

THE ROLE OF THE ASK GAP IN GENDER PAY INEQUALITY

NINA ROUSSILLE

The gender ask gap measures the extent to which women ask for lower salaries than comparable men. This article 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 résumé and ask salary. If the candidate is hired, a final salary is recorded. After adjusting for résumé 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 prefilled 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.*

* I am indebted to Patrick Kline, Hilary Hoynes, and Gabriel Zucman for their guidance and encouragement on this project. For very helpful comments, I thank David Autor, Anne Boring, David Card, Stefano DellaVigna, Larry Katz, Supreet Kaur, Juliana Londoño-Vélez, Claire Montialoux, Jesse Rothstein, Ben Scuderi, Chris Walters, and five anonymous referees, as well as numerous participants at the Summer Institute Labor Studies session and several other seminars. I also 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 finally 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.

© The Author(s) 2024. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2024), 1–54. <https://doi.org/10.1093/qje/qjae004>. Advance Access publication on February 14, 2024.

“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*

I. INTRODUCTION

Over the past several decades, the raw gender pay gap in the United States 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; Kiessling et al. 2024). 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; Biasi and Sarsons 2022).

This article 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 United States 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 are 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 provide information on only one side of the market: 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 data set 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 résumé 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 document a 6.6% raw ask gap on the platform. After controlling for all the candidates' résumé characteristics, the ask gap is still 2.9%. In other words, women ask for 2.9% less than men with comparable résumés. 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-Val, 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 have 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' résumé 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 résumé characteristics are not, this residual bid gap disappears. In other words, while accounting for résumé 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, résumé characteristics account for 3 percentage points 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' résumé 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 subsample of 7,582 hired candidates, gender differences in ask salaries explain nearly all of the gap in final offers.

In particular, while conditioning on résumé 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 prefilled 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 ask salary. Using an interrupted time series design, I show that the new framing of the ask salary elicitation eliminated the ask and bid gaps. These results are driven by women asking for higher salaries after the reform. Further, I find no discernible effect 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 résumés and the reform corrected them.

This article 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 (Olivetti and Petrongolo 2016; Blau and Kahn 2017), but a more recent strand of the literature has turned to investigate gender gaps in salary expectations (Reuben, Wiswall, and Zafar 2017; Kiessling et al. 2024). 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 nonresponse, and real labor market relevance. Finally, the recruitment process on the platform allows for the direct measurement of the effect 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 (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 (Goldin 2014; Bertrand 2017; Garbinti, Gouppille-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 (Jäger et al. 2024; Bennedsen et al. 2022; Baker et al. 2023; Cortés et al. 2023; Cullen and Pakzad-Hurson 2023) 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 this article 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 II provides details on the empirical setting. Section III presents a detailed description of the data. Section IV describes the empirical strategy to estimate the ask gap and documents its existence and magnitude. Section V provides evidence of the impact of the ask gap on the bid gap and final salary gap. Section VI details the reform of the elicitation of candidates' ask salaries and reports estimates of the effects of the reform; Section VII provides a framework to interpret the results of the reform. Section VIII concludes.

II. INSTITUTIONAL SETTING

II.A. Market Description

Several previous papers have studied online labor markets, such as Amazon MTurk, to explore the causes of the gender pay gap (Gomez-Herrera and Mueller-Langer 2019; Litman et al. 2020). These markets allow researchers to run experiments and precisely record the effects 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, on-site, high-wage engineering jobs based in the United States: 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 Payscale'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 Payscale'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

1. Payscale 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.

software engineers are distributed similarly to their equivalent found on Payscale.²

In addition, 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

II.B. 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 candidates apply to each posted job separately. Afterward, the company interviews a selection of applicants and decides whether and whom to hire. In contrast, on Hired.com, companies apply to candidates based on their profiles, and candidates decide whether 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' résumés and asks. Formally, the recruitment process can be divided into the following three sequential steps, also described in [Figure I](#):

1. *Supply Side*. Candidates create a profile that contains standardized résumé entries and the salary that the candidate wants to make in their next job: their ask salary.⁴ [Online Appendix Figure B.1](#) is a screenshot of a typical candidate's

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 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.

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 [Online Appendix Figure B.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.)

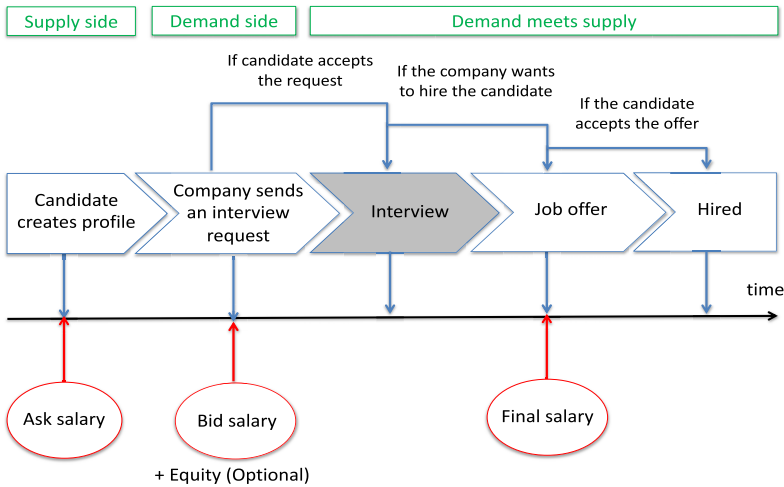


FIGURE I

Timeline of the Recruitment Process on Hired.com

The red circles at the bottom show the different salaries that are captured on the platform. The blue arrow-shaped boxes in the middle describe all the steps of recruitment on the platform from profile creation to hiring. The gray shading for the interview stage indicates that I do not have metadata from companies about their interview process. On the top of the diagram, in green boxes, is the classification of the recruitment process between the labor demand side (companies) and the labor supply side (candidates).

profile, and [Online Appendix Table A.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.

2. *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. [Online Appendix Figure B.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.

3. *Demand Meets Supply.* Hired.com records whether the candidate accepts or rejects the interview 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 nonbinding, 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

II.C. Relevance of the Recruitment Process to Other Wage Bargaining Settings

Although 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 with 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 nontrivial share of wage negotiations, candidates are asked for their desired salary before the company makes them an offer. For instance, in a

5. Although 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.

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).

III. DATA

III.A. 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 data set has 113,777 candidates, 39,839 jobs, and 6,532 firms 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, resulting in 7,582 final offers. The data span several recent years but per the research contract signed with the company, the exact start and end dates of the period cannot be disclosed.

III.B. Gender

Gender is an optional field on the profile, and only 50% of the candidates self-declared their gender. To obtain gender data for the other 50%, I use a standard prediction algorithm based on first names.⁶ Reassuringly, for the subsample 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, and men represent the remaining 80.2%.

