Skills, Job Application Behavior and the Gender Wage Gap: Evidence from Online Freelancing

Ole Teutloff* Eliza Stenzhorn† Otto Kässi‡

December 2023

Abstract

This paper investigates how workers' skills and job application behavior contribute to the gender wage gap using data from a leading online labor platform. We utilize machine learning models to quantify the value of workers' skills and estimate their impact on wages. We find a substantial raw gender wage gap of over 30%. However, the gender wage gap can be fully accounted for by three factors: differences in workers' skills, differences in the projects they apply to, and differences in asking wages. Our findings suggest no employer discrimination based on gender. Instead, the gender wage gap emerges because men and women seem to use the platform in different ways. Women, on average, prioritize consistent income, while men pursue higher-paying, occasional gigs. These differences likely stem from different constraints and labor market opportunities outside the platform. According to our results, the flexibility of the online gig economy is unlikely to favor women.

JEL Classification: J16, J24, J31

Keywords: gender wage gap, gig economy, skills, human capital, flexibility, job application behavior, online labor markets, random forest regression

* Copenhagen Center for Social Data Science, University of Copenhagen, Copenhagen, Denmark. E-mail: ole.teutloff@sodas.ku.dk.

† ZEW Mannheim, Germany. E-mail: eliza.stenzhorn@zew.de.

‡ Oxford Internet Institute, University of Oxford, Oxford, United Kingdom, Etla Economic Research Helsinki, Helsinki, Finland. E-mail: otto.kassi@etla.fi.

1 Introduction

In most countries, women still earn less than men. In the United States, full-time employed females earn roughly 80% of what is earned by full-time employed males (Goldin, 2021). In the European Union, the gender wage gap ranges from 99.3% to 77.7% (Eurostat, 2022). This gap persists even though women tend to acquire similar or higher levels of education than men in almost all high- and middle-income countries (Schofer and Meyer, 2005; Van Bavel et al., 2018).

Much of the observed gender wage gap can be attributed to parenthood, which leads women to prioritize career choices that offer flexibility (Adda et al., 2017; Angelov et al., 2016; Bertrand et al., 2010; Blau and Kahn, 2013; Kleven et al., 2019b). However, the preference for job flexibility comes at a significant cost to women. Many professions provide greater rewards to employees who can commit to extended and unpredictable hours. Goldin (2014) characterizes this as a convex relationship between working hours and wages: when workers are not easily interchangeable, those who can work extended and specific hours command a premium. This, in turn, results in a wage penalty for women. Interestingly, Goldin (1990) highlights that with 'pay-by-piece' payment systems prevalent in late 19$^{\text{th}}$-century manufacturing roles, there was minimal wage disparity between men and women performing identical jobs. In such payment systems, wages were primarily tied to individual productivity — which could easily be observed by the employer — rather than the number of hours worked. The shift to longer-term contracts altered this dynamic.

This paper studies the gender wage gap using data from an online freelancing labor platform, which employs a 'pay-by-piece' model, not unlike manufacturing over 150 years ago. Online labor platforms allow employers to divide work into distinct tasks, facilitating both precise oversight of worker outputs and a high degree of interchangeability between workers (Cook et al., 2021). In addition, work arrangements on the platform are completely remote. In principle, platform-mediated work is highly flexible: workers can choose when, where, and to what extent they wish to work. We demonstrate in this paper, that this flexibility does not result in the reduction of the gender wage gap. If anything, conditional on education, the gap in hourly wages is much higher than in traditional labor markets.

Our findings are consistent with prior research examining various facets of digitally mediated gig work. Specifically, previous research on ridesharing (Cook et al., 2021) and online clickwork (Adams-Prassl et al., 2023) reports qualitatively similar findings. Both ridesharing and clickwork are specific labor markets dominated by highly standardized tasks. These markets offer limited opportunities for skill differentiation, and wage negotiations are nonexistent. In contrast, our data originates from an online freelancing platform that facilitates transactions involving a diverse range of high-skill tasks and where wage negotiations play a central role.[1] In this paper, we quantify how gender differences in skills, application behavior, and wage-setting strategies shape the gender wage gap. Accounting for these three factors enables us to explain why the gender wage gap persists or even widens in platform-mediated work.

