Results

Graphical evidence Figure A.4 Panel (b) plots kernel density estimates of the distributions of bid salaries, separately by gender. This figure shows that women’s distribution is similarly shaped to men’s, but shifted to the left, such that women receive bids that are, on average, \$5,430 lower than men (\$115,290 vs. \$120,720). Further, comparing Panel (a) (the kernel density estimates of the distributions of ask salaries) to Panel (b) reveals that the ask and bid salary distributions are quite close. This is the first piece of evidence in a pattern I document in this section: firms’ bids closely track individuals’ asks.

Regression results The raw gender bid gap, as estimated by β_1 in Equation 4 and reported in Table III Column (1), is 3.4% and significant at the 1% level. Controlling for the resume characteristics in Column (2) of the same table only takes the gender pay gap down by 33%, to 2.2%.¹⁴ In other words, differences in resume characteristics, such as experience or coding skills, can only account for about a third of the gender bid gap. In contrast, controlling for the ask salary alone in Column (3) eliminates the gender bid gap: the coefficient on the female dummy even becomes positive, although very small (0.2%). This result persists when we add back all the candidate resume characteristics in Column (4): the coefficient on the female dummy remains very close to zero (-0.2%). Finally, we can test whether the effect of the ask salary on the bid salary differs by gender. To do so, Column (5) adds the interaction between the log ask salary and the female dummy. The point estimate of that interacted term is small and insignificant (0.1%), therefore failing to reject the null that men and women realize identical returns to asking for more.

A fundamental challenge in the gender pay gap literature is that the residual gap may not only capture wage differences between otherwise similar men and women, but also the fact that the econometrician is limited in her ability to control for the full information set available to firms. The recruitment process on Hired.com mitigates this concern because firms must formulate their initial bids to candidates before they are able to interact with them. Therefore, the bid salary is solely based on candidates’ resume characteristics and their ask salary and, as a result, having access to candidates’ profiles helps controlling for the firms’ information sets at the time they make

¹⁴ Appendix D.2 compares the residual bid gap to more traditional measures of the gender pay gap in alternative datasets.

their bids.¹⁵ The R^2 in Table III validates this overlap between Hired.com data and the firm’s information sets: the linear model conditioning on candidates’ resume characteristics explains 82% of the variation in bid salaries (Column (2)), while adding the ask salary to the controls raises the R^2 to 0.95 (Column (4)), leaving little room for omitted variable bias.

Figure III shows that the bid gap varies by experience and illustrates how differences in the ask salary can account for this heterogeneity. Figure III Panel (b) plots the coefficient on the female dummy in Equation 5 for different sub-groups of experience. The pattern in this figure mirrors Figure III Panel (a): the bid gap follows the ask gap and increases with experience. However, when we add the ask salary as an explanatory variable in Figure III Panel (c), the heterogeneity in experience disappears. Therefore, the difference in bid gap between more and less experienced women is entirely explained by differences in their asks.

There are two possible explanations for the gap in bid salaries. First, there may be *within-job* bid disparities, that is men and women are offered the same jobs, but women are extended lower bids for these jobs. Alternatively, the gap could come from *between-job* disparities: women, for a given resume, could be offered different, lower-paying jobs. In order to disentangle these channels, I run the same regressions as in the first five columns of Table III but add job fixed effects.

Column (6) of Table III shows that the raw bid gap within jobs is 4.8%. This estimate is larger than the raw bid gap without job fixed effects from Column (1). In other words, in this setting, it is not that women are being offered lower-paying jobs, but rather that, on average, they are offered lower pay for the same job. Once we add resume characteristics (Column (7)), the bid gap narrows to 1.8%. Therefore, for a given job, gender differences in resumes can only explain part of the within-job bid gap. Adding resume characteristics and the ask salary reduces the bid gap to a point estimate very close to zero (-0.3%). This result indicates that the bid gap does not operate through the composition of jobs for which women interview. Similar results hold when we control for firm fixed effects instead of job fixed effects in Table B.6.

Resume characteristics, such as experience, determine the type of jobs (and corresponding salary range) that individuals are selected for, but within jobs, they play a minor role in the determination of pay. This is illustrated by the evolution of the adjusted R^2 in the bid gap regression: while resume

¹⁵ It could still be that firms interpret and interact with the resume characteristics in ways that I cannot account for in this analysis. To get at the causal effect of the ask salary on the bid salary, in Section 6 I leverage a reform that can be interpreted, from the demand side, as an exogenous shift in the ask, and explore its effects on bids.

characteristics explain more than 80% of the total variation in the regressions without job fixed-effects (Table III Column (2)), they can only explain 33% of the total variation within jobs in Table III Column (7). In contrast, adding the ask salary increases the adjusted R^2 to 0.834 in Column (8) of Table III. Taken together, these results indicate that, for a given job, the ask salary plays a much larger role in the determination of the bids than resume characteristics.

5.3 Final offers: results

Given that bid salaries are non-binding, one may worry that the bid gap is not a relevant measure for the actual gender pay gap. To address this concern, Table V presents results on the final offer gap for the restricted sample of candidates that are hired by a company. The left-hand side variable is now $\text{Log}(Final_{ib})$, the salary at which candidate i was hired for the job corresponding to bid b . The right-hand side variables are the exact same as in Table III. The sample of final offers is much smaller than the sample of interview requests (463,860 interview requests were sent out and there were 7,582 final offers) but the point estimates are qualitatively similar. The raw final offer gap is 4.8% (Column (1)) and controlling for resume information leaves a significant 1.4% gap (Column (2)). After adding the ask salary to the resume controls, as in Column (4) of Table III, I find a point estimate for the gender pay gap that is close to zero (-0.9%). These results are insensitive to the addition of firm fixed effects in Columns (6) to (8).

5.4 Sensitivity analysis

In Table III, the relationship between the ask and the bid is estimated on the full sample of bids sent out by companies. However, only a sub-sample of the underlying jobs leads to a final hire. One may argue that only the bids from firms that end up hiring on the platform should be considered, since other firms may not be putting as much effort into their search and bid decisions. To address this concern, in Table B.7, I re-run the same regressions as in Table III but only keep the bids for jobs with a final hire. That corresponds to 43% of the total number of bids. The results are qualitatively the same as in Table III.

Another hypothesis is that there may be two types of firms: the ones that default to the candidate’s ask and the ones that price the job rather than the candidate. To test this idea, in Table B.8, I re-run the regressions from Table III but on the subset of bids that are different from

the ask, which represents 25% of the data. While the results on that sub-sample are qualitatively similar to Table III, the magnitudes vary in the direction predicted by the hypothesis. Indeed, the raw bid gap on that sub-sample is 3.9%, the adjusted gap is 1.6%, and adding the log ask salary narrows it further to 0.3%. In other words, for companies that do not default to the ask, the candidate’s resume explains more of the raw bid gap (59% vs. 33% on the full sample) but the gap remains large and significant, and adding the ask salary still narrows the bid gap to zero.

In addition to the (mandatory) bid, firms have the option to offer equity to the candidate. 44% of the interview requests also contain an equity offer. As evidenced in Table B.5, including equity as a control to the estimation of the bid gap does not alter any of the coefficients, in particular, it does not affect the coefficient on gender.¹⁶

5.5 Gender differences at the extensive margin

Selection into the interview pool The first five columns of Table VI explore whether there are gender differences in the number of bids received during a spell.¹⁷ In Column (1), I regress the number of bids received on a female dummy. Since the number of bids is count data, I also report the Average Marginal Effect (AME) in a Poisson regression on the female dummy at the bottom of each column. The coefficient is significantly negative: women receive about half an offer less than men. However, when adding candidates’ resume characteristics in Column (2), the coefficient on the female dummy flips and becomes small but significantly positive: women get on average 0.2 offers more than men. The fact that the coefficient changed significantly from Column (1) to Column (2) is mainly due to differences in the type of jobs that candidates of different genders are looking for: software engineering jobs, where there is a much higher concentration of men than women, are also the jobs that make a larger number of bids on average. Using a methodology developed in Roussille and Scuderi (2023) to rank firms, I also show, in Appendix G, that once we condition on observables, women and men receive bids from firms of the same rank (\approx quality). One could think that women are getting more bids because they are asking for less. However, Column

¹⁶ In Appendix E, I investigate racial differences in the ask, bid, and final salaries. Because race is self-reported and only a minority (27.6%) of candidates decide to declare it, I caution against drawing definitive conclusions.

¹⁷ Observations here are at the spell level rather than the candidate level. That is, if a candidate used the platform several times over the sample period, each spell is accounted for separately. The candidate controls are the same as in the ask salary estimations (Table II Column (5)), except that I add a control for the length of the spell, which varies between 2 and 6 weeks.

(3) in Table VI shows that adding the ask salary to the controls does not impact the coefficient on the female dummy much and, if anything, the coefficient is larger with the ask salary control. In fact, the ask salary has a small yet positive association with the number of interview requests received. This result may seem a priori surprising: for a given resume, candidates who ask for more are, on average, facing higher demand. Section 7.2 provides a rationale for this result. It's also worth noting that the coefficient on the square of the ask salary is negative (Column (4)). In other words, candidates cannot ask for infinitely more and face ever-growing demand: there is an inflection point after which a higher ask decreases the number of bids that they receive. Finally, Column (5) adds an interaction between the female dummy and the ask salary. The point estimate is insignificant and indistinguishable from zero. At the extensive margin, it is not the case that women are penalized or rewarded more than men for asking for more.

Selection into the final offer pool We now turn to testing whether, after an interview, firms are more or less likely to give the job to a comparable man or woman. In the last three columns of Table VI, the dependent variable is a dummy equal to 1 if a candidate was offered the job for which they interviewed. The gender gap in the probability of getting a final offer after interviewing is insignificant (Column (6)), and neither adding the ask salary (Column (7)) nor including job FE (Column (8)) affects this result. In a nutshell, conditional on interviewing, women are just as likely as men to get the job.

5.6 From descriptive to causal evidence

Introducing the ask salary as a control in Table III Column (4) brings the coefficient on the female dummy to zero. Is this result unique to the female dummy or does introducing the ask salary impact other coefficients? To answer this question, Table IV reports the coefficients on some of the other controls in the gender bid and final gap regressions. Specifically, Column (2) reports the coefficients on education, experience, and employment before adding the ask salary to explain the bid gap, and Column (3) reports them after adding it. Columns (4) and (5) do the same exercise for the final offer gap. This table shows that the coefficient on the female dummy is not the only one that shrinks to zero when adding the ask salary as a control. For instance, the coefficient on the employed dummy falls from 0.043 to 0.003 for the bid and from 0.031 to 0.007 for the final salary, and the magnitude of the coefficients' decrease is similar for education. The coefficients on

dummies for years of experience also decrease although some remain positive. For instance, the coefficient on 15+ years of experience drops from 0.291 to 0.031 for the bid.

This exercise highlights the limit to a causal interpretation of the ask salary on the bid salary in the cross-section analysis of Table III: we would not infer from the results described above that less educated or less experienced candidates are getting lower bids as a result of their lower asks. Rather, we would argue that they are able to command less in the labor market because of their lower education or skill, hence they ask for less. And since the bid and final salaries are highly correlated with the ask, part of the effect of controlling for the ask on resume characteristics such as education or experience is mechanical. With a similar reasoning, the effect of controlling for the ask on the female dummy could be partially mechanical or result from firms' read of the resume characteristics that I cannot fully account for with my resume controls.

In order to make progress on the causal effect of the ask on the bid, I now turn to the analysis of a change on the platform that affected the way some candidates were prompted to report their ask. Specifically, before the reform, the ask salary was an empty field. After the reform, the field was pre-filled with the median of the bid salary in the candidate's labor market cell (defined as the same experience, location, and job title). I leverage this reform for two distinct purposes. First, on the supply side, the reform allows me to investigate whether saliently providing candidates with the median salary in their labor market cell impacts their ask. I find that the reform closes the ask gap, mainly through an increase in women's ask.¹⁸ Second, from the demand side perspective, given that the reform was not announced to the firms, it provides for an exogenous shift in the ask salary of some of the candidates. Therefore, how this shift impacts the bid and final offers made by firms provides for a direct test isolating the impact of the ask on the bid and final salary offers.

6 Closing the gender gap

6.1 Description of the reform

To create their profiles, candidates have to answer the question: "what base salary are you looking for in your next role?". This is what I have referred to as the ask salary. From the first year

¹⁸ This demonstrates that a simple design change can have large effects and allows me to rule out a number of ex-ante plausible explanations for the ask gap, such as signalling different underlying preferences for non-wage amenities.

of data to mid-2018, the answer box for this question was an empty text entry. Starting in mid-2018, the answer box was pre-filled with the median bid salary on the platform over the past 12 months. The median that is shown to the candidate is specific to her combination of desired location, job title, and experience in that job. The change is illustrated in Figure A.5 with a screenshot of the ask salary elicitation web page before and after the reform. This change was motivated by the belief at Hired.com that the platform should provide candidates with a more transparent experience. Even before the reform, candidates could see a histogram of the salaries on the platform. However, the information was somewhat hard to interpret from the histogram since no scale was indicated on the y-axis, neither the median nor the mean were provided and, more substantially, the histogram bins were wide (\$10,000) and therefore did not provide very detailed information on salary choices. The change affected candidates who were either creating or updating a profile. The histogram and median salary were displayed only if Hired.com had enough data to make the calculations for the candidate’s combination of desired location, job title, and experience in that job. Unfortunately, the platform did not track what the threshold for computing the histogram and median was, so I cannot construct a control group for whom the information wasn’t shown. However, because San Francisco software engineer roles are the largest group (25% of the data has this single combination of occupation and location), I received confirmation that this population was fully treated. Therefore, the analysis focuses on San Francisco software engineer roles, comparing candidates who created or updated a profile before the reform with those who did so after the reform. This sample contains more than 40,000 candidates and 200,000 bids.

It is worth highlighting that the reform was not anticipated by either the candidates or the firms. Indeed, the company did not advertise the feature change externally, and therefore new candidates were not drawn to the platform by it. In addition, the feature change only impacted the candidates’ experience on the platform and the firms were not informed of this change at the time it was implemented. Hence, from the perspective of the demand side effects, we can interpret the reform as causing an exogenous shift in the ask of candidates.

6.2 The impact of the reform on the ask salary

Empirical strategy I compare individuals who created a profile before the change and after the change. I first explore the effect of the reform on the ask salary of men and women, as well as on

the ask gap. I follow the literature on Interrupted Time Series (ITS) designs by estimating:

$$\text{Log}(\text{Ask}_i) = \alpha + \beta_0 \text{After}_t + \beta_1 \text{Female}_i + \beta_2 \text{Female}_i \times \text{After}_t + \beta_3 X_i + \gamma_t + \epsilon_i \quad (8)$$

where $t = t(i)$ is the month in which candidate i created her profile, After_t is a dummy equal to 1 after the reform, Female_i is equal to 1 if the candidate is female, and X_i includes the candidate profile controls. γ_t includes a month FE (1 to 12) to capture seasonal effects and a linear time trend (t) to capture the growth of the platform over time. $\text{Log}(\text{Ask}_i)$ is measured at the beginning of the spell. β_0 estimates the effect of the reform on the male ask salary and $\beta_0 + \beta_2$ estimates the effect of the reform on the female ask salary. β_1 estimates the ask gap before the reform while $\beta_1 + \beta_2$ estimates the ask gap post-reform.

This interrupted time series analysis may be misleading if the selection into the platform changed as a result of the reform, in a way that would have led the ask gap after the reform to differ irrespective of the reform. To address this concern, I fit Equation 2¹⁹ in the pre-period to predict the ask salary of every candidate, controlling for all their resume characteristics. I then run this predicted ask against an interacted model of female and after dummies. Results are presented in Table B.9: the coefficient on the interaction between Female and After is exactly zero. In other words, the predicted ask gap is stable across periods. Table B.10 also provides summary statistics on candidates' resume characteristics before and after the change, illustrating the absence of differential selection of men and women onto the platform after the reform.

Graphical evidence Figure IV Panel (a) plots the time series of the mean ask salary for male and female separately, net of a rich set of controls, as in Chetty et al. (2011) and Yagan (2015). Within each month, I first regress the outcome variables on the candidates' resume characteristics. I then construct the two series (Male and Female) by setting each month's difference between the two lines equal to that month's regression coefficient on the female indicator, and setting the weighted average of that month's data points equal to the month's sample average. The figure shows that the female time series tracked the male time series of ask salaries closely in the several months before the feature change, suggesting that the two time series would have continued to evolve in parallel but at significantly different levels in the absence of the feature change. We then observe a

¹⁹ Except that instead of Month \times Year FE, there are just Month FE (1-12) and a monthly linear time trend.

clear jump in female ask salaries to the level of men’s salaries. The narrowing of the gap between the two lines persists several months after the change.

Regression results Table VII, Columns (1) and (2) formalizes the visual evidence in Figure IV Panel (a) by reporting the estimates of Equation 8. Column (1) shows that, in the pre-reform period, the ask gap was 2.9% (the coefficient on the female dummy). In the post period the ask gap, measured as the sum of the coefficient on the female dummy and on the interaction between Female and After, goes to zero. The reform also closes the gap when we consider the bid-weighted version in Column (2). This evolution in the ask gap is led by women asking for more, rather than by men asking for less. In particular, the reform led women to ask for 3.2% more while men continued asking for roughly the same as they would have otherwise. This is also graphically illustrated in Figure V Panels (a) and (b), which show the raw ask salary of candidates, separately by gender, pre-reform (in Panel (a)) and post-reform (in Panel (b)). It appears clearly on these graphs that the cdf of the ask salaries of women is much closer to that of men in the post than in the pre-reform period. Any remaining difference between the two can be explained by gender differences in observables, i.e. women have on average about 2 years less of experience compared to men. This is consistent with the gender imbalance on the platform (more than 80% of candidates are male) and therefore the median that all candidates saw is one of a male candidate.

The absence of bunching Finally, I explore whether candidates bunched at the default median that was suggested to them. Figure V Panels (c) and (d) plots the cumulative distribution function of ask salaries for the 4-6 years experience group,²⁰ respectively for men (Panel (c)) and women (Panel (d)), separately before the reform (solid lines) and after (dashed lines). All candidates in these two panels saw the same median, which is illustrated by the grey line. The first observation is that the distribution for men looks very similar pre and post-reform. Conversely, for women, the cumulative distribution function shifts to the right. The second observation is that the figure does not present clear evidence of bunching at a specific salary, suggesting that candidates did not massively resort to the default setting of the median salary after the reform. Section 7.1 explores the potential mechanisms behind these outcomes.

²⁰ I selected the 4-6 years experience group as an example, but similar patterns can be observed for other groups, with a larger shift for higher experience groups.

6.3 The impact of the reform on the bid salary

Empirical strategy I investigate the effect of the reform on the bid salaries sent by firms in Equation 9:

$$\text{Log}(\text{Bid}_{ib}) = \alpha + \beta_0 \text{After}_t + \beta_1 \text{Female}_i + \beta_2 \text{Female}_i \times \text{After}_t + \beta_3 X_{ib} + \gamma_t + \epsilon_{ib} \quad (9)$$

The controls here are the same as in Equation 8, except X_{ib} can now contain $\text{Log}(\text{Ask}_{ib})$, the ask salary of candidate i when he or she received her b 'th interview request. The dependent variable is the log of the bid salary sent to candidate i for her b 'th interview request. Similar to Equation 8, β_0 will document the effect of the reform on bids received by male candidates and $\beta_0 + \beta_2$ will document the effect of the reform on bids received by female candidates. β_1 estimates the bid gap before the reform while $\beta_1 + \beta_2$ estimates the bid gap after the reform. A similar analysis is then run on the final offers.

Results Table VII Columns (3) to (5) formalizes the visual evidence on the effect of the reform on the gender bid gap in Figure IV Panel (c) by reporting the estimates of Equations 8 and 9. Column (3) reports a 2.5% bid gap before the reform. This gap goes to -0.3% after the reform. This result is driven by the fact that women are offered 2.6% more and men are offered about the same as they would have been offered absent the reform. Controlling for the ask salary in Column (4) narrows the pre- and post-reform bid gaps to small point estimates. The results also hold when we add job fixed effects in Column (5): for a given job, the bid gap was 1.8% before the reform and fell to -0.4% after the reform. Finally, while underpowered, the analysis on the final salary also suggests that the reform closed the final offer gap (Column (6)).

Heterogenous effects of the reform Figure VI plots the effect of the reform on the ask (blue) and bid (red) gaps as a function of the pre-reform gaps, separately by experience groups. In line with the results in Section 4, the pre-reform gender gaps (on the x-axis) are much larger for candidates with more experience. For instance, while the ask and bid gaps before the reform are around 1% for candidates with 0-4 years of experience, they rise close to 5% for candidates with more than 10 years of experience in this occupation. Strikingly, the effect of the reform is also gradually increasing with experience such that changes in women's asks and bids essentially close

the ask and bid gap for all experience groups. The fact that the reform had an effect on the bid and ask gap that is proportional to the pre-reform gap is illustrated in Figure VI by the alignment of all the dots close to the 45 degree line.