III.C. Candidate Summary Statistics

[Table I](#), Panel A provides information on the résumé characteristics of the candidates. They have, on average, 11.3 years

6. The prediction can take five 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.”

TABLE I
DESCRIPTIVE STATISTICS ON CANDIDATES AND COMPANIES

	All	Male	Female	Difference	<i>p</i> -value
Panel A: Descriptive statistics on candidates					
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	.000
Probability of accepting an interview request	62.2	62.0	63.2	-1.2	.000
Education					
Share with a bachelor	97.6	97.3	98.7	-1.4	.000
Share with a master	41.4	40.3	45.2	-4.9	.000
Share with a CS degree	55.2	57.2	47.7	9.5	.000
Share with an IvyPlus degree	9.4	8.7	11.8	-3.1	.000
Preferences					
Share looking for full-time job	96.9	96.7	97.7	-1.0	.000
Share looking for a job in San Francisco	31.6	30.0	37.5	-7.5	.000
Share in need of visa sponsorship	13.6	13.0	15.7	-2.7	.000
Work history					
Years of total experience (mean)	11.3	11.7	10.1	1.6	.000
Share that worked at a FAANG	6.0	6.0	6.0	0.0	.679
Share leading a team	32.7	33.8	27.6	6.2	.000
Share employed	73.1	74.0	69.7	4.3	.000
Number days unemployed (median)	120	116	133	-17	.000
Occupation					
Share of software engineers	61.7	66.6	43.2	23.4	.000
Share of web designers	8.3	6.1	16.6	-10.5	.000
Share of product managers	8.3	7.5	11.4	-3.9	.000

TABLE I
CONTINUED

Panel B: Descriptive statistics on companies					
Number of:	Firms	Jobs	Bids sent	Final offers	Cities
	6,532	39,839	463,860	7,582	20
Revenue (yearly, million US\$)	1–25	26–100	101–500	501–1,000	1,000+
Share ($N = 962$)	47%	17%	12%	14%	10%
Firm age (years)	0–5	6–10	11–15	16–20	20+
Share ($N = 2,249$)	36%	45%	11%	4%	4%
Firm size (no. employees)	1–10	11–50	51–200	201–500	500+
Share ($N = 2,368$)	18%	29%	31%	11%	11%
Top three locations	SF	NY	LA		
Share ($N = 4,319$)	40%	24%	7%		
Top three industries	Software	Finance	Analytics		
Share ($N = 2,253$)	15%	10%	8%		

Notes. 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). FAANG is a dummy for whether the candidate has ever worked in one of Facebook, Amazon, Apple, Netflix, or Google. 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.

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 three out of four 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 versus 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% versus 47.7%). Finally, women are more likely to be looking for a job in San Francisco (37.5% versus 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. [Online Appendix Table A.2](#) presents gender differences in these preferences controlling for candidates' résumé 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 0.5 percentage points lower for women with the same résumé characteristics).

III.D. 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 five years of the end of the sample period, half of them report less than US\$25 million, in revenue, and almost half the firms enlist 1 to 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, 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 and finance (10%), as well as analytics (8%).

III.E. 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 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 versus 11.6). Similarly, candidates who get hired receive about 1.5 times as many interview requests as the average candidate (6.6 versus 4.5) and they are more likely to accept an interview request.

III.F. How Do the Ask and Bid Salaries Relate to More Traditional Salary Measures?

This article 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 with salary expectations or the reservation wage? Further, given that the bid is nonbinding, 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, that is, how much people expect to make in their next job (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

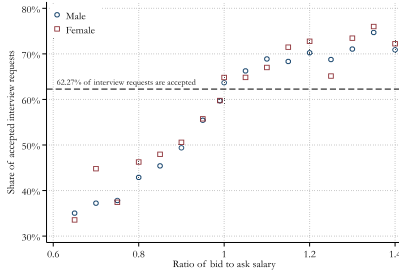
7. Candidates can request to remain visible for two to four additional weeks; 55% of the candidates are live for two weeks, 22% remain visible for four, and the remaining 23% for six.

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 to maximize their final offer.

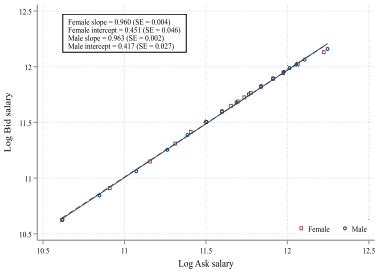
Candidates can adopt different strategies for the choice of the ask salary. Some may choose to record their reservation wage, that is, the lowest wage at which they would accept a job. Others may provide an estimation of their market value, and some may put the highest salary at which they think they can be hired. These possible interpretations are, to some extent, testable because 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 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.” [Online Appendix Figure B.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, although this reason is virtually never brought up when the ask is equal to or below 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.

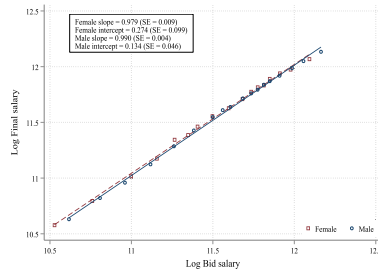
The bid salary is what firms declare they are willing to pay the candidate solely based on their profile, before any interaction with them. The final salary is offered to a candidate at the hiring stage. Given that companies are by no means contractually bound



(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

FIGURE II

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

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 Panel B and the log bid and log final salary offers in 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 US\$10k from the ask, while 36% of final salaries match the initial bid exactly and 78% of final salaries are within a range of US\$10k from the bid. Panel B includes the 463,860 observations with an associated bid and Panel C the 7,582 observations for which there is a final offer.

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. In addition,

36% of all final offers are identical to the bid, and 78% of all final offers are within \$10,000 of the bid.

IV. DOCUMENTING THE GENDER ASK GAP

IV.A. *The Gender Ask Gap: Methodology*

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

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

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. When collapsing the data to the candidate level, I select as Ask_i the first listed ask of candidate i .⁸

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

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

where the controls X_i are the candidates' résumé characteristics, as described in detail in [Online Appendix Table A.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 résumé characteristics capturing, for instance, education quality and work history. As in [equation \(1\)](#), Ask_i is the (first listed) 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.