Our data provides comprehensive information on job requirements, including desired worker experience, expected project duration, weekly hours, and skill prerequisites. It also captures worker characteristics such as formal education and specific skills. Both job-level and worker-level skill requirements come from the same standardized skill taxonomy consisting of more than 4,000 skill tags. In addition, we observe the application behavior of workers, capturing details such as asking wages and applied-for projects as well as the subsequent compensation for the work completed. Traditional labor market data typically lacks information on both skills and application behavior. In contrast, our data allow for an in-depth analysis of how men and women differ in what skills they possess, what jobs they apply for, and what they are paid for their work.

Our analysis is structured in three main steps. First, we quantify the raw gender wage gap, adjusting for typical background factors like education. In the next step, we use a random forest machine learning model (Breiman, 2001) to transform the extensive binary skill tags of workers into a single dimension that represents the value of their skill sets. We train this model to capture the relationship between project skill tags and their associated hourly wages. By applying this model to workers' self-listed skill tags, we can estimate the expected hourly wage of a worker based on their specific skills, helping us see if men and women have skills with different market values. Finally, we examine workers' application behaviors to understand how men and women choose jobs with different attributes. We consider several job amenities: expected project duration, workload, type of contract, and the experience level desired by the employer. Using our earlier model, we estimate the potential hourly wages of each project based on its skill requirements. This approach allows us to measure if women choose jobs with different characteristics and expected earnings compared to men, and how these choices impact the gender wage gap.

We find a raw gender wage gap in online freelancing, with women earning on average 33.6 log-points ($\approx 30\%$) lower hourly wages than men.[2] Our findings underscore that women, on average, have skills of lower market value than men, significantly contributing to the gender wage gap. By factoring in skills, the unexplained wage gap narrows to 12.3%. Accounting for the application behavior of workers explains the remaining wage

[1] Prior literature on gender wage gaps in online freelancing in particular, has concentrated either on asking wages (Foong et al., 2018) or on asking wages and job categories jointly (Gomez-Herrera and Müller-Langer, 2019). Our contribution is that we jointly account for skills, types of applied-for jobs, and asking wages.

[2] Our preferred measure of wages is the log-hourly wage. The gender gap is calculated as the difference in log-hourly wages, expressed in log-points.

disparity. In other words, while we find a considerable wage gap among platform workers, we can uncover that workers with similar skills, who apply for the same jobs, and who ask for the same wages earn, on average, similar wages regardless of gender. Consequently, we find no evidence of gender discrimination on the platform.

A natural follow-up question then is *why* do men and women apply for different jobs? Our data is less well suited to answer this question, but we probe for some potential mechanisms behind the differences in application behavior. Despite earning less, women tend to work on longer projects and have a higher likelihood of being hired compared to men. Within the limitations of our data, we find no evidence of compensatory behavior, where women might intentionally opt for lower-paying jobs to increase their hiring chances. We also find evidence against the hypothesis that differences in wages stem from disparities in return to experience. We generally find that women are over-represented among the more experienced workers and that the return to work experience is broadly similar between men and women. Instead, we conclude that the gender wage gap arises because men and women engage with the platform differently. Our results suggest that women may prioritize steady income or are less inclined to take risks, while men aim for higher-paying occasional jobs. While our data does not enable causal analysis, our evidence suggests that external factors, such as off-platform opportunities and personal preferences, are likely to drive the observed gender wage gap in online freelancing.

Our findings are in line with previous literature on how gender differences in constraints and preferences shape the wage gap.[3] Goldin (2014) suggests that women exhibit different preferences with regard to workplace arrangements because of motherhood and, as a result, suffer a wage penalty. Numerous studies find evidence in favor of this hypothesis (see, e.g., Azmat and Ferrer, 2017; Barth et al., 2021; Bertrand et al., 2010; Gallen, 2018; Goldin and Katz, 2016). To accommodate children and other care work, women trade off earnings for non-monetary job amenities such as increased flexibility and "schedule controllability" (Bolotnyy and Emanuel, 2022). We contribute to this stream of literature by showing that, even in the highly unregulated and flexible labor market of online freelancing, women appear to favor job amenities like predictability, even if it comes at the expense of hourly earnings.