The effect of the reform-induced change in ask on the bids Figure VII plots reduced form effects of the reform-induced change in (log) ask salaries on (log) bid salaries (y-axis) against first-stage effects of the reform on (log) ask salaries for gender-by-experience groups. Both sets of effects are estimated via regressions that control for the full vector of resume characteristics. As described in Angrist, Autor, and Pallais (2022), the slope of the line of best fit on this “Visual IV” plot is an IV estimate of the effect of increasing candidates’ asks on the bids they receive, where a dummy for the reform and its interactions with gender-by-experience bins are used as instruments for candidates’ asks. Strikingly, the slope of the fitted line (0.91) is very close to the OLS coefficient (0.85) on the ask salary when regressing the bid salary on the ask salary, controlling for resume characteristics (see Table III Column (4)). This suggests that there was indeed little room for omitted variable bias in the OLS regression, as argued in Section 5.2. In terms of generalisability of this IV slope, it’s important to keep in mind two contextual elements. First, women’s asks were only shifted by a few percentage points and didn’t surpass, on average, those of men. Firms’ response may have been different if the ask changes had been of a larger magnitude. Second, the reform applied to all candidates at the same time and firms’ response to this platform-level change may differ from their response to a single candidate’s ask change.²¹

Other variables The median salary shown to candidates accounts for the candidate’s experience but not for their other resume characteristics (for instance, their education). Therefore, candidates with different education levels but the same experience see the same suggestion. As a consequence, the reform could have impacted the role of other controls in the determination of the ask salary. Table B.11 reports the results of a regression of the log ask salary on all the resume characteristics controls, separately for the pre-reform period (Column (1)) and the post-reform period (Column (2)). It is worth noting that the coefficients of the variables used by Hired.com to determine the median suggested to the job seeker (e.g. experience) increase in the post-reform period. For instance, the coefficient on 2-4 years of experience goes from 0.091 to 0.111, and the one on 10-15

²¹ Another context in which the change in ask is arguably exogenous to firms is when candidates decide to update their ask during a spell. This is a case where a candidate unilaterally decides to change their ask, rather than a platform-level change. The effect on bids of such individual updating is analysed in Appendix F.

years of experience goes from 0.308 to 0.396. In contrast, the coefficients on the other controls, which are not used to compute the median, decrease in magnitude. For instance, the coefficient on Bachelor goes from 0.060 to 0.038, a decrease of similar magnitude is observed for the coefficient on Master. These changes are in line with the fact that candidates from different education levels or schools were exposed to the same median and therefore converged in their ask.

Extensive margin I have just shown that asking for more led to higher bids. However, it could be that this positive outcome comes at the expense of other dimensions in the recruitment process. For instance, women could get fewer interview offers as a result of the feature change.²² I explore several measures of the effect of the reform on: the number of bids received by a candidate during a spell k , the time it takes to receive the first bid during a spell k ,²³ the likelihood of getting a final offer, as well as the rank of the firms that bid and make offers to the candidate (see Appendix G for more details on how these ranks are computed). Table VIII presents the results of this analysis. First, Table VIII Column (1) runs the number of bids received by candidates on the Female dummy, the After dummy, and their interaction, as well as the same controls as in Column (1) of Table VI. The coefficient on the interacted term Female \times After is 0.19 (95% CI -0.17 to 0.55, mean = 4.8). Column (2) estimates the number of hours it takes for a candidate to get a first bid. Again, the point estimate for the coefficient Female \times After is very small (95% CI -8 to 9, mean = 62). Column (3) estimates the likelihood of getting an offer on the platform and, while admittedly imprecise as there are few final offers made, the point estimate for the coefficient on Female \times After is close to zero and insignificant (CI -0.011 to 0.024, mean = 0.09). In Columns (4) and (5) I show that the reform has not significantly altered the rank of firms that contact women (CI -0.3 to 0.9, mean = 62.5) or make a final offer to them (CI -1.8 to 1.8, mean = 62.9). Taken together, these results suggest that women face little or no penalty for demanding wages comparable to men’s.

²² Note that the total number of bids received at any given time depends on factors such as the growth of the platform and the demand for software engineers at that time. Therefore, the interrupted time series design is not well-suited to assess the general equilibrium effect of the reform on the total number of bids sent on the platform. However, I can still credibly observe whether the reform had a differential impact on several extensive margin variables.

²³ The specification is the same as in Equation 8 except the left side respectively becomes $Nbids_{ik}$ and $Hours_{ik}$, as defined in Section 5.5 and we add the length of the candidate’s spell (2 to 6 weeks) to the controls.

7 Discussion

The new ask elicitation framing led women to ask for more and firms to correspondingly bid more on them. Women also do not seem to be penalized, compared to men, at the extensive margin. Two questions arise from these results. On the candidate side, what mechanism could rationalize the fact that the new framing led women to ask for more? On the company side, why is it that firms are not decreasing their demand for female labor, compared to men?

7.1 Why do women ask for more in response to the default median?

Several reasons can be raised to explain why women were asking for less in the first place. The fact that the treatment closes the gender ask gap allows me to corroborate some of these reasons and eliminate others. Let us start with the possible explanations for the lower initial female ask salary that do not square with the reform effects. First, women could initially have actively been playing a different strategy than men. For instance, they could have been trying to signal different unobservables, such as the need for more flexible hours. Alternatively, they could have been asking for less so as to increase their chances of getting a job.²⁴ Finally, women could be less confident than men about their unobserved ability; therefore believing, for a given resume, that they are worth less than their male counterparts. But, if women were knowingly playing a different ask salary strategy then gender differences in ask salaries should have remained different even after the treatment. Further, the fact that men and women do not meaningfully differ in their preferences over firm characteristics (see Table B.2, described in Section 3.3) also casts doubt on a story where gender gaps in tastes for non-wage amenities drive differences in ask salaries. An alternative explanation for why women initially ask for less would be that they have downward biased beliefs about how much they can ask for, compared to men. This downward bias may have had two sources: (1) downward biased beliefs about the market wage for their resume, (2) anticipated gender discrimination, which would lead to lower asks to mitigate it. While I do not have definitive evidence to adjudicate between the reform having a (1) pure information channel vs (2) a norm-based explanation, one critical piece of evidence points toward the former rather than the latter. Indeed, the absence of bunching at the suggested ask, as illustrated in Figure V and discussed in the previous section,

²⁴ This would be in line with experimental evidence that women are more risk-averse (see Croson and Gneezy (2009)).

makes the norm-setting power of Hired.com an unlikely explanation: if we thought women used the suggested ask as a signal for what an “appropriate” ask is, we would have expected bunching at that number, but they do not. To understand how the treatment could generate this outcome, consider this simple heuristic: candidates form beliefs about their percentile in the quality distribution, then make assumptions and/or obtain information about the salaries in their field, and finally choose an ask in this distribution that corresponds to their quality percentile. The treatment effect would then be consistent with downward-biased beliefs about the median salary that the treatment corrected. Salary information, however, does not shift beliefs about the position in the quality, hence does not shift the variance in women’s asks.²⁵

Finally, if, as documented in the networks literature (McPherson, Smith-Lovin, and Cook (2001)), there is gender homophily in information networks and such group-specific homophily leads to frictions in the updating of beliefs (Golub and Jackson (2012)), we can explain two dimensions of heterogeneity in initial gender ask gaps. First, the fact that the gender ask gap is larger in labor markets (location \times job) where the share of women is smaller (as documented in Table B.12). Second, the fact that the gender ask gap is larger for more experienced women: the attrition in the share of women in manager positions will also restrict the pool from which experienced women get their information compared to men.

The fact that information asymmetries, rather than psychological traits, explain the initial ask gap is consistent with recent evidence from the behavioral literature. For instance, Dreber, Heikensten, and S  ve-S  derbergh (2022) run a survey on a representative sample of recent graduates in Sweden to shed light on the mechanism behind women’s lower ask. The paper finds suggestive evidence that beliefs about the wage an ideal candidate would ask for, but not perceived social cost or confidence, can explain most of the 2.5% gender gap in salary requests.

7.2 The ask salary as a signal of quality

A second question that the reform effects raises is the following: why are firms not decreasing their demand for female labor, compared to men, in response to the increase in women’s ask after the reform? This section provides a framework to better understand this ex-ante surprising result.

²⁵ Consistent with this interpretation, but in a different context, Coffman (2014) shows that a woman’s reluctance to contribute her idea to a group, especially in gender-incongruent areas, is largely driven by self-assessments, rather than fear of discrimination.

I first investigate, descriptively, the relationship between the number of bids received and the residual ask for all candidates. Figure A.6 documents a bell-shape relationship: For residual log ask salaries between -0.7 and 0.15, the number of bids received increases with the ask. Beyond 0.15, the relationship becomes negative, that is asking for more is associated with a lower number of bids received. The existence of an upward-sloping range can be rationalized by the following idea: firms interpret the ask salary as a signal of unobserved quality.²⁶ When deciding whether to send an interview request to a candidate, the firm considers the trade-off between the final salary it will have to pay the candidate and the expected return to the match. For a given set of resume characteristics, the expected return to the match is increasing in the quality of the candidate. While the firm cannot directly observe this quality before interviewing the candidate, the ask sends a positive signal about it.

The ask salary therefore plays an ambiguous role in the firm’s decision to interview the candidate. On the one hand, firms predict that a higher ask leads to a higher final offer. On the other hand, a higher ask is a signal of unobserved quality and therefore a higher return to the match. The relative size of these effects determines the sign of the relationship between the ask and the probability of getting an interview request from any given firm.

The idea of price as a signal of quality, while under-studied in the context of wage bargaining, has been theorized for consumer products in the fields of IO and game theory. Seminal papers in this literature (Wolinsky (1983), Milgrom and Roberts (1986)) study conditions under which product price or some combination of price and another quality signal, such as advertising, can effectively signal product quality when consumers are not fully informed.

In Appendix H, adapting Wolinsky (1983)’s model to the labor market, I propose a framework to explain how, in a context of imperfect information about a candidate, a separating equilibrium in which the candidate’s ask salary is a signal of their quality can exist. The intuition for the equilibrium in this model can be summarized as follows. For a given ask salary, firms expect a certain unobserved quality of the candidate. A candidate that asks for a given salary may be of

²⁶ Using the methodology developed in Roussille and Scuderi (2023) to calculate firm’s productivity, Figure G.1 plots the relationship between the average (normalized) productivity of firms and the residualized log ask salary of candidates. There is a clear, increasing relationship between the residual ask salary candidates list and the mean productivity of firms that bid on those candidates: candidates with higher residual asks tend to receive bids from more productive firms. This provides additional support for the idea that firms interpret the ask salary as a signal of unobserved quality.

lower quality, but information revealed during the interview will enable some prospective firms to find this out and, provided there are competing candidates, they will not hire this one. Therefore, in deciding whether to ask for a higher salary than what the firm expects given their quality, the candidate weighs the decrease in their chances of being hired against the gain in salary in the event they get an offer. If the chances of detection are large enough to outweigh the potential salary gains, it is best for the candidate to signal their true quality.

Firms differ in the candidates' quality-salary combination that maximises their expected profit. I model this as firms having a different match-productivity parameter: the match with a high-quality candidate has a higher return to the firm if the job involves complex tasks. In equilibrium, candidates receive interview requests from their ideal firm type, that is the type that is willing to pay them the most for their quality. Therefore, whether a higher candidate's ask is associated with more or less interview requests entirely depends on the empirical distribution of firm types on the platform. As explained in the model section H.5, we can approximate a given firm type by estimating the range of residual asks that it interviews in. Figure H.2 shows this relationship is also bell-shaped, providing further theoretical foundations to my empirical findings.

In this model, women have downward biased beliefs about the salary they can ask for that stem from inaccurate information about the equilibrium but firms do not learn about these biases because interviews go equally well for men and women. This feature comes from the signal design: it can only provide firms with a "red flag", that is whether the candidate is below her expected quality. But, in equilibrium, neither men nor women end up raising this flag since candidates of both genders either are of the quality they signal (men), or above (women).

We can now return to our initial question, namely why is it that firms are not decreasing their demand for female labor, compared to men, in response to the increase in women's ask after the reform? The model now provides an answer to this: if firms interpret women's higher ask as a signal of better quality, their demand for women does not necessarily decrease. Their demand for women may even increase if the women whose ask was shifted up by the reform are in the increasing region of Figure A.6. Columns (6) to (8) in Table VIII investigate this hypothesis. First, Column (6) adds the ask salary and ask salary squared to Column (1). This addition pushes the coefficient on the interaction between the Female and After dummy from 0.190 to 0.037. Therefore, the small estimated increase in the number of bids received by women post-reform is entirely explained by

their increased ask salary. The dependent variable in Column (7) is the predicted number of bids received using the specification in Column (1) on the pre-period. The coefficient on the interaction between the female and after dummies is 0.034 and insignificant. This confirms that, aside from their ask salary, women pre and post-reform do not differ in their likelihood of getting a bid based on their resume. Finally, the dependent variable in Column (8) is the predicted number of bids received using the specification in Column (6) in the pre-period. The coefficient on the interaction is now 0.178 (positive but insignificant). The fact that this coefficient is between that of Column (1) of that of Column (6) is consistent with the statistical relationship between the number of bids received and the ask salary being structural. Further, this positive coefficient indicates that the women whose ask was shifted up by the reform are in the increasing region of Figure A.6, which explains why they do not face a penalty for asking for more.

8 Conclusion

This paper introduces the gender ask gap to the gender pay gap literature. Using novel data from a leading recruitment platform, I document a 2.9% adjusted gender ask gap for a large sample of high-wage workers in the tech industry. This gap is statistically significant and economically meaningful: it represents, on average, \$3,830 in annual salary. The 3.4% raw bid gap can entirely be explained by the ask gap: solely controlling for the ask salary, the bid gap falls to 0.2%. Conversely, controlling for the candidates' resume characteristics only narrows the bid gap by 33%. These results qualitatively carry through to the 7,582 final salary offers for the sub-sample of hired candidates. On this platform, women are not discriminated against at the extensive margin. In particular, conditional on their resume characteristics, women in fact receive slightly more bids than men and, conditional on interviewing, women are just as likely as men to get a final offer. Finally, I show that a reform wherein candidates saw their ask salary field pre-filled with the median value of bids for similar candidates changed the adjusted ask gap from 2.9% to -0.6%, and similarly changed the adjusted bid gap from 2.5% to -0.3%. Yet the number of bids received by women, compared to men, or their likelihood of getting a final offer was not affected. This suggests that there is little penalty to asking for more. These results were obtained in the context of well-documented labor

supply shortages and high levels of competition between employers for qualified workers.²⁷ Given recent lab-based evidence that cautions against “lean in” recommendations (Exley, Niederle, and Vesterlund (2020)), a better understanding of the contexts and conditions under which asking for more benefits rather than harms women is an important avenue for future research.

References

- Agan, Amanda, Bo Cowgill, and Laura Katherine Gee. 2020. “Do Workers Comply with Salary History Bans? A Survey on Voluntary Disclosure, Adverse Selection, and Unraveling.” *AEA Papers and Proceedings* 110 (May): 215–219.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools.” *Journal of Political Economy* 113 (1): 151–184.
- Angrist, Joshua, David Autor, and Amanda Pallais. 2022. “Marginal Effects of Merit Aid for Low-Income Students.” *The Quarterly Journal of Economics* 137, no. 2 (May): 1039–1090.
- Babcock, Linda, and Sara Laschever. 2006. *Why Women Don’t Ask: The High Cost of Avoiding Negotiation and Positive Strategies for Change*. Bantam.
- Babcock, Linda, Sara Laschever, Michele Gelfand, and Deborah Small. 2003. *Nice Girls Don’t Ask*. Harvard Business Review.
- Baker, Michael, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar. 2023. “Pay Transparency and the Gender Gap.” *American Economic Journal: Applied Economics* 15 (2): 157–183.
- Barach, Moshe A., and John J. Horton. 2021. “How Do Employers Use Compensation History? Evidence from a Field Experiment.” *Journal of Labor Economics* 39 (1): 193–218.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014. “Inference on Treatment Effects after Selection among High-Dimensional Controls.” *The Review of Economic Studies* 81 (2): 608–650.

²⁷ The unemployment rate of U.S. tech workers had hit a record low in the study period: [The Unemployment Rate for U.S. Tech workers Just Hit the Lowest Number Ever Recorded](#).

- Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon. 2022. “Do Firms Respond to Gender Pay Gap Transparency?” *The Journal of Finance* 77 (4): 2051–2091.
- Bergerhoff, Jan, Lukas Kiessling, Pia Pinger, and Philipp Seegers. 2021. “Gender Differences in Wage Expectations: Sorting, Children, and Negotiation Styles.” *IZA Discussion paper*, no. 12522.
- Bertrand, Marianne. 2017. “The Glass Ceiling.” *Coase Lecture, London School of Economics, March 2017*.
- Biasi, Barbara, and Heather Sarsons. 2022. “Flexible Wages, Bargaining, and the Gender Gap.” *The Quarterly Journal of Economics* 137 (1): 215–266.
- Blau, Francine D., and Lawrence M. Kahn. 2017. “The Gender Wage Gap: Extent, Trends, and Explanations.” *Journal of Economic Literature* 55(3):789–865.
- Bowles, Hannah Riley, Linda Babcock, and Kathleen L. McGinn. 2005. “Constraints and Triggers: Situational Mechanics of Gender in Negotiation.” *Journal of Personality and Social Psychology* 89 (6): 951–965.
- Chamberlain, Andrew, and Jyotsna Jayaraman. 2017. “The Pipeline Problem: How College Majors Contribute to the Gender Pay Gap.” *Glassdoor Economic Research*.
- Chernozhukov, Victor, Ivàn Fernández, and Ye Luo. 2018. “The Sorted Effects Method: Discovering Heterogeneous Effects Beyond Their Averages.” *Econometrica* 86 (6): 1911–1938.
- Chetty, Raj, John N. Friedman, Emmanuel Saez Nathaniel Hilger, Diane Whitmore Schanzenbach, and Danny Yagan. 2011. “How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR.” *The Quarterly Journal of Economics* 126 (4): 1593–1660.
- Coffman, Katherine. 2014. “Evidence on Self-Stereotyping and the Contribution of Ideas.” *The Quarterly Journal of Economics* 129 (4): 1625–1660.
- Cortés, Patricia, Jessica Pan, Ernesto Reuben, Laura Pilossoph, and Basit Zafar. Forthcoming. “Gender Differences in Job Search and the Earnings Gap: Evidence from the Field and Lab.” *The Quarterly Journal of Economics*.

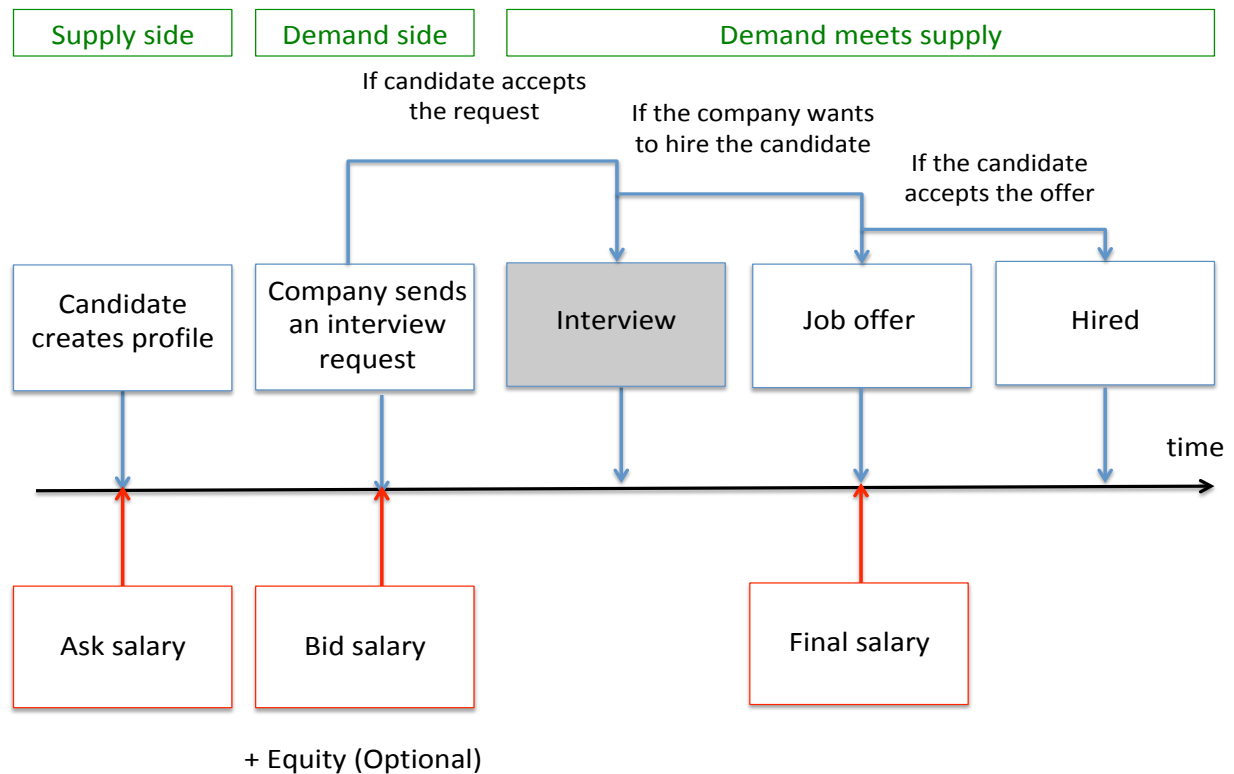
- Croson, Rachel, and Uri Gneezy. 2009. "Gender Differences in Preferences." *Journal of Economic Literature* 47 (2): 448–474.
- Cullen, Zoë, and Bobak Pakzad-Hurson. 2023. "Equilibrium Effects of Pay Transparency." *Econometrica* 91 (3): 765–802.
- Dreber, Anna, Emma Heikensten, and Jenny Säve-Söderbergh. 2022. "Why Do Women Ask For Less?" *Labour Economics* 78:102204.
- Exley, Christine L., Muriel Niederle, and Lise Vesterlund. 2020. "Knowing When to Ask: The Cost of Leaning In." *Journal of Political Economy* 128 (3): 816–854.
- Exley, Christine L, and Judd B Kessler. 2022. "The Gender Gap in Self-Promotion*." *The Quarterly Journal of Economics* 137, no. 3 (January): 1345–1381.
- Garbinti, Bertrand, Jonathan Goupille-Lebret, and Thomas Piketty. 2018. "Income inequality in France, 1900-2014: Evidence from Distributional National Accounts (DINA)." *Journal of Public Economics* 162:63–77.
- Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104 (4): 1091–1119.
- Golub, Benjamin, and Matthew O. Jackson. 2012. "How Homophily Affects the Speed of Learning and Best-Response Dynamics." *The Quarterly Journal of Economics* 127 (3): 1287–1338.
- Gomez-Herrera, Estrella, and Frank Mueller-Langer. 2019. "Is There a Gender Wage Gap in Online Labor Markets? Evidence from Over 250,000 Projects and 2.5 Million Wage Bill Proposals." *CESifo Working paper*, no. 7779.
- Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer. Forthcoming. "Worker Beliefs About Outside Options." *The Quarterly Journal of Economics*.
- Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet. 2021. "Gender Differences in Job Search: Trading off Commute against Wage." *The Quarterly Journal of Economics* 136 (1): 381–426.