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

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

where Ask_{ib} is the ask salary of candidate i when he or she receives her b th bid, 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 . Hence, in this specification, a candidate appears

8. The results are qualitatively the same if we opt for the last ask salary ([Online Appendix Table A.3](#)).

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	X	X	X	X
Month × 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
No. observations	113,777	113,777	113,777	113,777	113,777	63,916	463,860

Notes. This table presents estimates of β_0 from equations (1), (2), and (3). Column (1) controls for gender and time fixed effects at the month × year level and corresponds to equation (1). Columns (2) to (6) 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 Online Appendix Table A.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 their current job, a dummy for whether the candidate has ever worked in one of the FAANG (Facebook, Amazon, Apple, Netflix, Google), and a dummy for 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) 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 (in parentheses) are used in columns (1) to (6). Column (7) corresponds to equation (3), where the ask gap is estimated at the bid level (such that a given candidate appears as an observation as many times as the number of bids he/she received) and standard errors (in parentheses) are clustered at the candidate level. ^{*} $p < .1$, ^{**} $p < .05$, ^{***} $p < .01$.

as an observation as many times as the number of bids they received. A candidate's ask salary may differ across observations if the candidate updates it over time.⁹ The advantage of this specification is that the units of analysis are the same as those in [Table IV](#), which investigates the relationship between the ask and the bid gap.

IV.B. Results

1. *Graphical Evidence.* [Online Appendix](#) Figure B.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 versus \$121,942).¹⁰

2. *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 résumé 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 résumé characteristics detailed in [Online Appendix](#) Table A.1. This exercise identifies which résumé 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 (see [Online Appendix](#) Table A.4).¹¹ Since zero 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 à la [Belloni, Chernozhukov, and Hansen \(2014\)](#), with two- and three-way

9. A small share (7.4%) of candidates update their ask salary in a given spell. [Online Appendix](#) F discusses the behavior of these candidates.

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

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

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 percentage points. 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 [Online Appendix Table A.2](#) and described in [Section III.C](#), 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 down to 2.9%. This is mostly driven by the coding skills listed on candidates' profiles, not by differences in exposure to an "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 [Online Appendix D.1](#) discusses how the magnitude of the ask gap compares with other related salary measures such as salary expectations or reservation wages.

[Table III](#), 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 two to four years of experience in their current occupation tends to ask for 11.2% more than a candidate with zero to two 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

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 résumé. 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.

TABLE III
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)	(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 × 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)	

TABLE III
CONTINUED

Dep. var.:	Log ask salary	Log bid salary		Log final salary	
	(1)	(2)	(3)	(4)	(5)
Candidate's résumé characteristics	X	X	X	X	X
Month \times year FE	X	X	X	X	X
Adj. <i>R</i> -squared	0.708	0.816	0.954	0.827	0.920
No. observations	113,777	463,860	463,860	7,582	7,582

Notes. This table explores the role of controls other than gender in explaining ask, bid, and final salaries. Column (1) follows [equation \(2\)](#) and it is at the candidate level. The other columns are at the bid level. Column (2) corresponds to [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 set log final salary as the dependent variable. The number of observations is much smaller than in columns (4) and (5) as the unit of observation is restricted to candidates with final offers. 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. * $p < .1$, ** $p < .05$, *** $p < .01$.

degree is a PhD ask for 6.7% more than candidates whose highest degree is a master's.

In [Online Appendix C](#), a classification analysis using the sorted partial effect method of [Chernozhukov, Fernández-Val, and Luo \(2018\)](#) highlights that experience is the résumé 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 résumé characteristics but estimated separately for different experience groups. The ask gap increases considerably with experience: it is insignificant for the zero to four years and four to six years of experience groups and is only 1.5% for the six to eight years of experience group. It then jumps to 4% for the 8–15 years of experience group. The largest gap, for candidates with 15–20 years of experience, reaches 5.4%.¹³

13. While it is beyond the scope of this article to explain this gradient, my analysis of the reform described in [Section VI](#) 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.

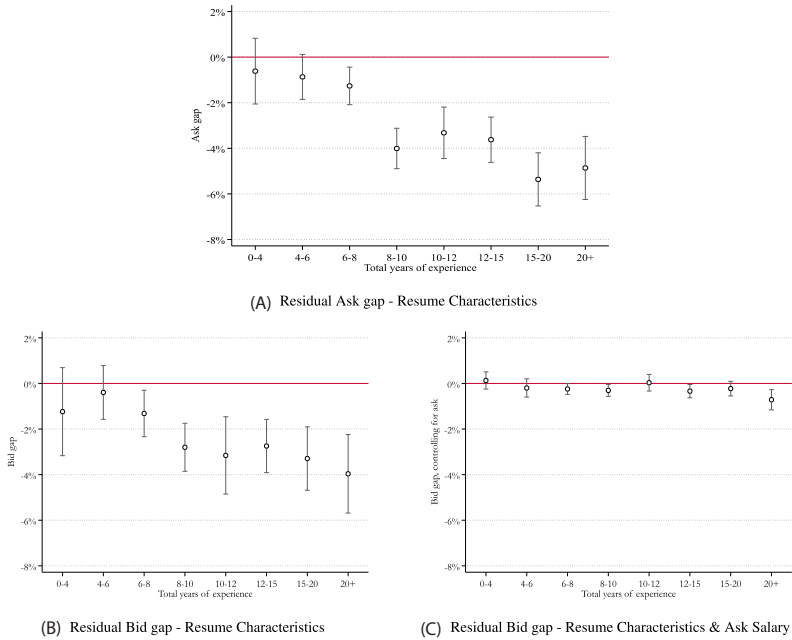


FIGURE III

Heterogeneity in the Ask and Bid Gap by Experience

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. Panel A plots the point estimate of the female dummy in equation (2), where the regression is run separately by total years of experience. Panel B plots the point estimate on the female dummy in equation (5), and 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 résumé characteristics I control for are all the variables described in Online Appendix Table A.1, except the total position experience since regressions are run separately for each total position experience group.

V. DESCRIPTIVE EVIDENCE ON THE ROLE OF THE ASK GAP IN GENDER PAY INEQUALITY

V.A. 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' résumé 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' résumé characteristics. Finally I estimate the effect of the ask salary on the bid gap, with and without the résumé characteristics controls. Formally, these three models can be written as:

Model 1:

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

Model 2:

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

Model 3a:

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

Model 3b:

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

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 résumé characteristics and log ask salary, when he or she receives the 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 . Observations are at the bid level such that, as in [equation \(3\)](#), a candidate appears as an observation as many times as the number of bids they received. A candidate's ask salary may differ across observations if the candidate updates it over time and, more systematically, bids may differ across observations (for a given candidate) since they are sent by different firms.

V.B. The Gender Bid Gap: Results

1. *Graphical Evidence.* [Online Appendix](#) Figure B.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 versus \$120,720). Furthermore, 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.