Another stream of research emphasizes the role of education, occupational sorting, and application behavior. Blau and Kahn (2017) underscore that even with similar levels of education, occupational sorting remains key for explaining gender wage differences. Women are still underrepresented in high-paying occupations and industries in the traditional labor market. While occupation differences remain central, differences in education levels have declined in importance. Women tend to acquire similar or higher levels of education than men in almost all high- and middle-income countries (Schofer and Meyer,

[3]For an overview of recent literature on gender differences in wages, we refer the reader to Blau and Kahn (2017) and Olivetti and Petrongolo (2016).

2005; Van Bavel et al., 2018). As a result, human capital differences have become less important in explaining the gender wage gap (Blau and Kahn, 2017). Nevertheless, women continue to choose different types of education and, therefore, acquire different hard skills. Gendered stereotypes, cultural norms, and a lack of role models result in women being underrepresented in science, technology, engineering, and mathematics (Kahn and Ginther, 2017; Card and Payne, 2021). This is in line with our findings. In the context of the gig economy, women also tend to sort into lower-paying projects, even conditional on their skills. While this is not the central focus of our investigation, we find that online freelance women are underrepresented in technical fields such as IT, data science, and engineering while being over-represented in writing and translation jobs.

Besides skills, gender differences in job search and application behavior contribute to the gender wage gap. Women work in different occupations, earning different wages, because they choose to apply for different jobs (Fluchtmann et al., 2021; Le Barbanchon et al., 2021). Using data on Danish unemployment insurance recipients, Fluchtmann et al. (2021) show that conditional on individual-level observable characteristics, women apply for jobs with 4.5 percent lower wages than men. In their analysis, differences in applied-for jobs explain a large share of the residual gender wage gap in wages. Our data corroborates these findings in the context of online freelancing. We find that women apply for jobs with different amenities. These differences in application behavior account for a substantial share of the gender wage gap. Moreover, female workers ask for lower wages reinforcing the findings in Roussille (2021) on the importance of asking wages for salary outcomes.

Additionally, our work links to a number of recent research articles focusing specifically on gender wage gaps in the gig economy. In low-skill location-based gig work involving tasks like delivery, shopping and carpentry, Cullen et al. (2018) find strong sorting by gender with women doing jobs that pay less and are associated with traditional female work. Even within the same job category, women do the lower-paying jobs. Cullen et al. (2018) argue that men can be more selective about which jobs to accept because they have better outside options. Their findings highlight the value of our method, which unpacks job postings into combinations of skills to gain a better understanding of gender-related wage disparities. In the context of ridesharing, a different form of location-based gig-work, Cook et al. (2021) estimate a gender wage gap of 7% amongst Uber drivers in the U.S.. Similar to our paper, they show that skills learned on the job — proxied by past work experience on the platform — and preferences for certain types of rides — driving speed and pickup areas — account for the entire observed gender wage gap.

Adams-Prassl et al. (2023) concentrates on a remote clickwork platform, Amazon Mechanical Turk. Tasks on this platform remain fairly standardized and often take only a few minutes to complete. On Mechanical Turk, workers can choose which task they work on from a list of available tasks, leaving little room for employer discrimination. According to her results, despite having, on average, the same platform experience and selecting similar tasks, women earn 20% less per hour than men. Via a survey, Adams-Prassl et al. (2023) demonstrated that the wage gap concentrates among women with children who report that domestic duties adversely affect their ability to plan and complete work on Amazon Mechanical Turk. In distinction to Adams-Prassl et al. (2023), we show how differences in skills and differences in applied-for jobs affect the gender wage gap.