- Leibbrandt, Andreas, and John A. List. 2015. "Do Women Avoid Salary Negotiations? Evidence from a Large-Scale Natural Field Experiment." *Management Science* 61 (9): 2016–2024.
- Litman, Leib, Jonathan Robinson, Zohn Rosen, Cheskie Rosenzweig, Joshua Waxman, and Lisa M. Bates. 2020. "The Persistence of Pay Inequality: The Gender Pay Gap in an Anonymous Online Labor Market." *PLoS ONE* 15(2).
- McPherson, Miller, Lynn Smith-Lovin, and James M Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27 (1): 415–444.
- Milgrom, Paul, and John Roberts. 1986. "Price and Advertising Signals of Product Quality." *Journal of Political Economy* 94 (4): 796–821.
- Murciano-Goroff, Raviv. 2022. "Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers." *Management Science* 68 (5): 3262–3281.
- Olivetti, Claudia, and Barbara Petrongolo. 2016. "The Evolution of Gender Gaps in Industrialized Countries." *Annual Review of Economics* 8 (1): 405–434.
- Reuben, Ernesto, Matthew Wiswall, and Basit Zafar. 2017. "Preferences and Biases in Educational Choices and Labour Market Expectations: Shrinking the Black Box of Gender." *The Economic Journal* 127:2153–2186.
- Rigdon, Mary L. 2012. "An Experimental Investigation of Gender Differences in Wage Negotiations" (October).
- Roussille, Nina, and Benjamin Scuderi. 2023. "Bidding for Talent: A Test of Conduct on a High-Wage Labor Market." *Working paper*.
- Small, Deborah, Michele Gelfand, Linda Babcock, and Hilary Gettman. 2007. "Who Goes to the Bargaining Table? The Influence of Gender and Framing on the Initiation of Negotiation." *Journal of Personality and Social Psychology* 93 (4): 600–613.
- Visier and Insights. 2017. "The Truth About Ageism in the Tech Industry." *Report*.
- Wolinsky, Asher. 1983. "Prices as Signals of Product Quality." *The Review of Economic Studies* 50 (4): 647–658.

Yagan, Danny. 2015. “Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut.” *American Economic Review* 105 (12): 3531–3563.

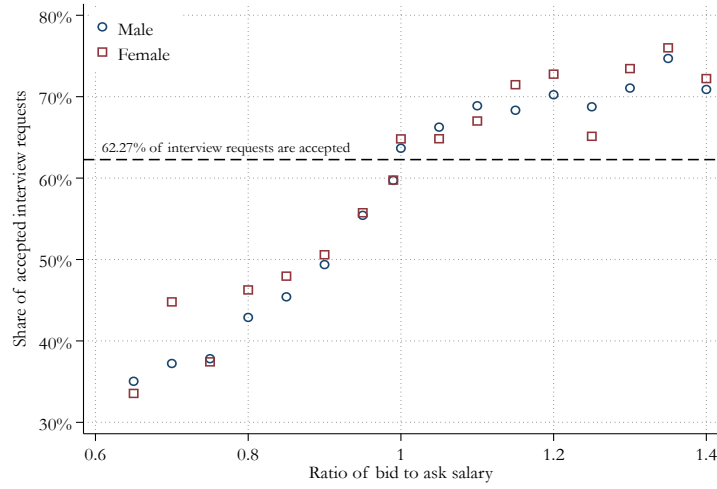
Figures

Figure I: Timeline of the Recruitment Process on Hired.com

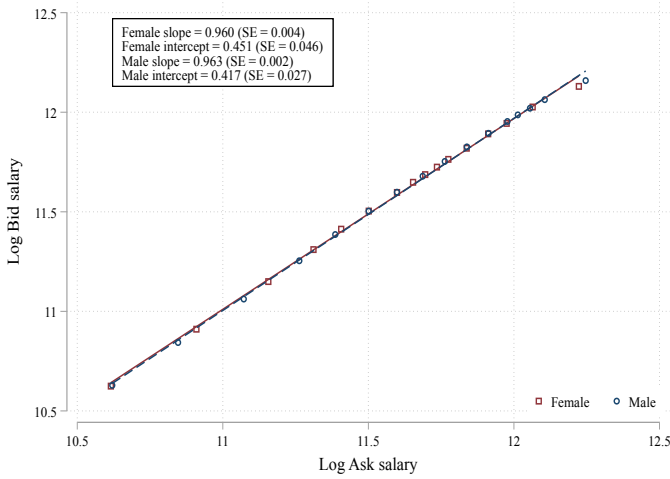


Note: This figure shows the timeline of recruitment on Hired.com. In red boxes are the different salaries that are captured on the platform. The blue boxes describe all the steps of recruitment on the platform, from profile creation to hiring. The grey shading for the interview stage indicates that I do not have metadata from companies about their interview process. In green is the classification of the recruitment process between labor demand side (companies) and labor supply side (candidates).

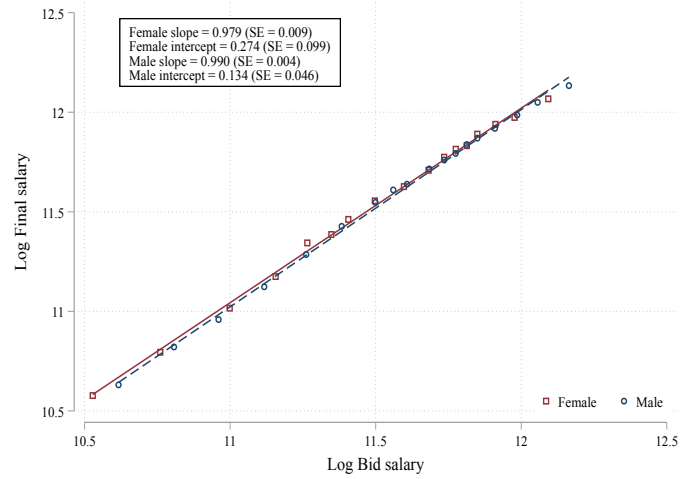
Figure II: Interview Request Acceptance Rate and the Relationship between Ask, Bid and Final Salary



(a) Interview request acceptance rate as a function of the bid to ask ratio



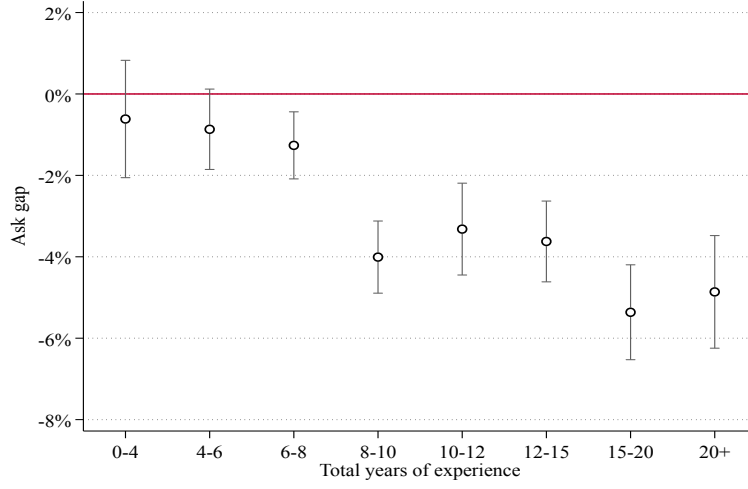
(b) The relationship between log bid and log ask salary



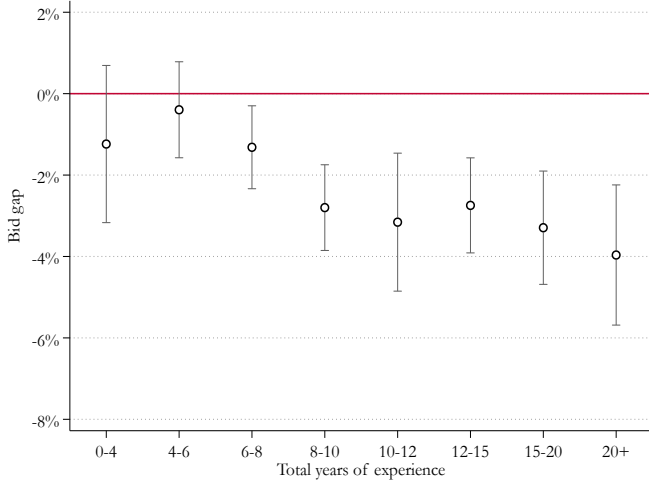
(c) The relationship between log final and log bid salary

Note: Figure II Panel (a) shows how the share of accepted interview requests changes with the ratio of bid to ask salary, separately for male and female candidates. Observations are grouped into bins of $\frac{bid}{ask}$ of length 0.05, except $\frac{bid}{ask} = 1$, which is plotted separately. This panel includes, for each candidate, the first five bids received to ensure that the candidate is active and available for interviews on the platform at the time he or she receives the request. This figure also shows the close relationship between the log ask and log bid salary in Figure II Panel (b) and the log bid and log final salary offers in Figure II Panel (c). They report these relationships separately for male (solid blue line) and female (dashed red line) candidates. The difference in the relationships between salaries is not significant by gender. Standard errors are clustered at job and individual levels and the binned scatter plots have 16 equally sized bins of observations. Overall, 77% of bid salaries are identical to the corresponding ask salary and 90% of bid salaries are within a range of 10k USD from the ask, while 36% of final salaries match the initial bid exactly and 78% of final salaries are within a range of 10k USD from the bid. Figure II Panel (b) includes the 463,860 observations with an associated bid and Figure II Panel (c) the 7,582 observations for which there is a final offer.

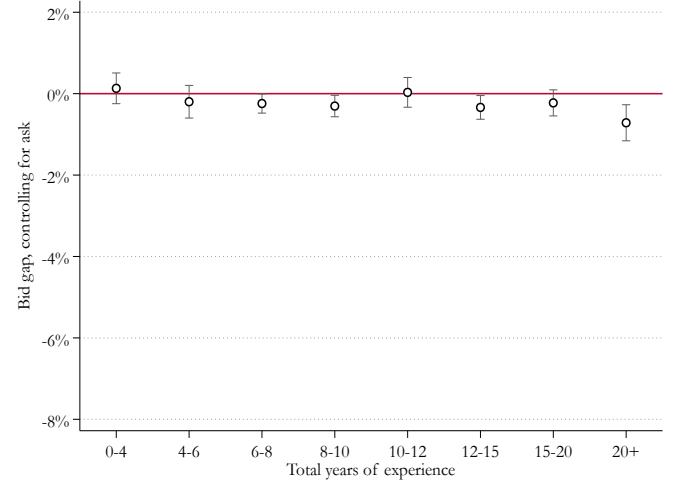
Figure III: Heterogeneity in the Ask and Bid Gap by Experience



(a) Residual Ask gap - resume characteristics



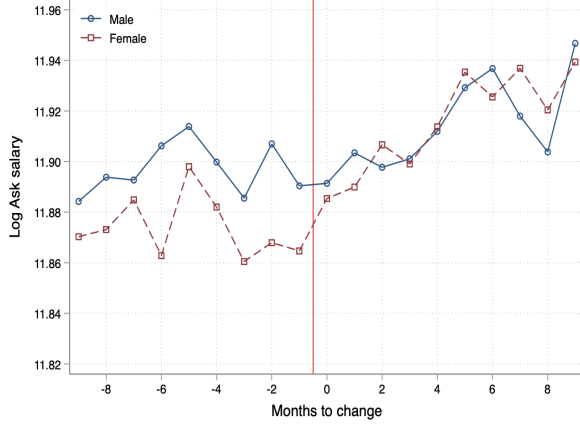
(b) Residual Bid gap - resume characteristics



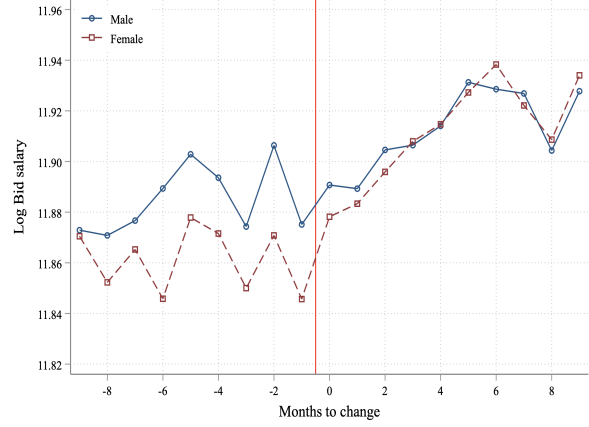
(c) Residual Bid gap - resume characteristics + ask salary

Note: These figures show the heterogeneity in the ask gap by experience as well as the importance of the ask salary in explaining the bid gap, separately by experience. Figure III Panel (a) plots the point estimate of the female dummy in Equation 2, where the regression is run separately by total years of experience. Figure III Panel (b) plots the point estimate on the female dummy in Equation 5 and Figure III Panel (c) plots the point estimate on the female dummy in Equation 7. In all figures, regressions are run separately for each group of total years of experience. The resume characteristics I control for are all the variables described in Table B.1, except the Total Position experience since regressions are run separately for each Total Position experience group.

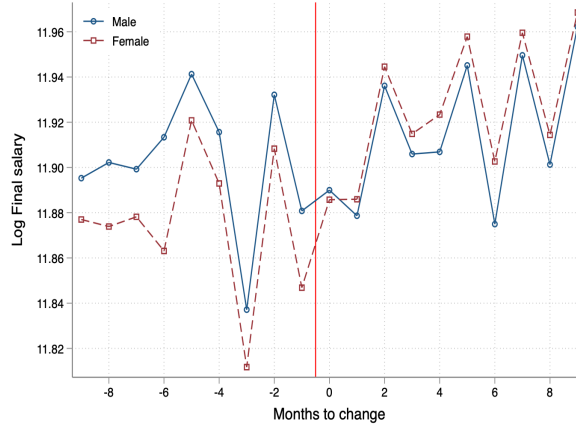
Figure IV: Effect of the Reform on the Gender Ask and Bid Gaps



(a) Log ask salary - all resume controls



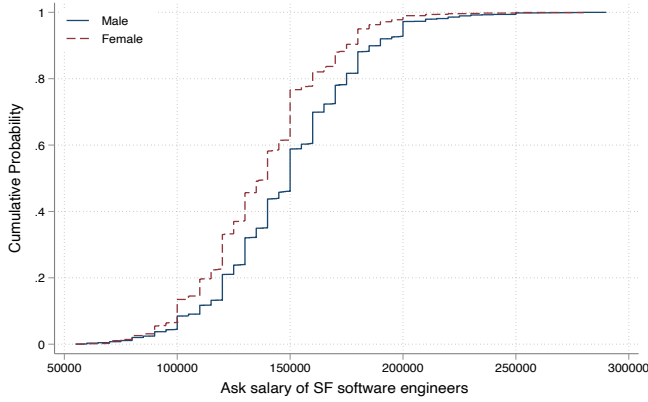
(b) Log bid salary - all resume controls



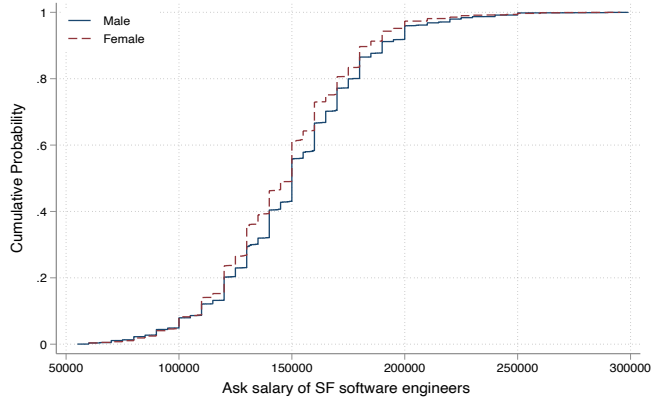
(c) Log final salary - all resume controls

Note: These figures plot the time series of annual mean salary for men and women, net of all resume characteristics. Each panel is constructed regressing the outcome variable (either log ask salary for Figure IV Panel (a), log bid salary for Figure IV Panel (b) or log final salary for Figure IV Panel (c)) within every month on a female indicator and the resume controls, requiring that the vertical distance between the two lines equals the regression coefficient on the female indicator and that the weighted average of the lines equals the sample average in that month. The ask salary regressions are bid-weighted (each observation is weighted by the number of bids received).

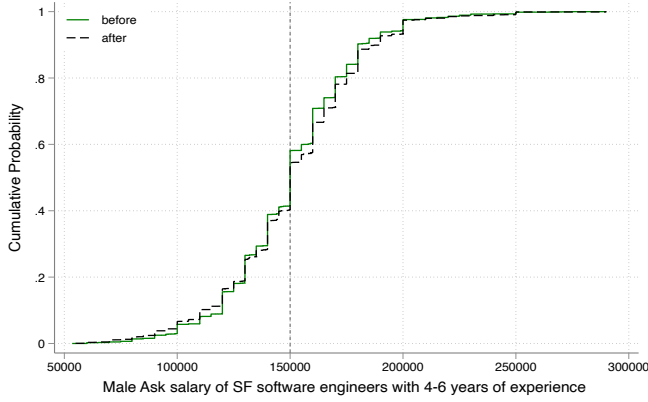
Figure V: Cumulative Distribution Function of Candidates' Ask Salaries before and after the Reform



(a) Pre-Reform Distribution of Ask Salaries by Gender



(b) Post-reform distribution of ask salaries by gender



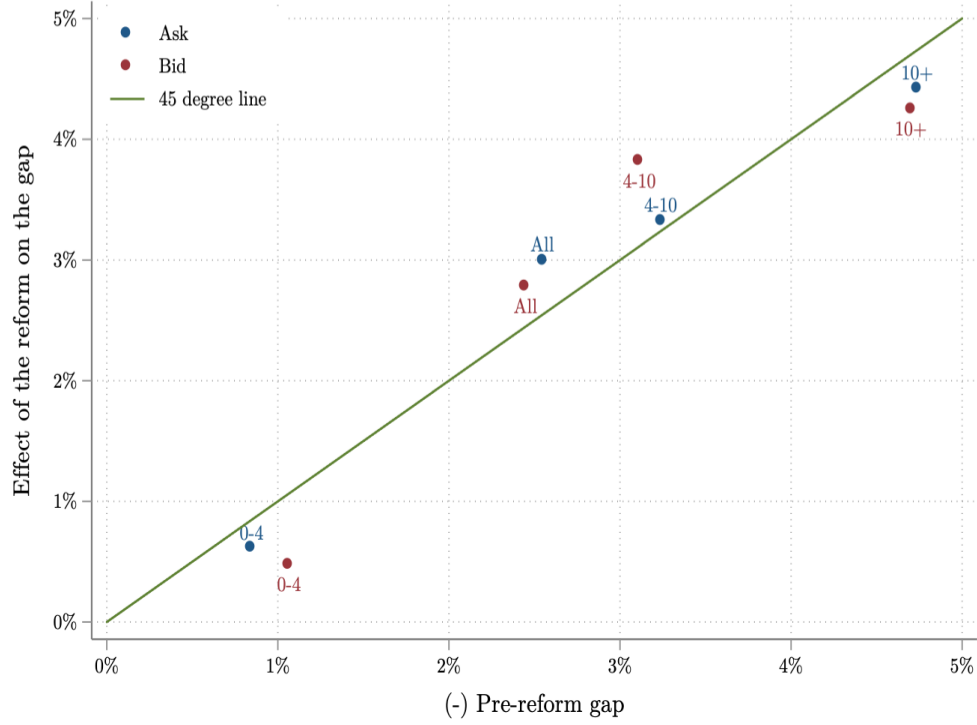
(c) Distribution of ask salaries for men in the 4-6 years of experience group



(d) Distribution of ask salaries for women in the 4-6 years of experience group

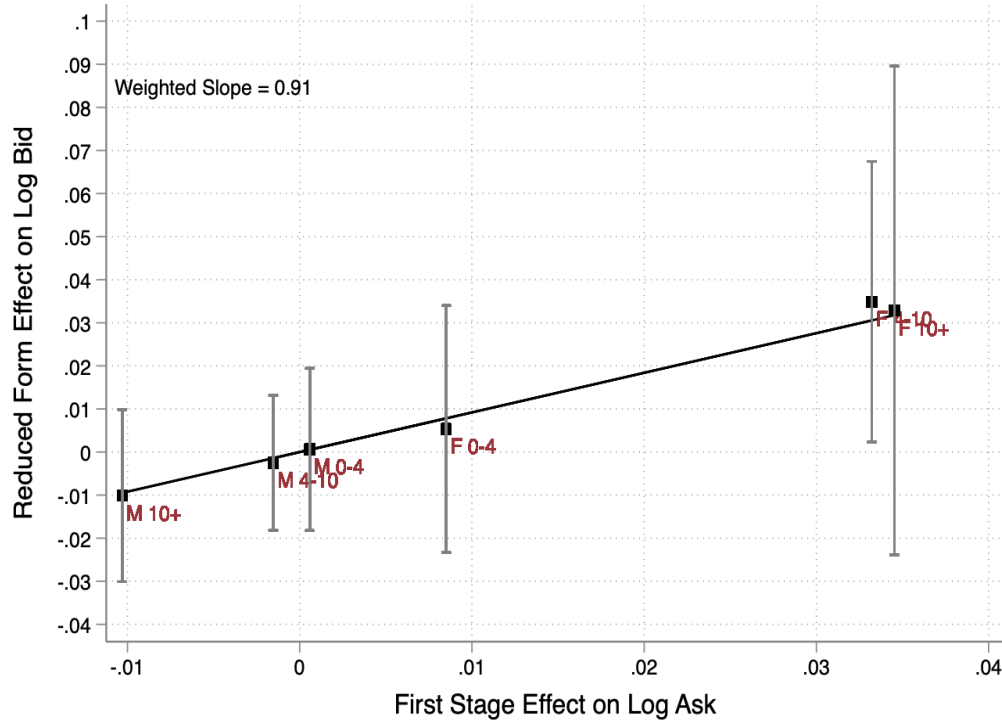
Note: Figure V Panels (a) and (b) show the raw distribution of the ask salary separately for male (in solid blue) and female candidates (in dashed red), respectively pre and post-reform. Figure V Panels (c) and (d) plot the cumulative density of ask salaries, separately, for male and female respectively, before (full green line) and after (dashed black line) the reform, for candidates in the 4-6 years of experience group. Given that salary suggestions are made at the experience level, all candidates with a given experience have seen the same suggestion. The exact median that was shown was not recorded but the grey line approximates it using the past 12 months of bids for the corresponding experience. The before period is limited to 12 months for better comparability of ask salaries.