2. *Regression Results.* The raw gender bid gap, as estimated by β_1 in equation (4) and reported in Table IV, column (1), is 3.3% and significant at the 1% level. Controlling for the résumé 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 résumé 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, but very small (0.2%). This result persists when we add back all the candidate résumé 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 capture not only wage differences between otherwise similar men and women, but also the fact that the econometrician is limited in the 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' résumé characteristics and their ask salary; as a result, having access to candidates' profiles helps to control for the firms' information sets at the time they make their bids.¹⁵ The R^2 in Table IV validates this overlap between *Hired.com* data and the firm's information sets: the linear model conditioning on candidates' résumé characteristics explains 82% of the variation in bid

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

15. It could still be that firms interpret and interact with the résumé 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 VI 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.

TABLE IV
THE ROLE OF THE ASK AND RÉSUMÉ CHARACTERISTICS IN BID SALARY GENDER DIFFERENCES

	No. job FE				Job FE			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var.: Log bid salary								
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 × 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 résumé characteristics								
Month × year FE	X	X	X	X	X	X	X	X
Job FE								
Adj. <i>R</i> -squared	0.007	0.816	0.950	0.954	0.954	0.014	0.329	0.834
No. observations	463,860	463,860	463,860	463,860	463,860	454,631	454,631	454,631

Notes. This table presents estimates of β_1 from equations (4) to (7), such that these regressions are at the bid level (i.e., the level of observation is the same as in Table II column (7)). All of them also include time fixed effects at the month × 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 Online Appendix Table A.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. * $p < .1$, ** $p < .05$, *** $p < .01$.

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. Panel B plots the coefficient on the female dummy in equation (5) for different subgroups 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 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 résumé, could be offered different, lower-paying jobs. To disentangle these channels, I run the same regressions as in the first five columns of Table IV but add job fixed effects.

Table IV, column (6) 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 that on average, they are offered lower pay for the same job. Once we add résumé characteristics (column (7)), the bid gap narrows to 1.8%. Therefore, for a given job, gender differences in résumés can only explain part of the within-job bid gap. Adding résumé 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 Online Appendix Table A.6.

Résumé 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 résumé characteristics explain more than 80% of the total variation in the regressions without job fixed effects (Table IV, Column (2)), they can only explain 33%

of the total variation within jobs in [Table IV](#), column (7). In contrast, adding the ask salary increases the adjusted R^2 to 0.834 in column (8). 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 résumé characteristics.

V.C. Final Offers: Results

Given that bid salaries are nonbinding, 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 variable is now $\text{Log}(Final_{ib})$, the salary candidate i was offered for the job corresponding to bid b . The right-hand variables are the same as in [Table IV](#). 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 résumé information leaves a significant 1.4% gap (column (2)). After adding the ask salary to the résumé controls, as in column (4), 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).

V.D. Sensitivity Analysis

In [Table IV](#), the relationship between the ask and the bid is estimated on the full sample of bids sent out by companies. However, only a subsample 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 [Online Appendix Table A.7](#), I rerun the same regressions as in [Table IV](#) 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 IV](#).

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 [Online Appendix Table A.8](#), I rerun the regressions from [Table IV](#) on the subset of bids that are different from the ask, which represents 25% of the data. Although the results on that subsample

TABLE V
THE ROLE OF THE ASK AND RÉSUMÉ CHARACTERISTICS IN FINAL-OFFER GENDER DIFFERENCES

	No firm FE			Firm FE				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var.: Log final salary								
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 × log ask salary					0.011 (0.011)			0.008 (0.013)
Constant	11.606 ^{***} (0.026)	23.392 ^{***} (1.379)	11.608 ^{***} (0.010)	14.346 ^{***} (0.940)	14.351 ^{***} (0.940)	21.707 ^{***} (1.451)	14.372 ^{***} (1.020)	14.370 ^{***} (1.019)
Candidate's résumé characteristics								
Month × year FE	X	X	X	X	X	X	X	X
Firm FE								
Adj. <i>R</i> -squared	0.012	0.827	0.903	0.920	0.920	0.515	0.762	0.762
No. observations	7,582	7,582	7,582	7,582	7,582	6,303	6,303	6,303

Notes: This table presents estimates of β_1 from equations (4) to (7), except the left side is $\text{Log}(\text{Final}_i)$ —the final salary that candidate i was offered for the job corresponding to bid b_i —instead of $\text{Log}(\text{Bid}_{i,p})$. Accordingly, in columns (1) to (5), controls are the same as the corresponding columns in Table IV. The regressions are run on the sample of final offers. Hence, the number of observations is much smaller than in Table IV as the unit of observation is restricted to candidates with 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. * $p < .1$, ** $p < .05$, *** $p < .01$.

are qualitatively similar to [Table IV](#), the magnitudes vary in the direction predicted by the hypothesis. Indeed, the raw bid gap on that subsample 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 résumé explains more of the raw bid gap (59% versus 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. Forty-four percent of the interview requests also contain an equity offer. As evidenced in [Online Appendix Table A.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

V.E. Gender Differences at the Extensive Margin

1. *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 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' résumé 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.

16. In [Online 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. Given that the data only covers a few years, the vast majority of candidates who have repeated spells have them within the same year. 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 two and six weeks.

TABLE VI
GENDER DIFFERENCES IN THE NUMBER OF BIDS RECEIVED AND THE PROBABILITY OF RECEIVING A FINAL OFFER AFTER AN INTERVIEW

Dep. var.:	No. 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 salary ²				-0.228*** (0.014)			0.001 (0.001)	-0.002*** (0.001)
Female × 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 résumé characteristics		X	X	X	X	X	X	X
Month × year FE		X	X	X	X	X	X	X
Job FE								
Adj. <i>R</i> -squared	0.015	0.240	0.244	0.245	0.244	0.008	0.008	0.038
No. observations	164,799	164,799	164,799	164,799	164,799	261,518	261,518	251,817

Notes. 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 × year level, column (2) adds résumé characteristics as controls (as in Table IV, column (2)). Column (3) adds the ask salary in US\$100,000 to column (2). Column (4) adds the square of the ask salary in US\$100,000 to column (3). Column (5) adds an interaction between the female dummy and the ask salary to column (3). Standard errors (in parentheses) for columns (1) to (5) 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 led 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) for columns (6) to (8) are clustered at the candidate and job ID level. * $p < .1$, ** $p < .05$, *** $p < .01$.

Using a methodology developed in [Roussille and Scuderi \(2023\)](#) to rank firms, I show in [Online 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, [Table VI](#), column (3) shows that adding the ask salary to the controls does not affect 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 résumé, candidates who ask for more are, on average, facing higher demand. [Section VII.B](#) 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 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.

2. *Selection into the Final Offer Pool.* I 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 one 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.

V.F. From Descriptive to Causal Evidence

Introducing the ask salary as a control in [Table IV](#), 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 affect other coefficients? To answer this question, [Table III](#) 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. 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-sectional analysis of [Table IV](#): 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. Instead, 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. Since the bid and final salaries are highly correlated with the ask, part of the effect of controlling for the ask on résumé 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 résumé characteristics that I cannot fully account for with my résumé controls.

To make progress on the causal effect of the ask on the bid, I turn to analyzing a change on the platform that affected how 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 prefilled 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 affects 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

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 signaling different underlying preferences for nonwage amenities.

an exogenous shift in the ask salary of some of the candidates. Therefore, how this shift affects 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.

VI. CLOSING THE GENDER GAP

VI.A. 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 refer to as the ask salary. From the first year of data to mid-2018, the answer box for this question was an empty text entry. Starting in mid-2018, the answer box was prefilled 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 [Online Appendix Figure B.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 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 the threshold for computing the histogram and median, so I cannot construct a control group for whom the information wasn’t shown. 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 it. This sample contains more than 40,000 candidates and 200,000 bids.

It is worth highlighting that the reform was not anticipated by the candidates or the firms. Indeed, the company did not advertise the feature change externally, and new candidates were not drawn to the platform by it. In addition, the feature change only affected the candidates' experience on the platform; 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.

VI.B. *The Impact of the Reform on the Ask Salary*

1. *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 designs by estimating:

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

where $t = t(i)$ is the month in which candidate i created her profile, After_t is a dummy equal to one after the reform, Female_i is equal to one 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 \(2\)](#) in the pre-period to predict the ask salary of every candidate, controlling for all their résumé characteristics.¹⁹ I run this predicted ask against an interacted model of female and after dummies. Results are presented in [Online Appendix Table A.9](#): the coefficient on the interaction between female and after is exactly zero.

19. Except that instead of month \times year FE, there are just month FE (1–12) and a monthly linear time trend.

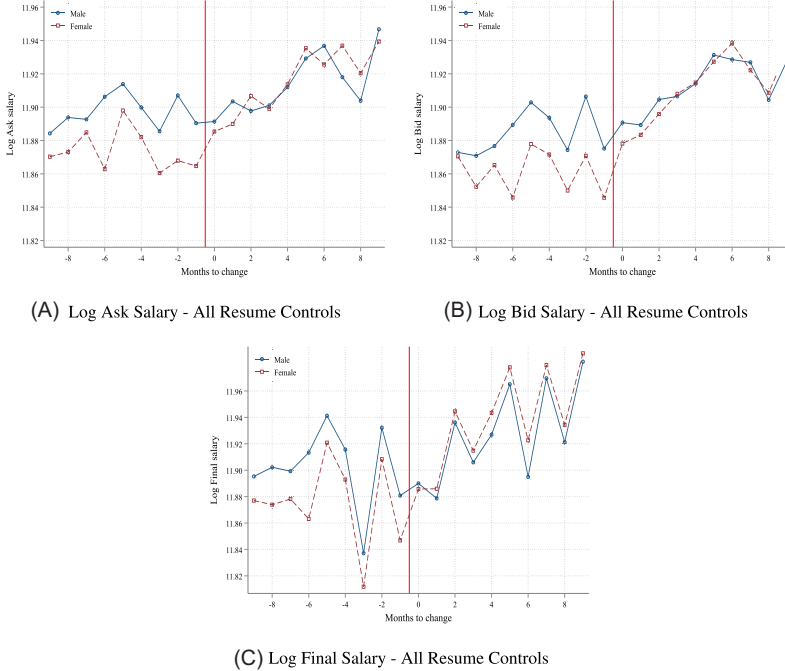


FIGURE IV

Effect of the Reform on the Gender Ask, Bid, and Final Gaps

These figures plot the time series of annual mean salary for men and women, net of all résumé characteristics. Each panel is constructed regressing the outcome variable (either log ask salary for Panel A, log bid salary for Panel B, or log final salary for Panel C) within every month on a Female indicator and the résumé 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).

In other words, the predicted ask gap is stable across periods. [Online Appendix Table A.10](#) also provides summary statistics on candidates' résumé characteristics before and after the change, illustrating the absence of differential selection of men and women onto the platform after the reform.

2. *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' résumé characteristics.

I 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 observe a 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.

3. *Regression Results.* Table VII, columns (1) and (2) formalize 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 in these graphs that the cumulative distribution function 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, that is, women have on average about two years less of experience compared with 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.

4. *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 plot the cumulative distribution

TABLE VII
THE EFFECT OF THE REFORM ON THE GENDER GAP

Dep. var.:	Log ask salary		Log bid salary		Log Final Salary	
	(1)	(2)	(3)	(4)	(5)	(6)
Female × 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 résumé characteristics	X	X	X	X	X	X
Job FE						
Adj. <i>R</i> -squared	0.517	0.493	0.514	0.865	0.252	0.519
No. observations	43,368	207,636	207,636	207,636	200,593	2,476

Notes. This table presents estimates of β_0 , β_1 , and β_2 from equations (8) and (9) and the corresponding coefficients for the log final salary controlling for all candidates' résumé 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. Robust standard errors (in parentheses) are used in columns (1). In column (2), standard errors (in parentheses) are clustered at the candidate level. From column (3) onward, standard errors (in parentheses) are two-way clustered at the candidate and job ID level. *, **, *** $p < .10$, $p < .05$, $p < .01$.

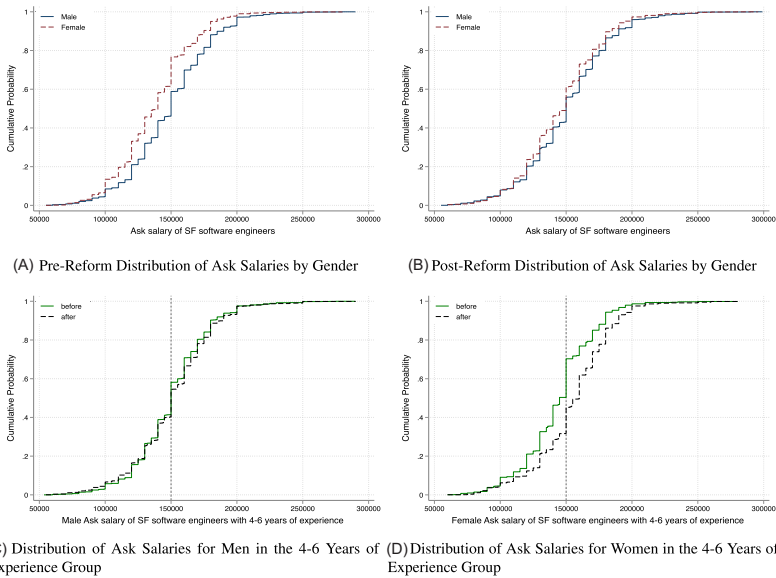


FIGURE V

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

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. Panels C and D plot the cumulative density of ask salaries, separately, for male and female, respectively, before (solid 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 gray 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.

function of ask salaries for the four to six years experience group for men (Panel C) and women (Panel D), separately before the reform (solid lines) and after (dashed lines).²⁰ All candidates in these panels saw the same median, which is illustrated by the short-dashed gray 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

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

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 VII.A](#) explores the potential mechanisms behind these outcomes.