Figure VI: The Effect of the Reform on the Bid and Ask Gap as a Function of the Pre-Reform Gaps



Note: This figure plots the effect of the reform on the bid and ask gaps as a function of the pre-reform gap, separately for three terciles of experience groups. The x-axis is the coefficient on the female dummy in Equation 8 for the ask - except the observations are weighted by the number of bids received - and Equation 9 for the bid. The y-axis is the coefficient on the Female \times After dummy in the same equations, respectively. Regressions are run separately for each experience group.

Figure VII: Estimates of the Effect of the Reform-Induced Change in Asks on the Bids



This figure plots reduced form effects of the reform-induced change in (log) ask salaries on (log) bid salaries (y-axis) against first-stage effects of the reform on (log) ask salaries for gender-by-experience groups. Both sets of effects are estimated via regressions that control for the full vector of resume characteristics. As described in Angrist, Autor, and Pallais (2022), the slope of the line of best fit on this “Visual IV” plot is an IV estimate of the effect of increasing candidates’ asks on the bids they receive, where a dummy for the reform and its interactions with gender-by-experience bins are used as instruments for candidates’ asks. The regression line is constrained to go through zero and estimated weighting by bid-level experience group-sizes. Whiskers mark 95% confidence intervals.

Tables

Table I: Descriptive Statistics on Candidates and Companies

<i>Panel A: Descriptive Statistics on Candidates</i>					
	All	Male	Female	Difference	p-value
Number of Candidates	113,777	76,223	19,998	56,225	
Average number of bids received per candidate	4.5	4.6	4.2	0.4	0.000
Probability of accepting an interview request	62.2	62.0	63.2	-1.2	0.000
Education					
Share with a bachelor	97.6	97.3	98.7	-1.4	0.000
Share with a master	41.4	40.3	45.2	-4.9	0.000
Share with a CS degree	55.2	57.2	47.7	9.5	0.000
Share with an IvyPlus degree	9.4	8.7	11.8	-3.1	0.000
Preferences					
Share looking for full time job	96.9	96.7	97.7	-1.0	0.000
Share looking for a job in San Francisco	31.6	30.0	37.5	-7.5	0.000
Share in need of visa sponsorship	13.6	13.0	15.7	-2.7	0.000
Work History					
Years of total experience (Mean)	11.3	11.7	10.1	1.6	0.000
Share that worked at a FAANG	6.0	6.0	6.0	0.0	0.679
Share leading a team	32.7	33.8	27.6	6.2	0.000
Share employed	73.1	74.0	69.7	4.3	0.000
Number days unemployed (Median)	120	116	133	-17	0.000
Occupation					
Share of software engineers	61.7	66.6	43.2	23.4	0.000
Share of web designers	8.3	6.1	16.6	-10.5	0.000
Share of product managers	8.3	7.5	11.4	-3.9	0.000
<i>Panel B: Descriptive Statistics on Companies</i>					
Number of:	Firms	Jobs	Bids sent	Final offers	Cities
	6,532	39,839	463,860	7,582	20
Revenue (yearly, in Million USD)	1-25	26-100	101-500	501-1,000	1,000+
Share (N = 962)	47%	17%	12%	14%	10%
Firm Age (in years)	0-5	6-10	11-15	16-20	20+
Share (N = 2,249)	36%	45%	11%	4%	4%
Firm Size (Nb. Employees)	1-10	11-50	51-200	201-500	500+
Share (N = 2,368)	18%	29%	31%	11%	11%
Top 3 Locations	SF	NY	LA		
Share (N = 4,319)	40%	24%	7%		
Top 3 Industries	Software	Finance	Analytics		
Share (N = 2,253)	15%	10%	8%		

Note: Panel (a) shows descriptive statistics for candidates in the sample (first column), separating them by gender (second and third columns) and reporting the difference between males and females (fourth and fifth columns). The average number of bids received and the probability of accepting are computed on the sample of candidates that receive at least one bid. The median number of days unemployed is computed conditional on being unemployed. Panel (b) shows descriptive statistics on the company side (number of firms, jobs, bids sent, and final offers sent) as well as on firm characteristics for a subsample of companies on Hired.com. The share of each category is reported.

Table II: Gender Differences in the Ask Salary

Dep. Var.: Log Ask salary	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female	-0.068*** (0.003)	-0.044*** (0.002)	-0.047*** (0.002)	-0.046*** (0.002)	-0.029*** (0.002)	-0.032*** (0.003)	-0.024*** (0.003)
Experience		X	X	X	X	X	X
City		X	X	X	X	X	X
Occupation		X	X	X	X	X	X
Education			X	X	X	X	X
Work preferences				X	X	X	X
Employment history					X	X	X
Recent company FE						X	
Month \times Year FE	X	X	X	X	X	X	X
Adj R-squared	0.010	0.656	0.668	0.678	0.708	0.601	0.809
Nb. obs	113,777	113,777	113,777	113,777	113,777	63,916	463,860

Note: This table presents estimates of β_0 from Equations 1 and 2, progressively adding the controls. Column (1) controls for gender and time fixed effects at the Month \times Year level and corresponds to Equation 1. The following columns correspond to Equation 2, progressively adding controls. Column (2) adds experience, location, and the field of occupation. The experience controls are a dummied out categorical variable for the number of years of experience in the preferred occupation (0-2, 2-4, 4-6, 6-10, 10-15, 15+) and the number of years of total experience (linear and square term), and a dummied out categorical variable for the candidates' experience on the platform measured as the number of previous spells and the length of the current spell. The location controls are both the current and desired city of the candidate. The occupation control is a (dummied out) categorical variable (e.g. Software engineering). Column (3) adds education controls as described in Table B.1. Column (4) adds work preferences expressed by the candidate such as remote work and sponsorship needs, Columns (5), (6) and (7) add controls for employment history, namely a dummy for whether the candidate is currently employed, the number of days of unemployment, the number of people managed by the candidate in her current job (1-5, 5-10 etc.), a dummied out categorical variable for the highest job title of the candidate (e.g. "manager"), the number of people managed in current job, a dummy for whether the candidate has ever worked in one of the FAANG (Facebook, Amazon, Apple, Netflix, Google), and a dummy on whether the candidate has included a link to a personal website or LinkedIn page on their profile. Finally, dummies for the skills that the candidate has (e.g. HTML, Python etc.) are included. Column (6) adds fixed effects for the candidates' most recent company id. For candidates with multiple spells on the platform, we select their first ask in Columns (1) to (6). Robust standard errors are used in Columns (1) to (6). In Column (7) the ask gap is estimated at the bid level and standard errors are clustered at the candidate level. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table III: The Role of the Ask and Resume Characteristics in Bid Salary Gender Differences

Dep. Var.: Log Bid salary	No Job FE					Job FE		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female	-0.034*** (0.007)	-0.022*** (0.003)	0.002** (0.001)	-0.002*** (0.001)	-0.002*** (0.001)	-0.049*** (0.002)	-0.018*** (0.002)	-0.003*** (0.001)
Log Ask salary			0.963*** (0.002)	0.849*** (0.008)	0.848*** (0.008)			0.774*** (0.009)
Female \times Log Ask salary					0.001 (0.004)			
Constant	11.658*** (0.012)	19.735*** (0.522)	11.588*** (0.003)	13.081*** (0.139)	13.080*** (0.139)	11.557*** (0.006)	17.787*** (0.398)	12.983*** (0.120)
Candidate's resume characteristics		X		X	X		X	X
Month \times Year FE	X	X	X	X	X	X	X	X
Job FE						X	X	X
Adj R-squared	0.007	0.816	0.950	0.954	0.954	0.014	0.329	0.834
Nb. obs	463,860	463,860	463,860	463,860	463,860	454,631	454,631	454,631

Note: This table presents estimates of β_1 from equations 4 to 7. All regressions include time fixed effects at the Month \times Year level, and Columns (2), (4) and (5) add the controls in Column (5) of Table II as well as candidates preferences over firm characteristics as described in Table B.1. Column (1) estimates the raw gender bid gap (Equation 4). Coefficients in Column (2) correspond to Equation 5. Column (3) only controls for gender and the mean-centered log ask salary (Equation 6) and the fixed effect as in Column (1). Column (4) presents estimates following Equation 7. Column (5) adds an interaction between the Female dummy and the mean-centered log ask salary. Columns (6), (7), and (8) add job fixed effects to Columns (1), (2), and (4) respectively and singleton jobs are dropped. For these three columns, the R-squared is adjusted within. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table IV: Estimates for Controls other than Gender in Equations 2 and 7 and for Final Offers

Dep. Var.:	Log Ask salary	Log Bid salary		Log Final salary	
	(1)	(2)	(3)	(4)	(5)
Female	-0.029*** (0.002)	-0.022*** (0.003)	-0.002*** (0.001)	-0.014** (0.006)	0.010** (0.004)
Employed	0.069*** (0.002)	0.043*** (0.002)	0.003*** (0.001)	0.031*** (0.005)	0.007* (0.004)
Log Ask salary			0.848*** (0.008)		0.709*** (0.028)
Female \times Log Ask salary			0.001 (0.004)		0.011 (0.011)
Years of experience in the desired occupation					
2-4	0.106*** (0.002)	0.093*** (0.003)	0.011*** (0.001)	0.104*** (0.008)	0.018*** (0.006)
4-6	0.199*** (0.003)	0.174*** (0.004)	0.020*** (0.002)	0.188*** (0.009)	0.038*** (0.007)
6-10	0.299*** (0.003)	0.245*** (0.004)	0.027*** (0.002)	0.252*** (0.010)	0.045*** (0.009)
10-15	0.345*** (0.004)	0.275*** (0.005)	0.031*** (0.003)	0.281*** (0.014)	0.044*** (0.012)
15+	0.378*** (0.005)	0.291*** (0.006)	0.031*** (0.003)	0.294*** (0.017)	0.043*** (0.015)
Education					
Bachelor	0.053*** (0.011)	0.026** (0.013)	0.004* (0.002)	0.012 (0.038)	-0.005 (0.016)
Master	0.086*** (0.011)	0.039*** (0.013)	0.006** (0.002)	0.034 (0.038)	0.002 (0.016)
PhD	0.151*** (0.012)	0.081*** (0.013)	0.011*** (0.003)	0.075* (0.040)	0.018 (0.019)
University Ranking					
21-100	0.002 (0.003)	-0.001 (0.004)	0.001 (0.001)	-0.003 (0.009)	-0.002 (0.007)
101-500	-0.021*** (0.003)	-0.019*** (0.004)	0.000 (0.001)	-0.013 (0.010)	-0.003 (0.007)
501-1,000	-0.038*** (0.004)	-0.027*** (0.005)	-0.000 (0.001)	-0.011 (0.011)	-0.014* (0.007)
1,001-5,000	-0.047*** (0.003)	-0.029*** (0.004)	-0.001 (0.001)	-0.016* (0.010)	-0.000 (0.006)
5,000+	-0.057*** (0.003)	-0.037*** (0.004)	-0.003*** (0.001)	-0.027*** (0.010)	-0.008 (0.007)
Candidate's resume characteristics	X	X	X	X	X
Month \times Year FE	X	X	X	X	X
Adj R-squared	0.708	0.816	0.954	0.827	0.920
Nb. obs	113,777	463,860	463,860	7,582	7,582

Note: This table explores the role of controls other than gender in explaining ask, bid, and final salaries. Column (1) follows Equation 2, and Column (2) Equation 5. Column (3) follows Equation 7 adding an additional interaction term between the Female dummy and the mean-centered log ask salary. Columns (4) and (5) use the same controls as Columns (2) and (3) respectively, but setting log final salary as the dependent variable. The omitted category for “Years of experience” is 0-2, for “Education” it is High School, and for “University Ranking” it is 1-20. In Column (1) standard errors are robust and in Columns (2) to (5) standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table V: The Role of the Ask and Resume Characteristics in Final Offer Gender Differences

Dep. Var.: Log Final salary	No Firm FE					Firm FE		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female	-0.049*** (0.014)	-0.014** (0.006)	0.023*** (0.004)	0.009** (0.004)	0.010** (0.004)	-0.018*** (0.006)	0.002 (0.004)	0.003 (0.004)
Log Ask salary			0.956*** (0.007)	0.712*** (0.026)	0.709*** (0.028)		0.617*** (0.026)	0.615*** (0.028)
Female \times Log Ask salary					0.011 (0.011)			0.008 (0.013)
Candidate's resume characteristics		X		X	X	X	X	X
Month \times Year FE	X	X	X	X	X	X	X	X
Firm FE						X	X	X
Adj R-squared	0.012	0.827	0.903	0.920	0.920	0.515	0.762	0.762
Nb. obs	7,582	7,582	7,582	7,582	7,582	6,303	6,303	6,303

Note: This table presents estimates of β_1 from Equations 4 to 7, except the left-hand side is $\text{Log}(\text{Final}_{ib})$ - the final salary that candidate i was offered for the job corresponding to bid b - instead of $\text{Log}(\text{Bid}_{ib})$. Accordingly, in Columns (1) to (5), controls are the same as the corresponding columns in Table III. The regressions are run on the sample of final offers. Columns (6), (7), and (8) add firm fixed effects to Columns (2), (4), and (5) respectively, and singleton firms are dropped. For these three columns, the R-squared is adjusted within. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table VI: Gender Differences in the Number of Bids Received and the Probability of Receiving a Final Offer after an Interview

Dep. Var.:	Nb. Bids Received					Final Offer Received		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female	-0.397*** (0.035)	0.227*** (0.032)	0.260*** (0.032)	0.271*** (0.032)	0.326*** (0.094)	0.001 (0.001)	0.001 (0.001)	-0.000 (0.001)
Ask salary			0.943*** (0.062)	1.926*** (0.068)	0.982*** (0.058)		-0.000 (0.003)	0.023*** (0.003)
Ask salary²				-0.228*** (0.014)			0.001 (0.001)	-0.002*** (0.001)
Female \times Ask salary					-0.059 (0.093)			
Constant	3.977*** (0.099)	-44.076*** (4.890)	-52.845*** (4.881)	-56.990*** (4.868)	-52.889*** (4.881)	-1.349*** (0.212)	-1.370*** (0.212)	-1.572*** (0.220)
Poisson / Logit AME on Female	-0.402	0.303	0.331	0.361	0.329	0.000	0.000	-0.018
Candidate's resume characteristics		X	X	X	X	X	X	X
Month \times Year FE	X	X	X	X	X	X	X	X
Job FE								X
Adjusted R-squared	0.015	0.240	0.244	0.245	0.244	0.008	0.008	0.038
Nb. obs	164,799	164,799	164,799	164,799	164,799	261,518	261,518	251,817

Note: This table assesses whether there are gender differences in the number of bids received during a candidate's spell on the platform and in the probability of getting a final offer after an interview. In the first five columns, it is important to note that regressions are at the spell level. Indeed, there are 113,777 candidates and, since some of them are on multiple spells, that sums up to 164,799 spells. Column (1) only controls for gender and time fixed effects at the Month \times Year level, Column (2) adds resume characteristics as controls (as in Table III Column (2)). Column (3) adds the ask salary in 100,000 USD to Column (2). Column (4) adds the square of the ask salary in 100,000 USD to Column (3). Column (5) adds an interaction between the Female dummy and the ask salary to Column (3). Standard errors (in parentheses) are clustered at the candidate id level. In the last three columns, each observation is one bid but the sample is restricted to bids that let to an interview. In other words, bids that were rejected by the candidate, so that there was no interview, are not in the sample. Columns (6) and (7) have the same controls as columns (2) and (4). Column (8) adds job fixed effects to Column (7). Standard errors (in parentheses) are clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table VII: The Effect of the Reform on the Gender Gap

	Log Ask salary		Log Bid salary			Log Final salary
	(1)	(2)	(3)	(4)	(5)	(6)
Female \times After	0.035*** (0.006)	0.030*** (0.009)	0.028*** (0.011)	0.002 (0.007)	0.022*** (0.009)	0.019 (0.035)
Female	-0.029*** (0.003)	-0.025*** (0.005)	-0.025*** (0.004)	-0.004*** (0.001)	-0.018*** (0.003)	-0.018 (0.012)
After	-0.003 (0.004)	-0.005 (0.006)	-0.002 (0.007)	0.002 (0.003)	0.001 (0.006)	-0.002 (0.019)
Log Ask Salary				0.816*** (0.008)		
Mean Dep. Var.	11.78	11.87	11.86	11.86	11.86	11.86
Candidate's resume characteristics	X	X	X	X	X	X
Job FE					X	
Adj R-squared	0.517	0.493	0.514	0.865	0.252	0.519
Nb. obs	43,368	207,636	207,636	207,636	200,593	2,476

Note: This table presents estimates of β_0 , β_1 and β_2 from Equations 8 and 9 as well the corresponding coefficients for the log final salary controlling for all candidate's resume characteristics as well as month dummies (1-12) for seasonal adjustments and a linear time trend. Columns (1) and (2) provide estimates for Equation 8. Column (1) uses the observations at the candidate level, while Column (2) shows the results bid-weighted. Columns (3) to (5) provide estimates for Equation 9, estimated at the bid level. Column (4) adds the log ask salary as a control to Column (3), while Column (5) adds job identifier fixed effects to Column (3). Column (6) provides estimates analogous to Column (3) but for log final salaries. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table VIII: The Effect of the Reform at the Extensive Margin

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Nb bids	Nb hours	Final offer	Firm rank (bid)	Firm rank (final)	Nb bids	Pred-Nb bids	Pred-Nb bids
Female \times After	0.190 (0.183)	0.741 (4.459)	0.006 (0.009)	0.309 (0.311)	0.018 (0.922)	0.037 (0.179)	0.034 (0.109)	0.178 (0.112)
Female	0.440*** (0.112)	1.115 (1.796)	-0.003 (0.005)	0.335 (0.251)	0.035 (0.812)	0.569*** (0.110)	-0.191*** (0.066)	-0.191*** (0.069)
After	-0.215* (0.124)	-4.501* (2.647)	-0.008 (0.007)	0.290 (0.220)	0.310 (0.710)	-0.252** (0.122)	0.149*** (0.049)	0.174*** (0.050)
Ask salary						0.047*** (0.003)		
Ask salary²						-0.000*** (0.000)		
Poisson AME on Female \times After	0.222	-1.738	0.009					
Mean Dep. Var.	4.79	62.33	0.09	62.5	62.9			
Candidate's resume characteristics	X	X	X	X	X	X		
Adj R-squared	0.227	0.109	0.033	0.242	0.032	0.092	0.002	0.002
Nb. obs	43,368	32,043	43,368	188,463	2,074	43,368	43,368	43,368

Note: This table estimates the effect of the reform at the extensive margin. Column (1) provides estimates for the regression of the number of bids on the Female dummy, the After dummy and their interaction, with controls for the candidate's resume characteristics. Column (2) provides estimates for the regression of the number of hours before a candidate's first bid and Column (3) presents estimates for the regression whether a final offer was provided on the same variables. The dependent variable in Column (4) is the Firm rank at the bid level as in Table G.1 Column (1) and the dependent variable in Column (5) is the Firm rank at the final offer level as in Table G.1 Column (3). Column (6) adds the ask salary and ask salary squared to Column (1). The dependent variable in Column (7) is the predicted number of bids received using Column (1) specification on the pre-period. The dependent variable in Column (8) is the predicted number of bids received using Column (6) specification on the pre-period. Observations are at the spell level. Standard errors (in parentheses) are clustered at the candidate id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Online Appendix

THE ROLE OF THE ASK GAP IN GENDER PAY INEQUALITY

Nina Roussille

This online appendix has eight sections. Appendix [A](#) includes extra figures, Appendix [B](#) includes extra tables. In Appendix [C](#), I explore the degree of heterogeneity in the ask gap. In Appendix [D](#), I discuss the external validity of both the ask and the bid gap. In Appendix [E](#), I investigate racial differences in the ask, bid, and final salaries. In Appendix [F](#), I focus on the subset of candidates updating their ask salaries during a spell, and describe the relationship between the ask and bid salaries for this specific cases. In Appendix [G](#), I explore the relationship between candidate characteristics and the quality of the bids they receive, and, in Appendix [H](#), I provide a framework formalizing the intuition behind interpreting the ask salary as a signal of quality. Finally, Appendix [I](#) outlines the main steps of the dataset construction.