VI.C. The Impact of the Reform on the Bid Salary

1. *Empirical Strategy.* I investigate the effect of the reform on the bid salaries sent by firms in [equation \(9\)](#):

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

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.

2. *Results.* [Table VII](#), columns (3) to (5) formalize 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)).

3. *Heterogeneous Effects of the Reform.* [Figure VI](#) plots the effect of the reform on the ask (blue squares) and bid (red dots) gaps as a function of the pre-reform gaps, separately by

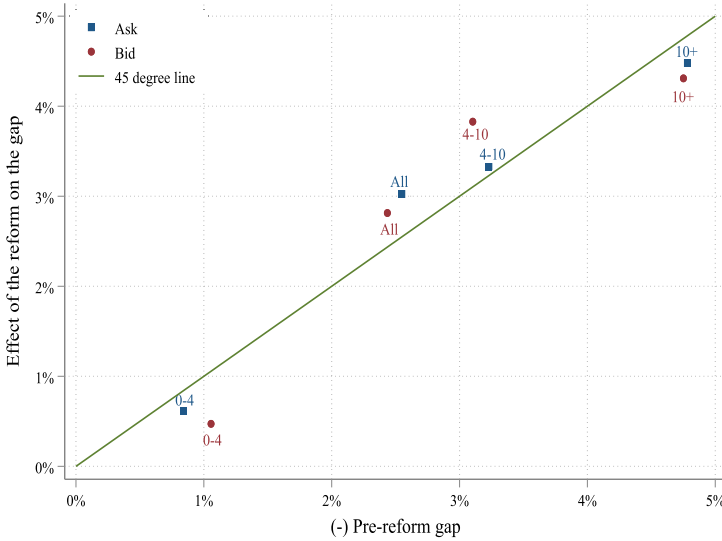


FIGURE VI

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

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.

experience groups. In line with the results in Section IV, 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.

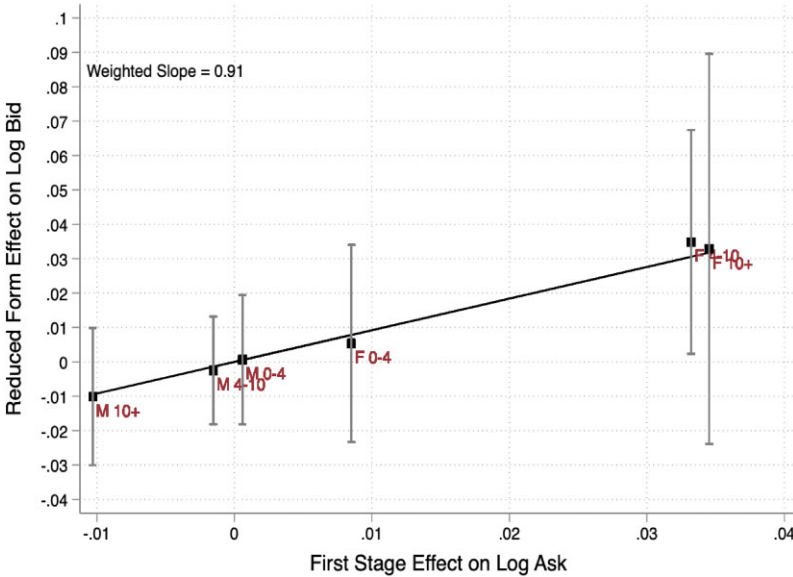


FIGURE VII

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 résumé characteristics. As originally described in [Holzer, Katz, and Krueger \(1991\)](#) and recently in [Angrist, Autor, and Pallais \(2022\)](#), the slope of the line of best fit in 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.

4. *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 résumé characteristics. As originally described in [Holzer, Katz, and Krueger \(1991\)](#) and recently in [Angrist, Autor, and Pallais \(2022\)](#), the slope of the line of best fit in 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 résumé characteristics (see Table IV, column (4)). This suggests that there was indeed little room for omitted variable bias in the OLS regression, as argued in Section V.B. In terms of generalizability 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 those of men on average. Firms' responses 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

5. *Other Variables.* The median salary shown to candidates accounts for the candidate's experience but not for their other résumé characteristics (e.g., their education). Therefore, candidates with different education levels but the same experience see the same suggestion. As a consequence, the reform could have affected the role of other controls in the determination of the ask salary. Online Appendix Table A.11 reports the results of a regression of the log ask salary on all the résumé 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 coefficient on 10–15 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

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 analyzed in Online Appendix F.

levels or schools were exposed to the same median and therefore converged in their ask.

6. *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 , the likelihood of getting a final offer, and the rank of the firms that bid and make offers to the candidate (see [Online Appendix G](#) for more details on how these ranks are computed).²³ [Table VIII](#) presents the results of this analysis. First, 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). The coefficient on the interacted term $\text{female} \times \text{after}$ is 0.19 (95% confidence interval -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 $\text{female} \times \text{after}$ is very small (95% confidence interval -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 $\text{female} \times \text{after}$ is close to zero and insignificant (confidence interval -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 (confidence interval -0.3 to 0.9 , mean = 62.5) or make a final offer to them (confidence interval -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 effect on several extensive-margin variables.

23. The specification is the same as in [equation \(8\)](#) except the left side becomes $Nb_{bids_{ik}}$ and $Hours_{ik}$, as defined in [Section V.E](#), and I add the length of the candidate's spell (two to six weeks) to the controls.

TABLE VIII
THE EFFECT OF THE REFORM AT THE EXTENSIVE MARGIN

Dep. var.:	No. bids (1)	No. hours (2)	Final offer (3)	Firm rank (bid) (4)	Firm rank (final) (5)	No. bids (6)	Pred. no. bids (7)	Pred. no. bids (8)
Female × 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 salary ²						-0.000 (0.000)		
Poisson AME on female × after	0.222	-1.738	0.009					
Mean dep. var.	4.79	62.33	0.09	62.5	62.9			
Candidate's résumé 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
No. observations	43,368	32,043	43,368	188,463	2,074	43,368	43,368	43,368

Notes: 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 résumé 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 sent on the same variables. The dependent variable in column (4) is the firm rank at the bid level as in [Online Appendix Table G.1, column \(1\)](#), and the dependent variable in column (5) is the firm rank at the final offer level as in [Online Appendix 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)'s specification on the pre-period. The dependent variable in column (8) is the predicted number of bids received using column (6)'s specification on the pre-period. Observations are at the spell level. Standard errors (in parentheses) are clustered at the candidate ID level in all columns but columns (4) and (5), where they are two-way clustered at the candidate and job ID level. * $p < .1$, ** $p < .05$, *** $p < .01$.