A Appendix Figures

Figure A.1: Mandatory Features of a Candidate Profile, at the Time of the Study

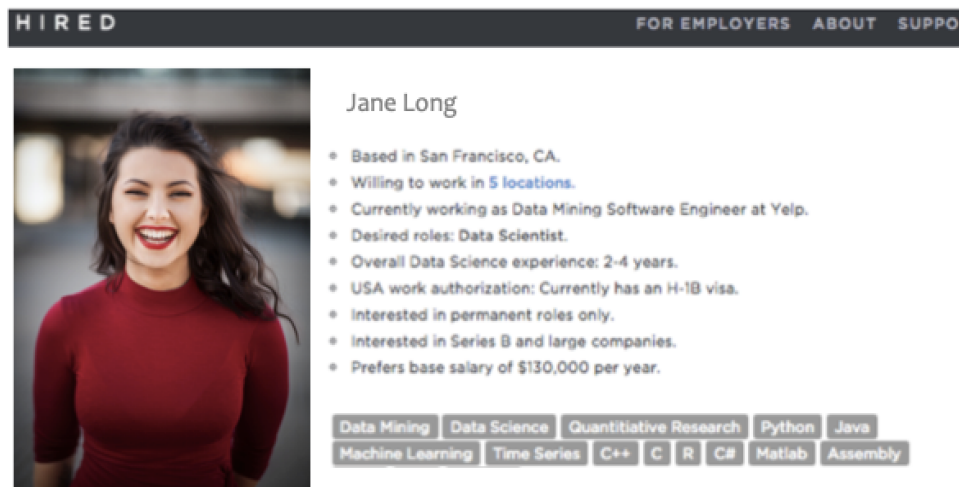


Figure A.2: Typical Interview Request Message Sent by a Company to a Candidate, at the Time of the Study

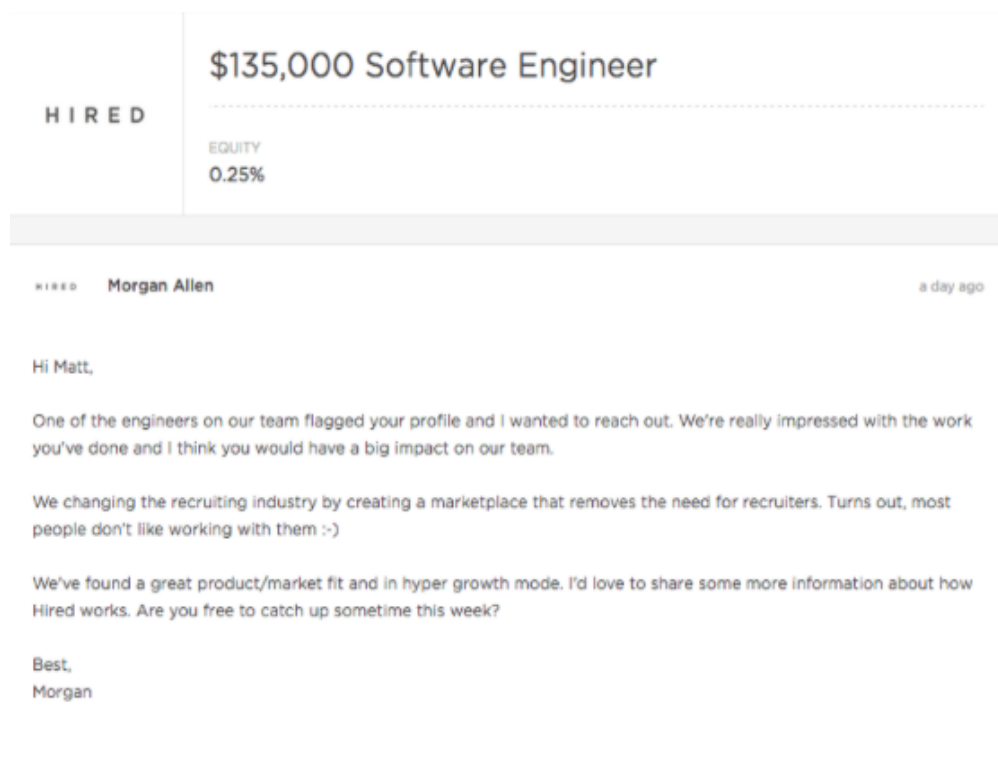
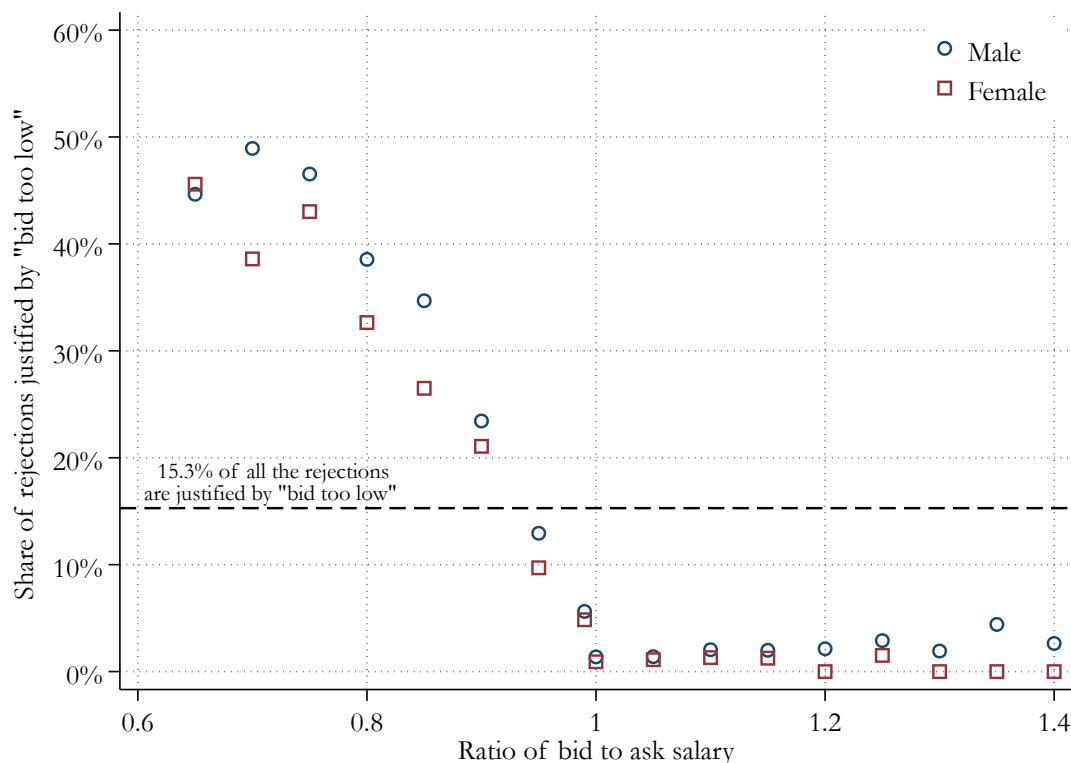
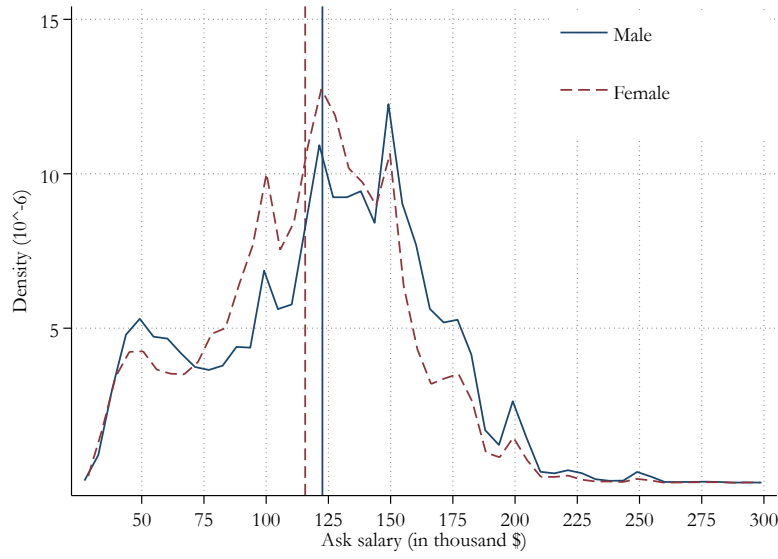


Figure A.3: Interview Request Rejection Reason as a Function of the Bid to Ask Ratio

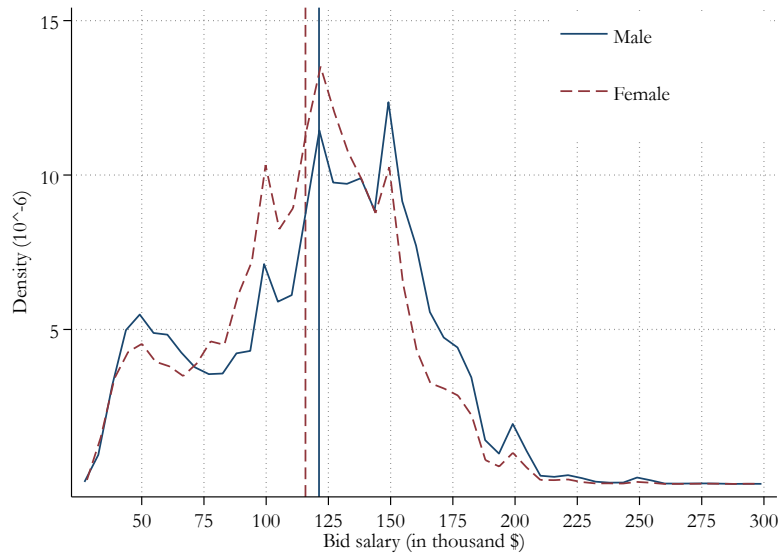


Note: This figure shows that the share of rejected interview requests - stating as the reason the bid was too low - decreases with the ratio of bid to ask salary, separately for male (blue circles) and female (red squares) candidates. When a candidate receives a bid, he can decide to reject it, that is he can refuse to interview with the company. They can choose from justifications such as “company culture”, “company size”, or “insufficient compensation”. The latter is the justification labeled as “bid too low”.

Figure A.4: Kernel Density of Ask and Bid Salaries



(a) Kernel density of ask salaries



(b) Kernel density of bid salaries

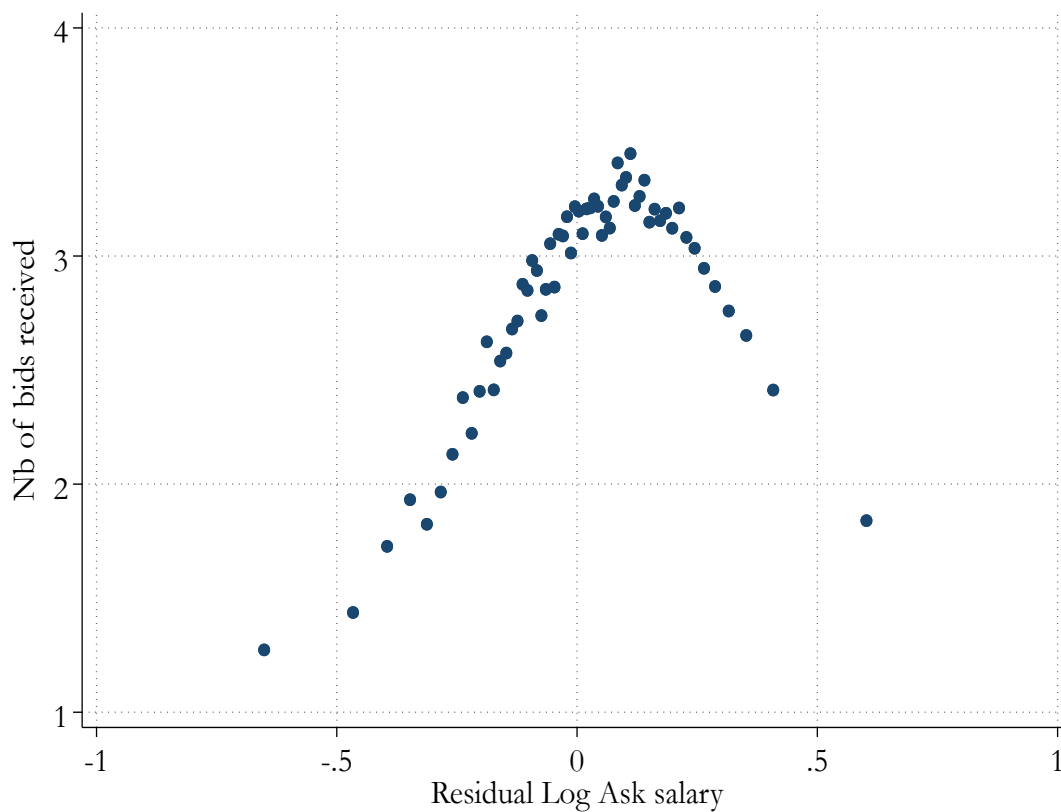
Note: These figures plot kernel density estimates of the distributions of ask salaries in Figure A.4 Panel (a) and bid salaries in Figure A.4 Panel (b), separately for male (solid blue line) and female (dashed red line) candidates. Vertical lines indicate the mean salary respectively for male and female. The kernel density estimates for bid salaries include all 463,860 bids and the kernel density estimates for ask salaries include all 113,777 ask salaries, weighted by the number of bids received by the candidate.

Figure A.5: Ask Feature Change on the Platform



Note: This figure shows the effect of the reform on the candidate's ask salary elicitation when they create their profile. In the top figure is the question design before the reform: the answer box is empty. In the bottom figure is the question design after the reform: the answer box is pre-filled with the median bid salary corresponding to the candidate's profile (here a software engineer in San Francisco with similar experience) and calculated over the last 12 months. Candidates were informed that the pre-filled number was the median salary corresponding to their location, job title, and experience.

Figure A.6: Binned Scatter Plot of the Number of Bids Received as a Function of the Residual Log Ask Salary



Note: This figure shows the relationship between the number of bids received by a candidate during a spell on the platform and the log ask salary of this candidate, residualized on all the resume characteristics on the candidate's profile. The underlying data contains the individual spells of all candidates.

B Appendix Tables

Table B.1: Fields on a Candidate’s Profile and Other Variables Used as Controls

Resume characteristics	Type of variable	Controls in the regression
Fields on a candidate’s profile		
What type of position do you currently have? (job title)	Categorical variable - drop down menu - single entry - mandatory	• Software Engineering • Engineering management • Design • Data Analytics • Developer Operations • Quality Assurance • Information Technology • Project management • Product management
Total Position experience (in years)	Categorical variable - drop down menu - single entry - mandatory	• 0-2 years • 2-4 years • 4-6 years • 6-10 years • 10-15 years • 15+ years
Skills : Rank your top 5 languages & skills	Categorical variable - drop down menu - multiple (up to 5 entries, at least 1)	Choice from many categories, the most cited (>10% of the time) are: • javascript • python • sql • c • nodejs • ruby • css • react. All CS skills that are cited by more than 0.05% of the sample are included as dummies in the regression. ²⁸
Where do you live?	Categorical variable - drop down menu - single entry - mandatory	• San Francisco • Los Angeles • San Diego • Seattle • Denver • Austin • Houston • Chicago • Boston • Washington D.C. • New York
Where do you want to work?	Categorical variable - drop down menu - multiple entry - mandatory	• San Francisco • Los Angeles • San Diego • Seattle • Denver • Austin • Houston • Chicago • Boston • Washington D.C. • New York
Are you interested in working remotely?	Categorical variable - drop down menu - single entry - mandatory	• Yes • No • Remote Only
What type of employment are you seeking?	Categorical variable - drop down menu - single entry - mandatory	• Full Time Only • Prefers Full Time • Full Time Only • Both equally • Prefers Contract • Contract Only
Preferred company size: I’d like to work at a company that has __ employees	Categorical variable - drop down menu - multiple entries - optional	Dummies on selected size: • 1-15 • 16-50 • 51-200 • 201-500 • 500+
Preferred industry: My ideal company would be in these industries:	Categorical variable - drop down menu - multiple entries - optional	Top ten most chosen industries: • bank, corporate finance, & investing • analytics & business information • e-commerce • health care technology & nursing • hardware, internet of things, & electronics • information systems • education • digital payments • social networking • digital communication
Preferred career path:	Categorical variable - drop down menu - multiple entries - optional	Dummies on selected path: • contributinal role • manager
Preferred career goal:	Categorical variable - drop down menu - multiple entries - optional	Dummies on selected goal: • leadership • great culture • mentorship • new technologies • socially conscious • large projects
Preferred skill to use on job:	Categorical variable - drop down menu - multiple entries - optional	Dummies on selected skill: top 5 most selected skills: Python, Java, JavaScript, Product Management, React. The top 30 skills cover 80% of all listed skills and are included as dummies in the regression.

²⁸ The full set of included dummies is: html, java, python, javascript, ios, pointnet, android, sql, c, ruby, dataanalysis, php, nodejs, css, react, go, r, saas, linux, agile, angular, swift, hadoop, scala

Where are you in your job search?	Categorical variable - drop down menu - single entry - mandatory	<ul style="list-style-type: none"> • not looking for new opportunities / just browsing • open to exploring new opportunities • actively looking for new opportunities • currently interviewing • have offers
Will you now or in the future require sponsorship for employment visa status (e.g. H-1B Visa)?	categorical variable - drop down menu - single entry - optional	<ul style="list-style-type: none"> • Sponsorship Required • Not Required
Firm history	Manual entry of the history of firms that the candidate worked at and when - optional	Here I built a dummy = 1 if the candidate has ever worked at an “elite” tech company (FAANG - Facebook, Amazon, Netflix, Google). I also identify the most recent company the candidate is working at for the most recent firm FE specification. I finally built two other variables: the average tenure at a job and the total number of months of work interruption (both enter the regression linearly and squared).
Job titles	Manual entry of the job title held at each firm the candidate worked at - optional	I created a categorical variable for the highest position held in a firm (“junior”, “senior”, “manager”, “lead”, “head”, “director”) as well as whether the candidate ever founded a company.
Number of people managed in current job	Categorical variable - drop down menu - single entry - optional	<ul style="list-style-type: none"> • 1-5 • 6-10 • 11-20 • 20+
Education	Manual entry of educational institution, degree and year - optional	Here I built 5 variables: categorical for highest degree achieved (high school, Associate, Bachelor, Master, MBA, PhD), the average global ranking as reported on Webometrics of all schools the candidate attended grouped into six categories (1-20, 21-100, 101-500, 501-1,000, 1,001-5,000, and 5,000+), a dummy for whether the candidate ever attended an IvyLeague+ school (as defined in Chetty et al. (2020), to which I added schools that are ranked in the top 5 programs in engineering by the annual U.S. News college ranking), a dummy for whether the degree is in computer science and the last graduation year.
External Websites	Links to external LinkedIn page or personal website - optional	Here I built two dummies: one dummy when a candidate has included at LinkedIn profile and one dummy when the candidate links to a personal website.

Other control variables

Equity	Optional	Dummy for whether equity of the company was included in the bid to the candidate.
Total experience	-	Number of years of experience, enters linearly and squared in the regression.
Number of jobs held	-	Number of jobs held, enters linearly and squared in the regression
Employed	Dummy variable	<ul style="list-style-type: none"> • Yes • No
Number of days searching for work	-	Number of days searching for work (enters the regression linearly).
Number of past spells on the platform	Categorical variable	<ul style="list-style-type: none"> • 1 • 2 • 3 • 4+
Length of spell on the platform	Categorical variable	Number of days the profile is live on the platform (15 - 22 - 29 - 36 - 43)

Table B.2: Relationship Between Gender and Expressed Preferences over Firm Characteristics

Dep. Var.:	No Preference	Pref. Company Size					Pref. Industry				Pref. Career Goal			
		1-15	16-50	51-200	201-500	500+	Hardware, IoT	Finance	Education	Health-Tech	New Technologies	Leadership	Mentorship	Socially Conscious
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Female	0.014*** (0.003)	-0.005 (0.003)	-0.030*** (0.003)	0.004 (0.003)	0.009*** (0.003)	0.020*** (0.003)	-0.011*** (0.001)	-0.007*** (0.001)	0.005*** (0.001)	0.007*** (0.001)	-0.013*** (0.003)	-0.005** (0.003)	0.007*** (0.002)	0.023*** (0.002)
Male mean	0.252	0.469	0.432	0.453	0.433	0.351	0.033	0.041	0.026	0.028	0.249	0.189	0.090	0.088
Candidate's resume characteristics	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Month \times Year FE	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Adj R-squared	0.211	0.404	0.362	0.383	0.366	0.289	0.126	0.148	0.132	0.122	0.348	0.275	0.157	0.149
Nb. obs	113,777	113,777	113,777	113,777	113,777	113,777	113,777	113,777	113,777	113,777	113,777	113,777	113,777	113,777

Note: This table presents estimates of a subset of candidates' expressed preferences over company size in Columns (2) to (6), industry in Columns (7) to (10), and career goals in Columns (11) to (14) on gender. Column (1) presents estimates for whether any preference was listed on Hired.com. Standard errors (in parentheses) are robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.3: The Last Ask Salary as a Function of Gender and Resume Characteristics

Dep. Var.: Log Ask salary	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female	-0.070*** (0.003)	-0.044*** (0.002)	-0.047*** (0.002)	-0.045*** (0.002)	-0.028*** (0.002)	-0.030*** (0.003)	-0.024*** (0.003)
Experience		X	X	X	X	X	X
City		X	X	X	X	X	X
Occupation		X	X	X	X	X	X
Education			X	X	X	X	X
Work preferences				X	X	X	X
Employment history					X	X	X
Recent company FE						X	
Month \times Year FE	X	X	X	X	X	X	X
Adj R-squared	0.009	0.657	0.669	0.680	0.712	0.609	0.809
Nb. obs	113,777	113,777	113,777	113,777	113,777	63,916	463,860

Note: This table presents estimates of β_0 from Equation 2, progressively adding controls. In contrast to Table II, the candidates' last ask is used here. Column (1) controls for gender and time fixed effects at the Month \times Year level. Column (2) adds experience, location, and the field of occupation. The experience controls are a dummied out categorical variable for the number of years of experience in the preferred occupation (0-2, 2-4, 4-6, 6-10, 10-15, 15+) and the number of years of total experience (linear and square term) and a dummied out categorical variable for the candidates' experience on the platform measured in the number of previous spells and length of the current spell. The location controls are both the current and desired city of the candidate. The occupation control is a (dummied out) categorical variable (e.g. Software Engineering). Column (3) adds education controls as described in Table B.1. Column (4) adds work preferences expressed by the candidate such as remote work and sponsorship needs, Columns (5), (6) and (7) add controls for employment history, namely a dummy for whether the candidate is currently employed, the number of days of unemployment, the number of people who report to the candidate in her current job (1-5, 5-10 etc.), a dummied out categorical variable for the highest job title of the candidate (e.g. Vice President), a dummy for whether the candidate has ever worked in one of the FAANG (Facebook, Amazon, Apple, Netflix, Google) and a respective dummy whether the candidate has included a link to a personal website or LinkedIn page on the profile. Finally, I add dummies for the skills that the candidate has (e.g. HTML, Python etc.). Column (6) controls for fixed effects of candidates' most recent company. Singleton firms are dropped and the R-squared is adjusted within. Robust standard errors are used in Columns (1) to (6). In Column (7) the ask gap is estimated on the bid level and standard errors are clustered at the candidate level. Standard errors in parentheses.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.4: Ask Gap Corrected for Unobservables following Altonji, Elder, and Taber (2005)

Treatment Variable:	Baseline Effect			Controlled Effect			Identified Set	R max	
	Coefficient	(Std. Error)	[R-sqrd]	Coefficient	(Std. Error)	[R-sqrd]	for $\beta = 0$, $\tilde{\delta} = 1$		
Female	-0.062	0.003	0.003	-0.028	0.002	0.712	-0.028	-0.011	1

Note: This table presents the results of the selection exercise on observable and unobservable variables as proposed by Altonji, Elder, and Taber (2005). $\beta = 0$ represents the null of a zero gender ask gap. As the identified set for the coefficient on the ask gap does not include 0, the null can be rejected.

Table B.5: The Role of the Ask Salary and Resume Characteristics in Bid Salary Gender Differences Including Equity

Dep. Var.: Log Bid salary	(1)	(2)	(3)	(4)	(5)
Female	-0.037*** (0.006)	-0.022*** (0.003)	0.002** (0.001)	-0.002*** (0.001)	-0.002*** (0.001)
Log Ask salary			0.964*** (0.002)	0.850*** (0.008)	0.849*** (0.008)
Female \times Log Ask salary					0.001 (0.004)
Constant	11.555*** (0.013)	19.738*** (0.521)	11.593*** (0.003)	13.071*** (0.139)	13.071*** (0.139)
Candidate's resume characteristics		X		X	X
Month \times Year FE	X	X	X	X	X
Adj R-squared	0.028	0.816	0.950	0.954	0.954
Nb. obs	463,860	463,860	463,860	463,860	463,860

Note: This table adds equity as a control to the estimates of β_1 from Equations 4 to 7 in Table III. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.6: The Role of the Ask Salary and Resume Characteristics in Bid Salary Gender Differences for a Given Firm

Dep. Var.: Log Bid salary	(1)	(2)	(3)	(4)	(5)
Female	-0.076*** (0.003)	-0.022*** (0.002)	-0.007*** (0.001)	-0.003*** (0.001)	-0.003*** (0.001)
Log Ask salary			0.859*** (0.006)	0.809*** (0.009)	0.808*** (0.009)
Female \times Log Ask salary					0.004 (0.004)
Candidate's resume characteristics		X		X	X
Month \times Year FE		X		X	X
Firm FE	X	X	X	X	X
Adj R-squared	0.028	0.472	0.869	0.874	0.874
Nb. obs	463,446	463,446	463,446	463,446	463,446

Note: This table adds firm fixed effects to the estimates of β_1 from Equations 4 to 7 in Table III. Singleton firms are dropped and the R-squared is adjusted within. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.7: The Role of the Ask Salary and Resume Characteristics in Bid Salary Gender Differences - Sample Restriction: Only Keep Bids for Jobs that Lead to a Hire on the Platform

Dep. Var.: Log Bid salary	(1)	(2)	(3)	(4)	(5)
Female	-0.035*** (0.008)	-0.022*** (0.003)	0.002** (0.001)	-0.002*** (0.001)	-0.002*** (0.001)
Log Ask salary			0.960*** (0.003)	0.842*** (0.010)	0.841*** (0.011)
Female \times Log Ask salary					0.002 (0.005)
Constant	11.686*** (0.016)	19.651*** (0.590)	11.595*** (0.004)	13.168*** (0.204)	13.169*** (0.204)
Candidate's resume characteristics		X		X	X
Month \times Year FE	X	X	X	X	X
Adj R-squared	0.004	0.809	0.949	0.953	0.953
Nb. obs	201,589	201,589	201,589	201,589	201,589

Note: This table presents estimates of the gender bid gap on the subset of bids for jobs that lead to a hire on the platform. For all columns, controls are the same as in Table III. The only difference is the sample since here we only keep bids for jobs that lead to a hire on the platform. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.8: The Role of the Ask Salary and Resume Characteristics in Bid Salary Gender Differences - Sample Restriction: Only Keep Bids that are Different from the Candidate's Ask

Dep. Var.: Log Bid salary	(1)	(2)	(3)	(4)	(5)
Female	-0.040*** (0.008)	-0.016*** (0.003)	0.007** (0.003)	-0.003 (0.002)	-0.004 (0.002)
Log Ask salary			0.856*** (0.007)	0.542*** (0.016)	0.540*** (0.017)
Female \times Log Ask salary					0.005 (0.011)
Constant	11.641*** (0.015)	20.284*** (0.784)	11.562*** (0.006)	15.692*** (0.518)	15.697*** (0.519)
Candidate's resume characteristics		X		X	X
Month \times Year FE	X	X	X	X	X
Adj R-squared	0.006	0.772	0.806	0.850	0.850
Nb. obs	105,144	105,144	105,144	105,144	105,144

Note: This table presents estimates of the gender bid gap on the subset of bids that are different from the candidate's ask salary. For all columns, controls are the same as in Table III. The only difference is the sample since here we only keep bids that are different from the candidate's ask. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.9: Predicted Ask Gap using a Model Fitted on the Pre-Reform Sample

	(1)
Female \times After	-0.003 (0.006)
After	0.002 (0.003)
Female	-0.080*** (0.003)
Adj R-squared	0.02
Nb. obs	43,368

Note: This table tests the stability over time of the predicted ask gap using a model fitted on pre-reform data. The sample is all San Francisco software engineers in the dataset. The predicted log ask salary (dependent variable) is obtained fitting Equation 2 on the pre-reform sample of SF software engineers, except that instead of Month \times Year FE, there are just Month FE (1-12) and a monthly linear time trend. Standard errors (in parentheses) are robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B.10: Summary Statistics on Candidates before and after the Reform

Variable	Female - After	Female - Before	Male - After	Male - Before
Nb. of bids	7,918	23,360	35,222	108,231
Nb. of candidates	1,699	5,028	7,242	22,376
Years of experience	9.0	9.9	11.0	11.4
Share with a bachelor	99.4	99.7	98.8	98.8
Share with a master	63.6	58.8	55.6	51.4
Share with a CS degree	69.9	71.6	68.5	71.2
Share with an IvyPlus degree	13.1	14.4	12.3	13.6
Share looking for full time job	98.8	99.0	98.3	97.8
Share in need of visa sponsorship	33.2	30.8	31.0	29.6
Share of remote only workers	0.4	0.3	2.5	1.9
Share employed	70.2	66.3	73.1	72.2
Share that worked at a FAANG	8.1	8.4	8.7	9.6
Share leading a team	18.1	25.0	24.8	29.2

Note: This table provides summary statistics on the number of bids and candidates as well as candidates' resume characteristics before and after the change, illustrating the absence of differential selection of men and women onto the platform after the reform. The "CS" in "CS degree" stands for computer science, "IvyPlus degrees", as defined by Chetty et al. (2020), including the eight Ivy League institutions + U. Chicago, Stanford, MIT, and Duke. Additionally, I include the schools with the top five highest-ranked programs in engineering on the annual U.S. News college ranking: UC Berkeley, California Institute of Technology, Carnegie Mellon University, and Georgia Institute of Technology. FAANG is a dummy for whether the candidate has ever worked in one of Facebook, Amazon, Apple, Netflix, or Google. Finally, note that summing the number of female and male candidates gives 36,345 observations, the remaining 7,023 observations in Table VII Column (1) and Table B.9 are candidates whose gender has not been identified.

Table B.11: Impact of the Reform on Controls other than Gender in the Ask Gap Estimation

	(1)	(2)
Dep. Var.: Log Ask Salary	Before	After
Female	-0.027*** (0.003)	0.004 (0.005)
Employed	0.061*** (0.003)	0.042*** (0.005)
Years of experience in the desired occupation		
2-4	0.091*** (0.004)	0.111*** (0.007)
4-6	0.174*** (0.004)	0.220*** (0.009)
6-10	0.264*** (0.005)	0.294*** (0.010)
10-15	0.308*** (0.007)	0.396*** (0.012)
15+	0.334*** (0.008)	0.446*** (0.014)
Education		
Bachelor	0.060* (0.034)	0.038 (0.055)
Master	0.078** (0.034)	0.036 (0.055)
PhD	0.116*** (0.034)	0.095* (0.056)
University Ranking		
21-100	0.011** (0.005)	0.001 (0.009)
101-500	-0.018*** (0.005)	-0.002 (0.010)
501-1,000	-0.022*** (0.005)	-0.007 (0.010)
1,001-5,000	-0.034*** (0.005)	-0.019** (0.009)
5,000+	-0.051*** (0.005)	-0.026*** (0.010)
Constant	16.292*** (0.715)	15.851*** (1.071)
Candidate's resume characteristics	X	X
Month \times Year FE	X	X
Adj R-squared	0.535	0.475
Nb. obs	32,649	10,719

Note: This table presents coefficients on a regression running the log ask salary on all resume characteristics controls, separately for the pre-reform period in Column (1) and the post-reform period in Column (2). Standard errors (in parentheses) are robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

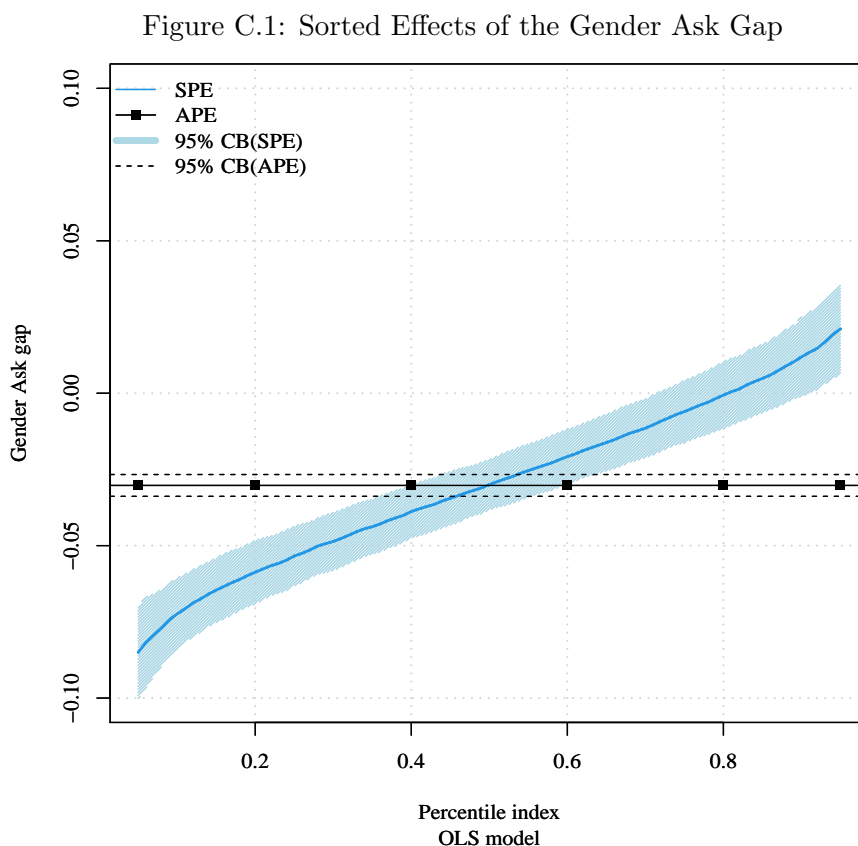
Table B.12: The Ask Gap by Share of Women in the Labor Market

Dep. Var.: Ask Gap	(1)	(2)
Share of Women	-0.087*	-0.072*
	(0.039)	(0.028)
Constant	0.102***	0.034***
	(0.004)	(0.003)
Candidate's resume characteristics		X
Adj R-squared	0.034	0.040
Nb. obs	203	203

Note: This table presents estimates for the relationship between the ask gap in a given labor market (location \times job title) and the share of female job seekers in this role. The gender ask gap is estimated on the location \times job title and then normalized as the difference between the male and female coefficient for the respective role. Column (1) reports the estimate for the raw ask gap only controlling for month \times year fixed effects, while Column (2) further controls for the candidate's resume characteristics. Standard errors in parentheses are robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

C Heterogeneity in the ask gap

To explore the degree of heterogeneity in the ask gap with respect to the underlying resume characteristics, I estimate a model that fully interacts the female dummy with the resume characteristics of the candidates. Following Chernozhukov, Fernàndez, and Luo (2018), I summarize these results using the sorted effect method for interactive linear models. This method reports the percentiles of the partial effects in addition to the average effect.²⁹ Figure C.1 plots the estimates and 95% confidence sets of the population average partial effect (APE) and sorted partial effect (SPE) for



Note: This figure shows the degree of heterogeneity in the ask gap by reporting percentiles of the sorted partial effects (SPE), in addition to the average partial effect (APE), from a regression model where the female dummy is fully interacted with the resume characteristics. The method is described in Chernozhukov, Fernàndez, and Luo (2018) and I used the corresponding `spe` package on R (Chen et al. (2019)) to implement the sorted method and graph this plot. 95% bootstrap uniform confidence bands (see derivation in paper) are shaded in blue.

²⁹ This method presents several advantages. First, the sorted partial effects are derived from an interactive model where the selection of controls is semi-automated (i.e. the researcher's degrees of freedom in the selection of controls are reduced), such that the estimation is relatively immune to suspicions of p-hacking. Second, the sorted effect curve is an efficient way to represent the range of heterogeneous effects as it allows to visualize them in one single plot. Last, the classification analysis allows to identify the characteristics of the most or least affected groups.

the ask gap. The estimates range from 8.5% to -2.1%. In other words, there exists a subgroup of female candidates for whom the ask gap is close to 3 times as large as the APE, and there is a subgroup of women who actually ask for higher salaries than similar men. Table C.1 reports the results of the classification analysis, comparing the resume characteristics of people with the highest and lowest SPEs. In line with previous findings on the gender pay gap over the lifecycle (Goldin et al. (2017)), I find that the group with the largest ask gap (8.5%) is more experienced (13 years vs. 7 years of total experience). I also find that they are more likely to be unemployed, with longer unemployment spells, less likely to have a computer science or an IvyPlus degree, and less likely to list highly-demanded coding skills.

Table C.1: Classification Analysis - Averages of Characteristics of the Women with the Smallest and Largest Ask Gap

	5% Smallest ask gap	SE	5% Highest ask gap	SE
Total years of experience	7.31	0.42	12.85	0.51
Position experience = 2-4 years	0.38	0.05	0.12	0.02
Position experience = 4-6 years	0.14	0.03	0.19	0.03
Position experience = 6-10 years	0.08	0.02	0.41	0.04
Position experience = 10-15 years	0.01	0.01	0.13	0.03
Position experience = 15+ years	0.01	0.01	0.10	0.03
Employed	0.73	0.04	0.63	0.04
Days unemployed	48.11	11.96	247.31	50.04
Ivy League Plus	0.20	0.04	0.07	0.02
CS degree	0.55	0.05	0.41	0.04
Java	0.25	0.04	0.18	0.03
HTML	0.17	0.03	0.10	0.02
Python	0.28	0.04	0.11	0.02
JavaScript	0.34	0.04	0.11	0.02
SQL	0.18	0.03	0.27	0.04
Data analysis	0.13	0.03	0.08	0.02
pointnet	0.05	0.02	0.01	0.01
C	0.16	0.04	0.03	0.01
Node JS	0.08	0.02	0.05	0.01
CSS	0.16	0.03	0.08	0.02
React	0.17	0.04	0.01	0.01

Note: This table presents partial effects estimated from a linear model with interactions between the female dummy and all other resume characteristics. The classification analysis is performed using the Chernozhukov, Fernández, and Luo (2018) procedure. The procedure is implemented on the sample of all candidates' first ask salary in a spell on the platform.

D External validity

D.1 External validity of the ask gap

While there is no direct evidence on the ask gap in other datasets, the 6.8% raw ask gap I observe is comparable to the raw gender pay gap among computer engineers, who comprise much of the Hired sample. Specifically, the gender pay gap in computer engineering calculated using U.S. Census Bureau’s 2016 American Community Survey is 8%.

Further, to benchmark the adjusted ask gap estimate of 2.9%, it is useful to compare the ask salary to related concepts like survey expectations or reservation wages. The 2.9% estimated gap is on the lower end if I compare it to studies based on survey data. For instance, Krueger and Mueller (2016) found an 8.3% reservation wage gap in their survey of unemployed workers in New Jersey. However, recent papers using large administrative datasets have found similar estimates for closely related gender gaps. For instance, Le Barbanchon, Rathelot, and Roulet (2021) found a 3.6% residual gender reservation wage gap in France, using administrative data from unemployment insurance claimants. Fluchtmann et al. (2021), using data from Danish UI recipients, showed that after conditioning on a rich set of observables, women apply to jobs with wages that are on average 1.9 percent lower than men. Similar to my analysis, both papers employ a large and reliable set of observations and controls. This likely explains why Le Barbanchon, Rathelot, and Roulet (2021) find an R^2 in their gender reservation wage gap regressions that is similar to the R^2 in my gender ask gap regressions (0.73 for them vs 0.71 in Column (5) of Table II). In comparison, most gender wage gap studies have R^2 in the range of 0.4-0.5 (for a review of R^2 in gender pay gap studies, see Table 10 in O’Neill and O’Neill (2006)).

D.2 External validity of the bid gap

The adjusted bid gap on the platform (2.2%) is smaller than the residual pay gap found in population surveys such as the PSID or the CPS. For instance, Blau and Kahn (2017) found an 8.4% adjusted gap in the 2010 PSID. However, when focusing on similar populations (i.e., similar datasets with granular resume information and/or engineering majors), studies have found pay gaps closer to my estimate. For instance, Chamberlain, Zhao, and Stansell (2019) reported a 5.4% adjusted pay gap in the Information Technology industry on Glassdoor in 2019. Using administrative UI datasets with granular resume information, Fluchtmann et al. (2021) and Le Barbanchon, Rathelot,

and Roulet (2021) respectively found a 1.9% residual wage gap in Denmark in 2015-2017 and a 3.7% residual wage gap in France between 2006 and 2012. This result further aligns with Goldin (2014), who postulated that among top earners, the wage gap is smaller in tech occupations, which do not require the long and unpredictable hours of workers such as lawyers or doctors, for whom the return on extra hours is much higher.³⁰ Additionally, it could be that the estimates of the bid gap are a more accurate rendition of the residual pay gap, since I can explain close to 100% of the variations in bids with my controls, leaving little room for omitted variable bias.

³⁰ Mas and Pallais (2017) also highlighted the fact that women, particularly those with young children, have a higher willingness to pay to work from home and to avoid employer scheduling discretion.

E Racial gap

When creating their profiles, candidates are invited to voluntarily disclose their race. This information is not displayed on the profile that companies see.³¹ 27.6% of the sample (i.e., about 31,200 candidates) decided to report their race. In this sub-sample, 48.3% are White, 40.1% are Asian, 4.6% are African American, and 7.4% are Hispanic.³²

Column (1) of Table E.1 reports estimates of the raw race ask gap and Column (2) reports estimates of the adjusted race ask gap. Once we control for resume characteristics, there is a small ask gap between candidates who identify as White and those who identify as African American (2.1%), Asian (-0.6%), or Hispanic (1.5%). Columns (3) and (4) provide estimates of the raw and adjusted race bid gaps, respectively. Resume characteristics explain the majority of the race bid gap. The bid gap between White and other races is insignificant and highest for African American at 0.6%. Adding the log ask salary as a control in Column (6) brings all coefficients further down to zero. Similar to the gender bid gap, the coefficients on the interactions between race variables and the log ask salary in Column (7) are all insignificant. Columns (8) to (12) provide estimates of the gap in final offers by race. Due to the restricted sample size estimates are noisy. Similar to the gender final offer gap: the resume characteristics in Column (9) can only explain part of the raw final offer gap of Column (8). However, adding the ask salary as a control in Column (11) does not fully close the final offer gap. In particular, Hispanics stand out with a 2.2% final offer gap. One notable difference is the large, although imprecise, coefficient on African American (-0.027) and the interaction term between African American and the log ask salary in Column (12) (-0.236), suggesting that African American candidates are getting lower final returns to asking for more than White candidates do. Another difference with the main paper results is that, conditional on resume characteristics, the gender gap in final salaries is insignificant and women earn significantly more (1.7%) controlling for their ask salary. However, this is not due to controlling for race (this result looks the same if we do not control for race), but rather a feature of this selected subsample.