VII. 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 with 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?

VII.A. 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 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 résumé, 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 [Online Appendix Table A](#)) also casts doubt on a story where gender gaps in tastes for nonwage 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: (i) downward-biased beliefs about the market wage for their résumé, and (ii) anticipated gender discrimination, which would lead to lower asks to mitigate it. While I

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

do not have definitive evidence to adjudicate between the reform having a (i) pure information channel versus (ii) 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, makes the norm-setting power of *Hired.com* an unlikely explanation. If we thought women used the suggested ask as a signal for an “appropriate” ask, 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 which the treatment corrected. Salary information, however, does not shift beliefs about the position in the quality distribution, 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 [Online Appendix Table A.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 with 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 Säve-Söderbergh \(2022\)](#) run a survey on a representative sample of recent graduates in Sweden to shed light on the mechanism behind women’s lower ask. The article finds

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.

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.

VII.B. *The Ask Salary as a Signal of Quality*

A second question that the reform effects raise is the following: why are firms not decreasing their demand for female labor, compared with 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.

I first investigate, descriptively, the relationship between the number of bids received and the residual ask for all candidates. [Online Appendix](#) Figure B.6 documents a bell-shaped 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 résumé 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.

26. Using the methodology developed in [Roussille and Scuderi \(2023\)](#) to calculate firms' productivity, [Online Appendix](#) 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.

The idea of price as a signal of quality, while understudied in the context of wage bargaining, has been theorized for consumer products in the fields of industrial organization 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 [Online Appendix H](#), adapting [Wolinsky's \(1983\)](#) 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 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-ask combination that maximizes 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 [Online Appendix H.5](#), we can approximate a given firm type by estimating the range of residual asks in which it interviews. [Online Appendix 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 their expected quality. But in equilibrium, neither men nor women end up raising this flag because candidates of any gender either are of the quality they signal (men) or above (women).

We can now return to our initial question: why is it that firms are not decreasing their demand for female labor, compared with 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 [Online Appendix Figure B.6](#). [Table VIII](#), columns (6) to (8) 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 résumé. 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) and that of column (6) is consistent with the statistical relationship between the number of bids received and the ask salary being structural. Furthermore, this positive coefficient indicates that the women whose ask was shifted up by the reform are in the increasing region of [Online Appendix Figure B.6](#), which explains why they do not face a penalty for asking for more.

VIII. CONCLUSION

This article 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 3.3% 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' résumé characteristics only narrows the bid gap by 33%. These results qualitatively carry through to the 7,582 final salary offers for the subsample of hired candidates. On this platform, women are not discriminated against at the extensive margin. In particular, conditional on their résumé 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 prefilled 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.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY, UNITED STATES

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

DATA AVAILABILITY

The code underlying this article is available in the Harvard Dataverse, <https://doi.org/10.7910/DVN/HTATUN> (Roussille 2024).

27. The unemployment rate of U.S. tech workers had hit a record low in the study period (Warner 2019).

REFERENCES

- Agan, Amanda, Bo Cowgill, and Laura Katherine Gee, "Do Workers Comply with Salary History Bans? A Survey on Voluntary Disclosure, Adverse Selection, and Unraveling," *American Economic Association Papers and Proceedings*, 110 (2020), 215–219. <https://doi.org/10.1257/pandp.20201123>
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber, "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy*, 113 (2005), 151–184. <https://doi.org/10.1086/426036>
- Angrist, Joshua, David Autor, and Amanda Pallais, "Marginal Effects of Merit Aid for Low-Income Students," *Quarterly Journal of Economics*, 137 (2022), 1039–1090. <https://doi.org/10.1093/qje/qjab050>
- Babcock, Linda, and Sara Laschever, *Why Women Don't Ask: The High Cost of Avoiding Negotiation and Positive Strategies for Change* (New York: Bantam, 2006).
- Babcock, Linda, Sara Laschever, Michele Gelfand, and Deborah Small, "Nice Girls Don't Ask," *Harvard Business Review*, October 2003.
- Baker, Michael, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar, "Pay Transparency and the Gender Gap," *American Economic Journal: Applied Economics*, 15 (2023), 157–183. <https://doi.org/10.1257/app.20210141>
- Barach, Moshe A., and John J. Horton, "How Do Employers Use Compensation History? Evidence from a Field Experiment," *Journal of Labor Economics*, 39 (2021), 193–218. <https://doi.org/10.1086/709277>
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen, "Inference on Treatment Effects after Selection among High-Dimensional Controls," *Review of Economic Studies*, 81 (2014), 608–650. <https://doi.org/10.1093/restud/rdt044>
- Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon, "Do Firms Respond to Gender Pay Gap Transparency?," *Journal of Finance*, 77 (2022), 2051–2091. <https://doi.org/10.1111/jofi.13136>
- Bertrand, Marianne, "The Glass Ceiling," Coase Lecture, London School of Economics, March 2017. <https://doi.org/10.2139/ssrn.3191467>
- Biasi, Barbara, and Heather Sarsons, "Flexible Wages, Bargaining, and the Gender Gap," *Quarterly Journal of Economics*, 137 (2022), 215–266. <https://doi.org/10.1093/qje/qjab026>
- Blau, Francine D., and Lawrence M. Kahn, "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 55 (2017), 789–865. <https://doi.org/10.1257/jel.20160995>
- Bowles, Hannah Riley, Linda Babcock, and Kathleen L. McGinn, "Constraints and Triggers: Situational Mechanics of Gender in Negotiation," *Journal of Personality and Social Psychology*, 89 (2005), 951–965. <https://doi.org/10.1037/0022-3514.89.6.951>
- Chamberlain, Andrew, and Jyotsna Jayaraman, "The Pipeline Problem: How College Majors Contribute to the Gender Pay Gap," Glassdoor Economic Research, 2017.
- Chernozhukov, Victor, Iván Fernández-Val, and Ye Luo, "The Sorted Effects Method: Discovering Heterogeneous Effects Beyond Their Averages," *Econometrica*, 86 (2018), 1911–1938. <https://doi.org/10.3982/ECTA14415>
- Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan, "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR," *Quarterly Journal of Economics*, 126 (2011), 1593–1660. <https://doi.org/10.1093/qje/qjr041>
- Coffman, Katherine Baldiga, "Evidence on Self-Stereotyping and the Contribution of Ideas," *Quarterly Journal of Economics*, 129 (2014), 1625–1660. <https://doi.org/10.1093/qje/qju023>