These results are suggestive of a larger role for the ask gap, not only in women’s salary determination but more broadly in minority groups’ negotiations. However, self-selection into the sample that declares race and noisier estimates due to the restricted sample size prevents me from drawing definitive conclusions.

³¹ Candidates also have the option to upload a picture of themselves, from which companies can make racial inferences.

³² This sums to 100.4% instead of 100% because a few candidates in the sample declared more than one race.

Table E.1: The Racial Ask, Bid and Final Salary Gap

Dep. Var.:	Log Ask Salary		Log Bid Salary					Log Final Salary				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Female	-0.110*** (0.008)	-0.031*** (0.005)	-0.109*** (0.018)	-0.029*** (0.007)	0.000 (0.002)	-0.000 (0.002)	-0.001 (0.001)	-0.083*** (0.019)	-0.007 (0.011)	0.017*** (0.007)	0.016** (0.008)	0.017*** (0.008)
African American	-0.091*** (0.012)	-0.021*** (0.008)	-0.043* (0.023)	-0.006 (0.009)	0.001 (0.002)	-0.000 (0.002)	-0.001 (0.002)	-0.021 (0.040)	-0.030 (0.021)	-0.028 (0.027)	-0.027 (0.019)	-0.031 (0.020)
Asian	0.040*** (0.005)	0.006 (0.004)	0.090*** (0.009)	-0.004 (0.005)	0.008*** (0.001)	0.001 (0.001)	0.002* (0.001)	0.065*** (0.017)	-0.001 (0.011)	0.018*** (0.007)	0.001 (0.007)	0.002 (0.007)
Hispanic	-0.036*** (0.009)	-0.015** (0.006)	0.020 (0.016)	-0.004 (0.008)	0.004* (0.002)	0.001 (0.002)	0.000 (0.002)	0.033 (0.030)	-0.018 (0.016)	-0.013 (0.011)	-0.022** (0.010)	-0.020** (0.010)
Log Ask salary					0.951*** (0.003)	0.869*** (0.008)	0.869*** (0.008)			0.900*** (0.022)	0.671*** (0.050)	0.711*** (0.033)
Female \times Log Ask salary							0.010** (0.005)					0.002 (0.025)
African American \times Log Ask salary							0.003 (0.010)					-0.236* (0.136)
Asian \times Log Ask salary							-0.009 (0.007)					-0.058 (0.039)
Hispanic \times Log Ask salary							0.006 (0.010)					-0.013 (0.055)
Constant	11.705*** (0.019)	20.146*** (0.700)	11.688*** (0.026)	19.154*** (0.962)	11.578*** (0.009)	12.533*** (0.198)	12.571*** (0.198)	11.614*** (0.068)	24.672*** (2.667)	11.562*** (0.030)	16.046*** (2.029)	15.579*** (1.854)
Candidate's resume characteristics		X		X		X	X		X		X	X
Month \times Year FE	X	X	X	X	X	X	X	X	X	X	X	X
Adj R-squared	0.038	0.623	0.063	0.731	0.935	0.939	0.939	0.057	0.755	0.856	0.880	0.882
Nb. obs	31,241	31,241	106,274	106,274	106,274	106,274	106,274	1,852	1,852	1,852	1,852	1,852

Note: This table assesses racial differences in the ask salary as their role in conjuncture with resume characteristics in the determination of bid and final salaries. The omitted category is the White dummy. It is equal to 1 if the candidates self-identify as White, 0 otherwise. This regression is run on the sub-sample of candidates who self-report their race (27.6% of candidates self-report their race). The non-reported controls in Column (2) are the same as in Column (5) in Table II. The candidate's resume controls of Columns (3) to (7) are the same as, respectively, Columns (1) to (5) in Table III, and Columns (8) to (12) have the same controls as the respective Columns (1) to (5) in Table V. Standard errors are robust for Columns (1) and (2) and are two way clustered at candidate and job id level in Columns (3) to (12). Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

F Updaters analysis

Candidates have the opportunity to update the ask salary displayed on their profile at any time during their spell. Spells on the platform usually only last two weeks, but 7.4% of the candidates still update their ask salary within a spell. Therefore, we can observe, for a given candidate, how bids change when the candidate updates her ask salary. Table F.1 reports the results of a regression of the log bid salary on the log ask salary with individual spell fixed effects, restricting the sample to people who update during a given spell.

It is important to acknowledge that this analysis suffers from a selection problem: candidates do not decide to update at random. In particular, candidates who raise their ask may be reacting to high demand from companies, while candidates updating downwards may be reacting to low demand. This is evident from the gap in offers before the update: candidates who update upwards already have on average seven bids before they update, compared to four for the ones who update downwards, and the average spread between their ask salary and bid salary, before the update, is -\$735, compared to \$-5,132 for the ones who update downwards. However, the exercise can still be informative, as one can read the coefficient on the log ask salary in this context as a lower bound for the true effect of the ask on the bid, since previous bids already partially adjusted for the quality of the candidate. Keeping that in mind, a coefficient of 0.483 (Column (1)) is still relatively close to the 0.849 estimate in Table III Column (4). When splitting the sample, we find that there is an asymmetry: bids increase more when the candidate updates upward (the coefficient on the log ask salary is 0.537 in Column (3)) than when he or she updates downward (the coefficient on the log ask salary is 0.398 in Column (5)). It may seem a priori counter-intuitive that candidates gain more when they increase their ask than they lose when they decrease it. The selection issue can explain this phenomenon: candidates updating downward are reacting to a lack of demand and bids lower than their ask, while candidates updating upward are reacting to a high demand, yet they were, on average, not receiving bids higher than their ask before their update.

Table F.1: The Within-Candidate Effect of a Change of the Ask Salary on the Bid Salary

	All Updaters		Upward Updaters		Downward Updaters	
Dep. Var.: Log Bid salary	(1)	(2)	(3)	(4)	(5)	(6)
Log Ask salary	0.483*** (0.034)	0.484*** (0.041)	0.537*** (0.039)	0.554*** (0.044)	0.398*** (0.056)	0.383*** (0.066)
Female \times Log Ask salary		-0.059 (0.094)		-0.100 (0.112)		-0.012 (0.152)
Candidate's resume characteristics	X	X	X	X	X	X
Month \times Year FE	X	X	X	X	X	X
Adj R-squared	0.954	0.954	0.958	0.958	0.945	0.945
Nb. obs	39,930	39,930	26,113	26,113	13,817	13,817

Note: This table shows the effect of a within-candidate, within-spell change in the ask salary on the bid salary. This model is run on the sub-sample of candidates who update their ask salary during their spell. Columns (1) and (2) contain the sample of all updaters, Columns (3) and (4) only the candidates that update upward and Columns (5) and (6) only the candidates that update downward. Columns (1) (3) and (5) otherwise have the same specification as in Column (4) of Table III and Columns (2) (4) and (6) otherwise have the same specification as in Column (5) of Table III. There are individual spell fixed effects, that is a different dummy for each candidate \times spell. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. There are no other controls since resume characteristics do not vary within a candidate \times spell cell. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

G Extensive margin: other dimensions

Another dimension of the selection into the interview pool is the quality, rather than the quantity, of bids received. To proxy for the quality of the firms that contact candidates, we develop in Roussille and Scuderi (2023) a new methodology to infer candidates' preferences over firms by aggregating interview acceptance and rejection decisions. The utility afforded by an interview request made by firm j to candidate i with bid salary b_{ij} is modelled as the sum of monetary/wage and non-wage components:

$$V_{ij} = \underbrace{u_i(b_{ij})}_{\text{monetary comp.}} + \underbrace{\Xi_{ij}}_{\text{non-wage amenities}}.$$

The non-wage component of utility is assumed to be the sum of a systematic component A_j and an idiosyncratic component ξ_{ij} ,

$$\Xi_{ij} = A_j + \xi_{ij},$$

where $\xi_{ij} \stackrel{iid}{\sim} EV_1$. A_j reflects the component of valuations over firm j 's non-wage amenities that are shared across all candidates, while ξ_{ij} reflects idiosyncratic differences in those valuations. For instance, differences in A_j might reflect whether firm j 's offices are located in a more or less desirable location, while differences in ξ_{ij} might reflect whether candidate i 's commute to firm j would be relatively long or short. Estimates of A_j are constructed by isolating a set of interview acceptances and rejections for which candidates' revealed preferences may be cleanly inferred. In particular, suppose candidate i has offers from both firms j and k *at the same bid salary*. If i accepts j 's interview request, but rejects k 's interview request, we may infer:

$$A_j - A_k \geq \xi_{ik} - \xi_{ij}.$$

Roussille and Scuderi (2023) builds on Sorkin (2018) to develop a maximum likelihood procedure to estimate the mean non-wage amenity value of each firm, A_j , by aggregating these revealed preference inequalities. For the purposes of the analysis here, the estimated amenity values are converted into percentile ranks, where those ranks are increasing in the estimated non-wage amenity value A_j .

Table G.1 reports regressions that predict the quality of bids and final offers received by candidates, as proxied by firms' estimated rankings, as a function of candidate gender and ask salary. All regressions include month-by-year fixed effects. Column (1) illustrates that unconditionally,

women receive bids from firms with somewhat worse amenities than men: the coefficient on the female dummy is roughly -1.8. This unconditional relationship also holds true for final offers (Column (4)). Adding the full set of resume controls, however, eliminates differences between men and women in the quality of the bids (Column (2)) and final offers (Column (5)) they receive. Indeed, additionally controlling for the ask salary actually reverses the gap: conditional on the ask (plus other resume characteristics), women receive bids from firms with slightly better amenities than men, although the coefficient on the female dummy is not statistically significant. Even more starkly, the gap in the amenity values of final offers is completely reversed (Column (6)): unconditionally, women receive final offers from firms ranked 1.24 percentiles lower than men, but conditional on ask, women receive bids from firms ranked 1.25 percentiles *higher* than men.

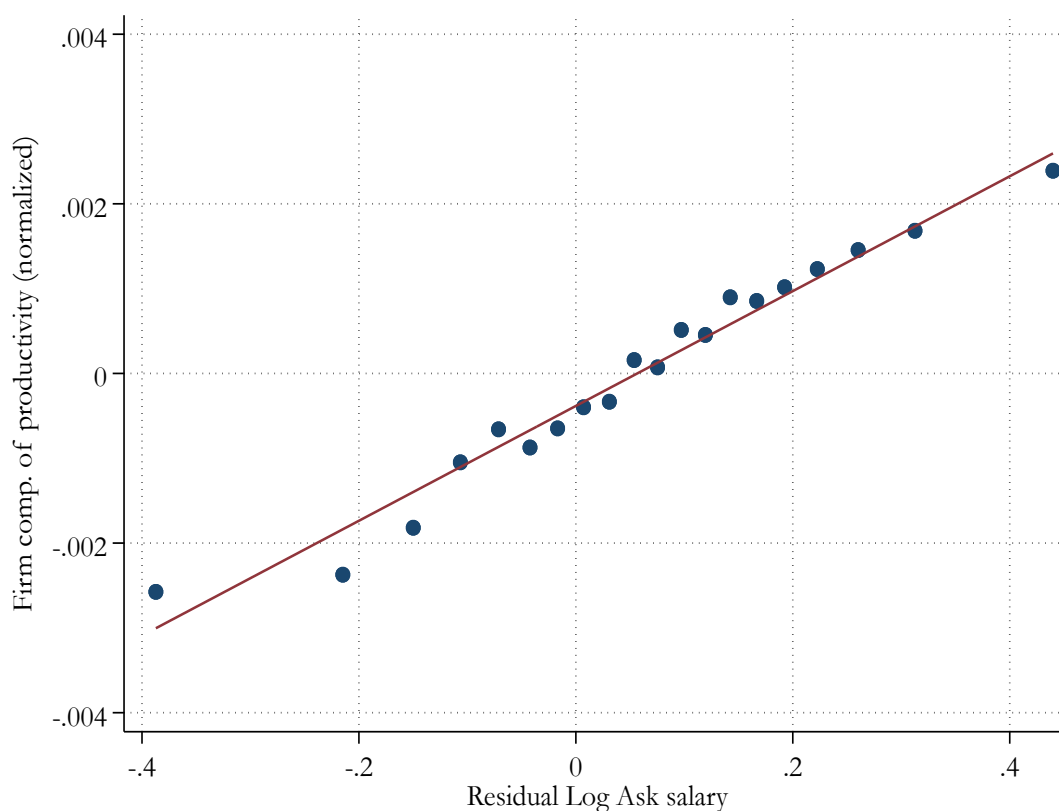
Table G.1: The Ask Salary as Signal of Quality: Relationship between Firm Rank and Residual Log Ask

Dep. Var.:	Firm Rank (Bid)			Firm Rank (Final)		
	(1)	(2)	(3)	(4)	(5)	(6)
Female	-1.795*** (0.406)	-0.042 (0.283)	0.163 (0.278)	-1.241 (1.181)	0.856 (1.167)	1.254 (1.170)
Log Ask salary			8.785*** (0.755)			13.157*** (2.427)
Constant	60.840*** (1.337)	227.970*** (36.277)	167.015*** (35.369)	60.013*** (2.993)	311.712 (253.111)	146.620 (253.526)
Mean rank percentile	62.5	62.5	62.5	64.3	64.3	64.3
Resume Characteristics		X	X		X	X
Month \times Year FE	X	X	X	X	X	X
Adj R-squared	0.004	0.042	0.045	0.005	0.088	0.096
Nb. obs	259,749	259,749	259,749	3,454	3,454	3,454

Note: This table assesses whether there are gender differences in the ranking of firms sending interview requests to candidates. It reports estimates for the subsample of data for which firms can be ranked (see Roussille and Scuderi 2023). The ranking estimates are normalized as percentiles from 0 to 100 where 100 is the best possible ranking. Column (1) controls for gender and time fixed effects at the Month \times Year level. Column (2) adds resume characteristics as controls and Column (3) includes candidates' mean-centered log ask salary. Column (4) to (6) progressively add controls in the same way as Column (1) to (3) but on the subset of observations that include a final offer. Standard errors (in parentheses) are two-way clustered at the candidate and job id level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Finally, in addition to providing estimates of the non-wage components of utility associated with each firm, Roussille and Scuderi (2023) also estimate a structural model of firm demand. Given structural estimates of labor demand, the expected match productivity of any candidate can be inferred from firms' bids (see Roussille and Scuderi (2023) for detail on the estimation of firms' labor demand). Figure G.1 plots the relationship between the average (normalized) productivity of firms and the residualized log ask salary of candidates in the sample of bids made by firms in the San Francisco Bay area. There is a clear, increasing relationship between the (residual) ask salary candidates list and the mean productivity of firms that bid on those candidates: candidates with higher (residual) asks tend to receive bids from more productive firms.

Figure G.1: Binned scatter plot of firm productivity as a function of the residual log ask salary



Note: This figure plots the relationship between (a normalized measure of) the firm component of match productivity (as in Roussille and Scuderi 2023) and candidates' log ask salaries (residualized on candidates' resume characteristics) in the sample of bids made to candidates in the San Francisco Bay Area.

H Appendix Model

This section provides a framework that formalizes the insight that the ask salary can be a signal of quality.

Wolinsky (1983) shows that, in a context of imperfect information about product quality, there exists a separating fulfilled expectations equilibrium: each price signals a unique quality level and if all agents expect that a product at price p has quality q , then this expectation is in turn fulfilled by agents' behavior.

In this section, I adapt Wolinsky (1983)'s model to the labor market and show that there exists an *at least*-fulfilled expectations separating equilibrium: each ask salary signals a unique candidate quality and if all firms expect that a candidate with ask salary a has quality q , then this expectation is in turn either fulfilled or surpassed by the candidate.

The intuition for the equilibrium in this model can be summarized as follows. At a given ask salary, firms expect a certain quality of candidate. A candidate who asks for a given salary may turn out to be of lower quality, but information revealed during the interview will enable some potential firms to discover it. Therefore, in deciding whether to ask for a higher salary than what the firm expects given her quality, the candidate weighs the decrease in her chances of being hired against the gain in salary that results if she gets an offer. If the chances of detection are large enough to outweigh the potential salary gains, it is best for the candidate to signal her true quality with her ask salary.

In this model, women have downward biased beliefs about the salary they can ask for that stems from inaccurate information about the equilibrium. There is no mechanism in the model for firms to learn about these biases because interviews go equally well for men and women. This feature of the model comes from the way the signal is designed: the candidate's quality signal revealed during the interview can be interpreted as "red flag", that is whether the candidate falls below her expected quality. However, in equilibrium, neither men nor women end up triggering this flag. When women are debiased they ask for more but interviews still lead to a hire because they remain strictly above the firm's minimal quality threshold.

The key testable prediction that comes out of this model is a link between the number of bids and the ask salary ala Figure A.6. In the model, firm types are characterized by the range of ask salaries they are willing to interview. I use my data to measure the interview strategies of large employers and show that demand for ask salaries is upward-sloping over the same range as found

in Figure A.6 and downward-sloping afterwards, thereby corroborating the ask salary as a signal of quality mechanism.

H.1 Sequence of actions

There are two types of agents, the firms, denoted by $j \in J$, and the candidates, denoted by $i \in I$. First, Nature draws a discrete quality $q \in Q$ and gender $g \in (m, f)$ - male or female - for candidate i . Second, candidate i sets her ask salary. Third, the firm decides whether to ask a candidate for an interview. During the interview, the firm pays a small cost c and receives a signal about the quality of the candidate. Based on this signal, the firm decides whether or not to hire this candidate.

H.2 Firms

Each firm seeks to hire one candidate. Firms can observe candidates' ask salary a_i but not their quality q_i so they have to form expectations about the latter. $q^e(a_i)$ are common point expectations that assign a single quality level, $q \in Q$, to each candidate ask salary a_i , where $q^e(\cdot)$ is increasing in a_i . $q^e(a_i)$ is labelled common point expectations because all firms believe that quality $q^e(a_i)$ is offered by candidates that ask for a_i . Firms differ in the returns they derive from the expected quality of candidate i , $q^e(a_i)$. In particular, a firm of type k ($k = A, \dots, Z$) derives a benefit of $m_k(q^e)$, where $m_k(\cdot)$ is strictly increasing in q^e , but the slope differs by type k . One can think of m_k as a match-productivity parameter: the highest quality types have a higher return if the job includes complex tasks. A firm decides whether to hire the candidate only after interviewing him/her. In the course of the interview, firm j gets some information on q_i . This information is represented by a signal d_i^j which depends on q_i and random factors. Without loss of generality, assume that the conditional distribution of d_i^j is identical for all j . Let $D(t, q)$ denote the distribution function of d_i^j conditional on q_i .

$$D(t, q) = \mathbb{P}(d_i^j < t | q_i = q)$$

It is assumed that there is a positive probability that d_i^j will enable firm j to establish with certainty that candidate i 's quality is not greater than q_i . Formally:

- For every q there exists a t such that $D(t, q) = 0$.
- Let $t(q) = \sup\{t | D(t, q) = 0\}$, then $q_1 < q_2 \longrightarrow t(q_1) < t(q_2)$.

Similar to Wolinsky (1983), the ask salary of the candidate is modelled as a price: to interview

the candidate, the firm commits to paying them their ask salary if the candidate is hired. In other words, we do not consider that the firm can bid or make a final offer below or above the ask salary of the candidate.³³

Before hiring, a firm can interview as many candidates as it likes. Any interview, whether or not it ends in a hire, involves a fixed cost of interviewing c . If firm j goes to n candidates and ends up hiring candidate i at their ask salary a_i , its expected profit will be $m_k(q^e(a_i)) - a_i - nc$. The firm's goal is to maximize this expression.

H.3 Candidates

Candidates choose their ask salary so as to maximize their expected benefit, that is the product of their ask salary and the probability that they get hired.

Candidates of type q want to maximize the salary offers they receive for a given q^e . We assume Bertrand competition between firms. Therefore, candidates set their ask salary so that the expected benefit to a firm of type $k_i^* = \arg \max [k' \in K | m_{k'}^q(q^e(a_i))]$ is zero and $g \in (m, f)$. This means that for each expected quality level, candidates target their ideal firm, that is the one that has the most returns, and therefore the highest willingness to pay for, that expected quality level. In equilibrium, as described in Section H.5, this condition provides a mapping between a candidate's quality and their ask salary.

The cost of interviewing to the firm, c , is common knowledge. Firms assume $m_k(\cdot)$ is common knowledge as well. However, I model the fact that women ask for less on average as resulting from women receiving downward-biased information about $m_k(\cdot)$. This assumption is in line with the information channel that is put forward to explain the effect of the reform in Section 6. For tractability we model this downward-bias by assuming that women observe m_k^f instead of the true m_k , where $m_k^f(q^e) = m_k(q^e - 1)$, while men received information about the true match-productivity parameter, that is $m_k^m(q^e) = m_k(q^e)$.

For a given type k , it is assumed that firms request an interview from all the (yet unsampled) candidates whose ask salary is in the range that maximises the firm's expected profit. Candidates agree to interview with one of the available firms they received an interview request from, at random.³⁴

³³ This simplification allows the model to focus on its main narrative, which is that the ask salary is a signal of quality. It also does not depart radically from the empirical evidence: 78% of the bids are made exactly at the ask salary of the candidate.

³⁴ For Bertrand competition between firms to hold we need to assume that, for each candidate, there are at least two firms willing to hire them based on their ask salary.

After an interview, the probability that the candidate is hired is 1 if their ask salary signals at least their true quality q_i and $1 - D[t(q'_i), q_i]$ if their true quality q_i is less than the quality q'_i their ask salary suggests.

The goal of candidate i with quality q is to maximize her expected benefit. If the candidate signals her true quality q_i the expected benefit is $a_i(q_i)$. If the candidate signals a quality q'_i above the true quality q_i , the expected benefit is $a_i(q'_i) \times (1 - D[t(q'_i), q_i])$. Here we assume the candidate can interview with at most one firm: if he/she does not get hired, the other firms will know this candidate lied about their quality and therefore will not want to interview him/her.

H.4 Separating equilibrium conditions

Definition: A separating equilibrium is characterized by common point expectations $q^e(a)$, candidates I with ask salary choices \bar{a}_i , $\bar{A} = (\bar{a}_i)$, and firm strategies $\bar{S} = [(\bar{s}_j)]$, such that the conditions (C-1)-(C-4) hold:

(C-1) Candidate expected-benefit maximization: $\bar{a}_i \in \bar{A}$ maximizes i 's expected benefit, for all $i \in I$

(C-2) Firm's profit maximization: Given \bar{A} and $q^e(\cdot)$, strategy \bar{S} maximizes j 's expected benefit, for all $j \in J$

(C-3) Credibility of firm's strategy: Let i be the first candidate sampled by j . The response prescribed by \bar{s}_j to the event that $d_i^j < t[q^e(a_i)]$ remains optimal after the event has occurred.

(C-4) At least-fulfilled expectations: $q_i \geq q^e(\bar{a}_i)$ for all $i \in I$.

Condition (C-2) requires \bar{S} to be optimal but not necessarily in the face of unexpected deviations. Condition (C-3) extends it by requiring optimality of \bar{S} in the event that firm j encounters unexpectedly a deviating candidate and gets a signal $d_i^j < t[q^e(a_i)]$. The rationale for this requirement is that the equilibrium choices of candidates will depend on how firms "threaten" to respond to deviations, and for these "threats" to convince the candidates they have to be credible in the sense of condition (C-3). Condition (C-4) is an equilibrium concept rather than a condition of individual rationality.

H.5 Solving for the separating equilibrium

The point expectations $q^e(\cdot)$ imply a simple form of firm strategy. Every firm will go to a candidate whose ask salary maximises its profit, it will hire that candidate unless it realizes that this candidate's quality is lower than expected; in that event, it will go to another candidate. On the

candidate side, the key is to find the sufficient condition under which (C-4) holds, that is candidates signal at least their true quality. Signalling a quality q' higher than the true quality q has two opposing effects. On the one hand, the candidate's expected benefit increases because of the salary increase in the event the candidate is hired ($a(q') > a(q)$). On the other hand, the candidate's expected benefit decreases because of the increased risk of being detected as lower quality and therefore not being hired. In particular, the probability of getting hired when signalling one's true quality is 1, while it is only $(1 - D[t(q'), q])$ when signalling q' . Putting these together, the expected benefit from declaring q' is $a(q')(1 - D[t(q'), q])$ while that of declaring q is $a(q)$. The sufficient condition under which the candidate does not signal a quality above his true quality is that for all $q, q' > q \in Q$:

$$\begin{aligned} a(q')(1 - D[t(q'), q]) &< a(q) \Leftrightarrow \\ \frac{a(q')}{a(q)} &< \frac{1}{(1 - D[t(q'), q])} \end{aligned} \quad (1)$$

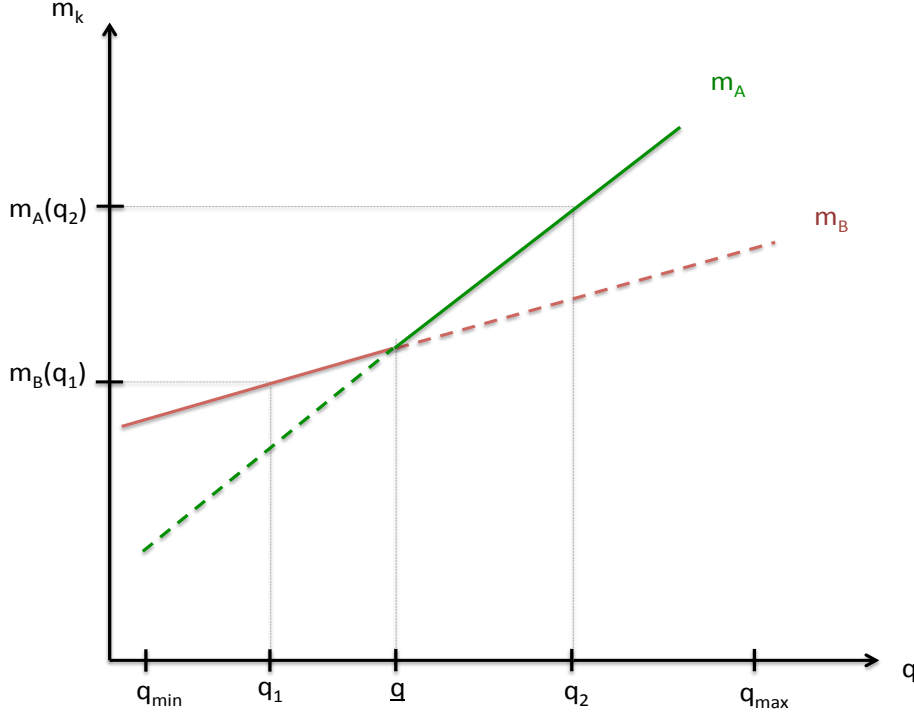
This condition means that for any $q' > q$, the relative increase in the salary in the event of a hire has to be lower than the relative decline in the probability of being hired. Note that there is no point in signalling a quality q'' lower than the true quality q since that reduces the ask salary but does not change the likelihood of hire.

If condition 1 is satisfied, then given firms' behaviour, a candidate i expects to be hired only by firms who happen to choose them for their first interview. This is because the candidate believes that all other candidates ask for salaries matching the quality expected from them and therefore always get hired after an interview. In other words, $n = 1$.

We now have all the elements to construct the exact mapping between a candidate's quality and her optimal ask salary. From Section H.3 and the fact that in equilibrium $q^e = q$, we know that candidates set their ask salary so that the expected benefit of the firm of type $k_i^* = \arg \max [k' \in K | m_{k'}^g(q(a_i))]$ is zero, with $g \in (m, f)$. An example with just two firms and male candidates provides intuition for this condition. Assume firm A and B respectively have match-productivity parameter $m_A(q)$ and $m_B(q)$, where $m'_a(q) > m'_b(q)$ and the two functions cross at \underline{q} as in Figure H.1. Assume $q \in (q_{min}, \underline{q})$, then the optimal strategy is to target Firm B with an ask salary that sets firm B's profit to zero, that is $a = m_B(q) - c$. Conversely, if $q \in (\underline{q}, q_{max})$, Firm A will be targeted and the ask salary will be $a = m_A(q) - c$. From the firm's perspective, this provides a one-to-one mapping from firm types to ask salary ranges: firms of type A will interview

candidates whose ask belongs to $(q_{min} - c, \underline{q} - c)$ and firms of type B will interview candidates whose ask belongs to $(\underline{q} - c, q_{max} - c)$.

Figure H.1: Candidate's Target Firm Choice - Two Firm Example



Note: This figure represents, in equilibrium, the relationship between the quality of candidates and the match productivity parameters of two firms with different types, A and B. Firm A's $m_A(q)$ is in green and Firm B's $m_B(q)$ is in red. The solid line is the upper envelope of the two firms' matching parameters and indicates the candidate's target firm at each quality level. Firm A is targeted for quality in the (q_{min}, \underline{q}) range and firm B is targeted for for quality in the (\underline{q}, q_{max}) range.

Given their information about m_k , the optimal ask salary for women is $a_{if}^* = m_k^f(q_i) - c = m_k(q_i - 1) - c$ and for men is $a_{im}^* = m_k^m(q_i) - c = m_k(q_i) - c$. Given these ask salaries, firms predict that women are of quality $q_i - 1$ and offer them $m_k(q_i - 1) - c$, while they predict men are of quality q_i and offer them $m_k(q_i) - c$. The key to this being an equilibrium outcome is that while women are indeed of higher quality for a given ask salary, firms do not know or learn about it in this setting. This feature of the model comes from the fact that, in equilibrium, the probability of detecting that men are below the expected quality is the same as for women (zero) and therefore the recruitment process per se does not reveal gender-specific information about quality. Conversely, women with quality q_i would benefit from asking for more but, given their beliefs, asking for a_{if}^* is the expected

benefit-maximising choice.

H.6 Reform impact for women

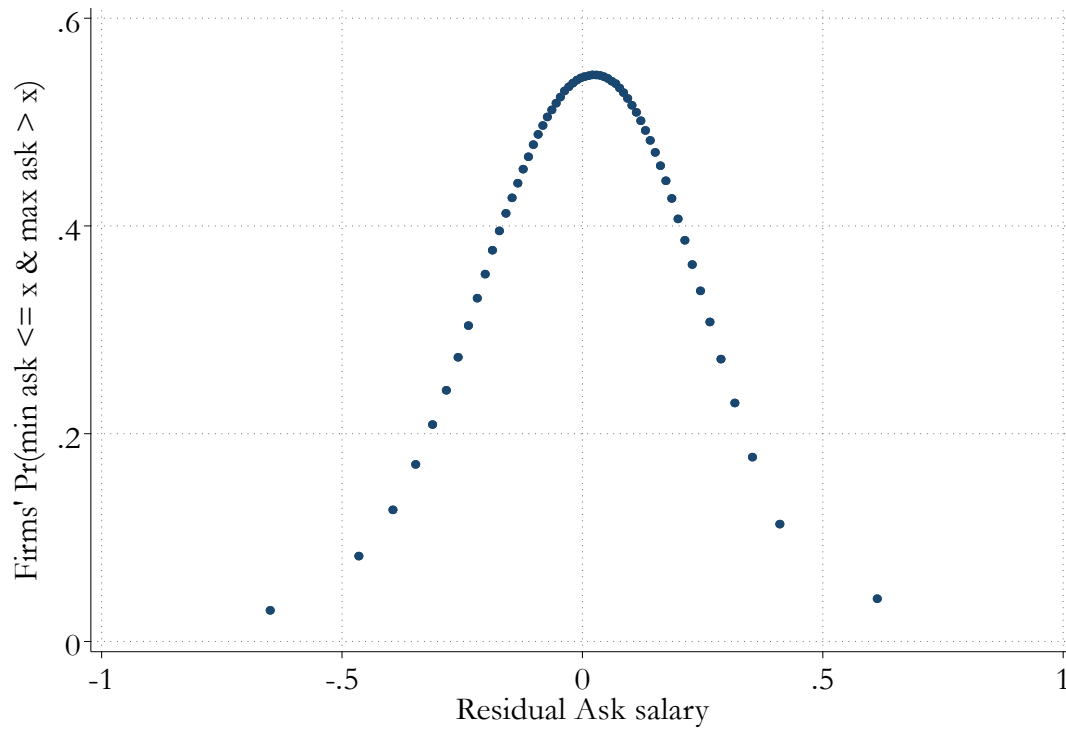
Assume women were provided with information on $m_k(\cdot)$, so that their beliefs now match those of men.

Impact on salary: Before the information provision, women asked for $m_k(q - 1) - c$, firms predict that women are of quality $q_i - 1$ and offer them $m_k(q_i - 1) - c$.³⁵ After the information provision, women of quality q shift their beliefs to $m_k^{lf} = m_k(q^e)$ and, in equilibrium, they ask for $m_k(q) - c$, companies think they are of quality q and offer women $m_k(q) - c > m_k(q - 1) - c$. In sum, women ask for more - the same as men - and correspondingly get more.

Impact on the number of interview requests received: For a given type k , firms request an interview from all the (yet unsampled) candidates asking for a salary that is in the range that maximises the firm’s expected profit. As described earlier, this provides a one-to-one mapping between firm types and a profit-maximising ask salary range in which each firm type interviews. Therefore, the model predicts that the impact of the change in women’s ask on the number of interview requests received depends on two elements: (1) the empirical distribution of the number of firms willing to interview at a given ask salary and (2) where in this distribution women lie before and after the information provision. We can define as “min ask” the lowest ask salary at which a given firm interviews and “max ask” as the highest one. Figure H.2 then describes, for a given ask salary, the share of firms on the platform for which this ask salary fits into their interview range, [min ask, max ask]. Given the mapping between the profit-maximising ask salary range and firm types, this figure provides the empirical distribution of firm types. From this bell-shape relationship we can conclude that a sufficient condition for the information provision to have had a positive impact on the number of interview requests received by women is that enough of them had an initial ask salary that was in the increasing range of the figure.

³⁵ This assumes that firms do not know that women, for a given ask, are better than men. This assumption is in line with recent evidence from Feld et al. (2022). This paper tests whether there is statistical discrimination in the tech section. In a field experiment, they measure the programming skills of job applicants for a programming job and recruit HR tech professionals to assess the performance of these applicants. They find that while there are no significant gender differences in performance, employers believe that female programmers perform worse than male programmers.

Figure H.2: Empirical Distribution of Firm Types



Note: The model predicts that, in equilibrium, a firm's type is defined by the range of ask salaries at which it interviews, $[\min \text{ ask}, \max \text{ ask}]$. This figure therefore provides, for each ask salary, the share of firms for which this salary falls in their type's interview range. The sample is restricted to firms that send more than 20 interview requests.

I Data build

I.1 Sample restrictions

The sample was restricted to profiles with an ask salary between \$30,000 and \$999,999. The bid and final salaries were also restricted between \$30,000 and \$999,999. A manual check of a random subset of the profiles and bids beyond this range suggested that the candidate or the firm had made a typo when typing something above \$999,999. For profiles with salaries below \$30,000 they often indicate what seems to be a per hour rate and correspond to candidates looking for part-time work or consulting missions, which is not the aim of the platform. We drop 2.6% of the raw data by using these salary restrictions.

I.2 Gender

Gender is an optional field on the profile and only 50% of the candidates self-declared their gender. In order to obtain a gender for the other 50%, I use a prediction algorithm based on first names. The algorithm can be found on [this website](#). The prediction can take 5 values: male, mostly male, ambiguous, mostly female and female. When available, I used the self-declared gender of the candidate, otherwise I used the predicted gender only if it predicts that the person is male or female. Reassuringly, for the sub-sample that self-declares their gender (i.e. 50% of the full sample), I verify that the algorithm guesses their gender incorrectly only 0.6% of the time. 14.6% of the profiles remain without an assigned gender. These profiles are coded as “unknown” in the categorical female variable. Coefficients on this category are not displayed in the main tables¹ but the “unknown” gender observations remain in the sample for estimation precision.

I.3 Education

Candidates provide four different information items for each education institution they list: the degree they received (B.A., M.A. etc.), the graduation year, the study field and the university they went to. All this information is manually entered by the candidate and therefore requires some cleaning. I first implemented the standard cleaning procedures for manual data (remove punctuation, spaces, capitalize etc.). I then looked into each field separately:

Education level: I created five groups of education level: high school, associate (two-year de-

¹ They usually lie somewhere between the Male and Female coefficients, often closer to the Male one given the gender imbalance of the sample.

grees), bachelor (4 year degrees), master (2 year post-bachelor), mba, phd. Data was then matched to those groups by using the comprehensive list of names that could match each of them. For instance, PhD would include all with observations that include “ phd ” or “ph ” or “doctor ” or “ dphil ”.² I then selected the highest degree achieved by the candidate for the regression analysis.

Study field: For the study field, I identified whether the candidate had a computer science background using the following list of keywords in each of their education institutions: “comput”, “cs”, “software”, “programm”, “ web ”, “ informatic”, “developer”, “ systems ”, “ it ”, “ information tech”.

Education quality: In order to get standardised names for the university that the candidates attended I used an open-source software developed by Google called “Open Refine”. The algorithm matches hand written university name to the university name on [the wikidata education project](#). This requires two steps: Step 1) The names were clustered using fingerprint and 2-gram fingerprint. Step 2) open refine then matched standardised university names to the wiki database. The “reconciliation process” is described in detail in [this post](#). I indicated that I was looking to match “educational institutions” to narrow the search on the database. The software then returns a potential match with a match score (the method used is Dice coefficient). Observing the scores and the corresponding universities, I decided that all entries with scores above 50 (which represents the likelihood that this is match) would be matched to wikidata. This leaves only about 12% of universities unmatched. I then used this dataset of college names to check whether the candidate ever attended an Ivy Plus League school, that is Ivy Leagues + U. Chicago, Stanford, MIT, and Duke. I also added the schools that are ranked in the top 5 programs in engineering by the annual U.S. News college ranking. Specifically these schools are UC Berkeley, California Institute of Technology, Carnegie Mellon University and Georgia Institute of Technology. I considered that this dummy would be 0 for those with unmatched university names. Adding this dummy did not have an impact on the gender ask gap. As a robustness test, for the subset of U.S. universities, I added categorical controls for the tier of the school, Barron’s Selectivity Index and Average SAT scores in 2013 using the dataset on College Level Characteristics from the IPEDS Database and the College Scorecard from Chetty et al. (2020). Because it didn’t change the coefficients on the female dummy and restricted the sample due to the U.S. education limitation and missing rankings, this variable

² Meticulous attention was paid to the inclusion of spaces in order to select the right observations in the education level groups. For instance I would include those who have “ph + a space” in their degree, so as not to include all the ones that mistakenly list a major that starts or contains “ph”, such as philosophy, in that field. I manually checked a large sample of these education levels. The same process was implemented for the subsequent education and firm history categories.

is left out of the current version of the models. I also tried a version with school Fixed Effects which had no impact on my results and restricted the sample. **Year of graduation:** I selected the year of graduation from college as a proxy for age for the (large) subset that has a college degree. Because it didn't change the coefficients on the female dummy and restricted the sample due to missing year of graduation for some candidates, this variable is left out of the current version of the models.

I.4 Work Experience

Candidates manually enter the history of firms that they worked at, the job titles they held and the duration of their work at each firm. I use this data to construct a dummy for whether the candidate has ever worked in one of the FAANG (Facebook, Amazon, Apple, Netflix, Google) companies, a variable for how many jobs the candidate ever held and the average tenure at each job. I also use this data to compute the highest job title the candidate ever held. This first-order approximation for work experience does not impact the coefficients on the female dummy or the bell shape of the relationship between the residual ask salary and the number of interview requests received.

In the specification where I control for most recent firm fixed effect (Table II Column (6)), I first clean the name of the firms. That is I normalise company names to ensure minimal existence of singleton groups by stripping organization types such as “Ltd.” or “Inc.” from the name, converting all names to lowercase, and removing punctuation as well as special characters. I then match the company names with (identically cleaned) data from Burning Glass containing around two million company names. To account for slight deviations in the way candidates entered their recent employer on their profile, I use a fuzzy matching approach based on a variant of Levenshtein distance and probabilistic record linkage.

References

- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy* 113 (1): 151–184.
- Blau, Francine D., and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature* 55(3):789–865.
- Chamberlain, Andrew, Daniel Zhao, and Amanda Stansell. 2019. "Progress on the Gender Pay Gap: 2019." *Glassdoor Economic Research*.
- Chen, Shuowen, Victor Chernozhukov, Iván Fernández-Val, and Ye Luo. 2019. "Sorted Effects: Sorted Causal Effects in R." *arXiv preprint arXiv:1909.00836*.
- Chernozhukov, Victor, Iván Fernández, and Ye Luo. 2018. "The Sorted Effects Method: Discovering Heterogeneous Effects Beyond Their Averages." *Econometrica* 86 (6): 1911–1938.
- Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan. 2020. "Income Segregation and Intergenerational Mobility Across Colleges in the United States." *The Quarterly Journal of Economics* 135 (3): 1567–1633.
- Feld, Jan, Edwin Ip, Andreas Leibbrandt, and Joseph Vecchi. 2022. "Identifying and Overcoming Gender Barriers in Tech: A Field Experiment on Inaccurate Statistical Discrimination." *CESifo Working Paper*, no. 2205.
- Fluchtmann, Jonas, Anita Glenny, Nikolaj Harmon, and Jonas Maibom. 2021. "The Gender Application Gap: Do Men and Women Apply for the Same Jobs?" *IZA Discussion Paper*, no. 14906.
- Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104 (4): 1091–1119.
- Goldin, Claudia, Sari Pekkala Kerr, Claudia Olivetti, and Erling Barth. 2017. "The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census." *American Economic Review: Papers and Proceedings* 107(5):110–114.
- Krueger, Alan B., and Andreas I. Mueller. 2016. "A Contribution to the Empirics of Reservation Wages." *American Economic Journal: Economic Policy* 8:142–179.

- Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet. 2021. “Gender Differences in Job Search: Trading off Commute against Wage.” *The Quarterly Journal of Economics* 136 (1): 381–426.
- Mas, Alexandre, and Amanda Pallais. 2017. “Valuing Alternative Work Arrangements.” *American Economic Review* 107 (12): 3722–3759.
- O’Neill, June, and Dave O’Neill. 2006. “What Do Wage Differentials Tell Us about Labor Market Discrimination?” *The Economics of Immigration and Social Policy* 24:293–357.
- Roussille, Nina, and Benjamin Scuderi. 2023. “Bidding for Talent: A Test of Conduct on a High-Wage Labor Market.” *Working paper*.
- Sorkin, Isaac. 2018. “Ranking Firms Using Revealed Preference.” *The Quarterly Journal of Economics* 133 (3): 1331–1393.
- Wolinsky, Asher. 1983. “Prices as Signals of Product Quality.” *The Review of Economic Studies* 50 (4): 647–658.