- Cortés, Patricia, Jessica Pan, Laura Pilossoph, Ernesto Reuben, and Basit Zafar, "Gender Differences in Job Search and the Earnings Gap: Evidence from the Field and Lab," *Quarterly Journal of Economics*, 138 (2023), 2069–2126. <https://doi.org/10.1093/qje/qjad017>
- Crosan, Rachel, and Uri Gneezy, "Gender Differences in Preferences," *Journal of Economic Literature*, 47 (2009), 448–474. <https://doi.org/10.1257/jel.47.2.448>
- Cullen, Zoë B., and Bobak Pakzad-Hurson, "Equilibrium Effects of Pay Transparency," *Econometrica*, 91 (2023), 765–802. <https://doi.org/10.3982/ECTA19788>
- Dreber, Anna, Emma Heikensten, and Jenny Sävje-Söderbergh, "Why Do Women Ask for Less?," *Labour Economics*, 78 (2022), 102204. <https://doi.org/10.1016/j.labeco.2022.102204>
- Exley, Christine L., and Judd B. Kessler, "The Gender Gap in Self-Promotion," *Quarterly Journal of Economics*, 137 (2022), 1345–1381. <https://doi.org/10.1093/qje/qjac003>
- Exley, Christine L., Muriel Niederle, and Lise Vesterlund, "Knowing When to Ask: The Cost of Leaning In," *Journal of Political Economy*, 128 (2020), 816–854. <https://doi.org/10.1086/704616>
- Garbinti, Bertrand, Jonathan Goupille-Lebret, and Thomas Piketty, "Income inequality in France, 1900–2014: Evidence from Distributional National Accounts (DINA)," *Journal of Public Economics*, 162 (2018), 63–77. <https://doi.org/10.1016/j.jpubeco.2018.01.012>
- Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 104 (2014), 1091–1119. <https://doi.org/10.1257/aer.104.4.1091>
- Golub, Benjamin, and Matthew O. Jackson, "How Homophily Affects the Speed of Learning and Best-Response Dynamics," *Quarterly Journal of Economics*, 127 (2012), 1287–1338. <https://doi.org/10.1093/qje/qjs021>
- Gomez-Herrera, Estrella, and Frank Mueller-Langer, "Is There a Gender Wage Gap in Online Labor Markets? Evidence from over 250,000 Projects and 2.5 Million Wage Bill Proposals," Max Planck Institute for Innovation & Competition Research Paper No. 19-07, 2019. <http://dx.doi.org/10.2139/ssrn.3427012>
- Holzer, Harry J., Lawrence F. Katz, and Alan B. Krueger, "Job Queues and Wages," *Quarterly Journal of Economics*, 106 (1991), 739–768. <https://doi.org/10.2307/2937925>
- Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer, "Worker Beliefs about Outside Options," *Quarterly Journal of Economics*, 139 (2024). <https://doi.org/10.1093/qje/qjae001>
- Kiessling, Lukas, Pia Pinger, Philipp Seegers, and Jan Bergerhoff, "Gender Differences in Wage Expectations and Negotiation," *Labour Economics*, 87 (2024), 102505. <https://doi.org/10.1016/j.labeco.2024.102505>
- Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet, "Gender Differences in Job Search: Trading Off Commute against Wage," *Quarterly Journal of Economics*, 136 (2021), 381–426. <https://doi.org/10.1093/qje/qjaa033>
- Leibbrandt, Andreas, and John A. List, "Do Women Avoid Salary Negotiations? Evidence from a Large-Scale Natural Field Experiment," *Management Science*, 61 (2015), 2016–2024. <https://doi.org/10.1287/mnsc.2014.1994>
- Litman, Leib, Jonathan Robinson, Zohn Rosen, Cheskie Rosenzweig, Joshua Waxman, and Lisa M. Bates, "The Persistence of Pay Inequality: The Gender Pay Gap in an Anonymous Online Labor Market," *PLoS One*, 15 (2020). <https://doi.org/10.1371/journal.pone.0229383>
- McPherson, Miller, Lynn Smith-Lovin, and James M. Cook, "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology*, 27 (2001), 415–444. <https://doi.org/10.1146/annurev.soc.27.1.415>
- Milgrom, Paul, and John Roberts, "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 94 (1986), 796–821. <https://doi.org/10.1086/261408>
- Murciano-Goroff, Raviv, "Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers," *Management Science*, 68 (2022), 3262–3281. <https://doi.org/10.1287/mnsc.2021.4077>

- Olivetti, Claudia, and Barbara Petrongolo, "The Evolution of Gender Gaps in Industrialized Countries," *Annual Review of Economics*, 8 (2016), 405–434. <https://doi.org/10.1146/annurev-economics-080614-115329>
- Reuben, Ernesto, Matthew Wiswall, and Basit Zafar, "Preferences and Biases in Educational Choices and Labour Market Expectations: Shrinking the Black Box of Gender," *Economic Journal*, 127 (2017), 2153–2186. <https://doi.org/10.1111/eoj.12350>
- Rigdon, Mary L., "An Experimental Investigation of Gender Differences in Wage Negotiations," SSRN Working Paper, 2012. <https://doi.org/10.2139/ssrn.2165253>
- Roussille, Nina, "Replication Data for: 'The Role of the Ask Gap in Gender Pay Inequality,'" (2024), Harvard Dataverse, <https://doi.org/10.7910/DVN/HTATUN>.
- Roussille, Nina, and Benjamin Scuderi, "Bidding for Talent: A Test of Conduct on a High-Wage Labor Market," Working Paper, 2023. <https://doi.org/10.2139/ssrn.4536978>
- Small, Deborah, Michele Gelfand, Linda Babcock, and Hilary Gettman, "Who Goes to the Bargaining Table? The Influence of Gender and Framing on the Initiation of Negotiation," *Journal of Personality and Social Psychology*, 93 (2007), 600–613. <https://doi.org/10.1037/0022-3514.93.4.600>
- Visier and Insights, "The Truth about Ageism in the Tech Industry," Report, 2017.
- Warner, Bernhard, "The Unemployment Rate for U.S. Tech Workers Just Hit the Lowest Number Ever Recorded," *Fortune*, 2019. <https://fortune.com/2019/06/12/unemployment-tech-workers-lowest-ever/>
- Wolinsky, Asher, "Prices as Signals of Product Quality," *Review of Economic Studies*, 50 (1983), 647–658. <https://doi.org/10.2307/2297767>
- Yagan, Danny, "Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut," *American Economic Review*, 105 (2015), 3531–3563. <https://doi.org/10.1257/aer.20130098>