In contrast to clickwork, jobs on online freelancing platforms are longer, more diverse, and complex and generally require higher skills. To our knowledge, Gomez-Herrera and Müller-Langer (2019) and Foong et al. (2018) are the only two studies focusing on gender wage gaps in online freelancing. Using a global dataset, Gomez-Herrera and Müller-Langer (2019) identify a 4% gender wage gap that can be entirely attributed to the bidding behavior of workers. According to their results, women tend to bid on projects with smaller declared budgets. Moreover, conditional on bidding on a project, women ask for lower wages. Gomez-Herrera and Müller-Langer (2019) additionally document that female freelancers have a higher probability of winning projects and seem to make up for their lower pay per project by completing more projects. However, their research design remains silent on whether the lower declared budgets and asking wages are due to differences in skills, negotiation strategies, or other factors. In contrast, our approach allows us to disentangle how workers' skills and application behavior contribute to the observed gender wage gap. Foong et al. (2018), on the other hand, focuses on how asking wages shape the gender wage gap in online freelancing. Yet, besides asking wages, the analysis in Foong et al. (2018) does not include other dimensions of application behavior. The major contribution of our study is to combine granular information on skills with detailed information on application behavior and important worker- and job-level background characteristics.

While our data is from a rather specific setting, our contribution is more general. The approach we take could readily be implemented outside of the context of digital labor platforms if granular data on skills and application behavior are available, for example, in human resources departments of large corporations or employment offices.

The paper is structured as follows: Section 2 starts with a description of the online freelancing market and provides details about our data set. Then we present descriptive statistics and quantify the raw gender wage gap in Section 3. In Section 4, we disentangle how workers' skills and job-seeking patterns contribute to the gender wage gap. In Section 4.3, we employ the Gelbach (2016) decomposition method to determine the extent to which each factor included in our analysis affects the gender wage gap. We discuss our findings and potential mechanisms driving our results in Section 5. Section 6 concludes the paper with a discussion of the implications and limitations of our work.

2 Background and Data

2.1 Online Freelancing Platforms

Online labor platforms are digital marketplaces connecting buyers and sellers of remotely deliverable work. These platforms can be subdivided into microtask platforms, such as *Amazon Mechanical Turk*, where tasks are split into small pieces and freelancing platforms, such as *Upwork*, *Fiverr*, or *Freelancer* which host bigger and more complex projects. We use data from one prominent online freelancing platform based in the United States, which wished to remain anonymous. This platform hosts millions of workers who bid on thousands of new projects posted daily by employers.[4]

Employers range from individuals and startups to Fortune 500 companies (Corporaal and Lehdonvirta, 2017). Workers are decentralized individuals worldwide who transact their work digitally over the Internet. In this way, online freelancing differs from location-based gig work (such as ride-sharing or food delivery). However, just like in location-based gig work, online workers operate as independent contractors. As a result, they have (at least in theory[5]) full flexibility concerning when, how much, and for whom to work. The workers act as independent contractors without a formal employment relationship. This implies that standard labor market regulations on working hours, minimum wages, or equal pay legislation do not apply. Moreover, since the workers are self-employed, they are not entitled to employer-paid family leave.

Projects on the platform in question span a wide range of activities including data entry and administrative support, design, writing and translation, marketing, accounting, human resources, software development, and legal counseling. The employer initiates the hiring process by posting a vacancy on the platform, which includes a description of the job, the expected duration of the contract, preferred worker characteristics (such as experience and time commitment), the contract type (fixed sum or hourly pay rate), and project-specific skill requirements. When creating a project, employers choose from a dictionary of approximately 4,000 skills to define the skill requirements of their job posting. Workers select skills from the same dictionary and display them on their personal profiles to showcase their expertise.

Our data contains information on completed project transactions and worker profiles. The project-level data includes information on who applied to a project, who was selected, and how much was paid for the work. Such granular data on the demand and supply side of skills combined with price information as well as application behavior is a major advantage for studying gender wage disparities. Our data allows for a fine-grained analysis of how men and women differ in what skills they possess, which projects they apply

[4]For further details, see Ghani et al. (2014); Kässi and Lehdonvirta (2022); Lehdonvirta et al. (2019).

[5]It is beyond the scope of this paper to whether online freelancers should be classified as independent contractors or not, or to assess whether their actual flexibility is hindered by competition from other workers.

for, and what they are paid for their work, conditional on project- and worker-level control variables, such as project duration, contract type, or worker education. A possible downside of our data is that we do not observe the family arrangement (such as the marital status and number of children) of the workers in our data.

2.2 Collecting Online Freelancing Data

The data was collected as part of the Online Labor Index project (Kässi and Lehdonvirta, 2018), which tracks daily new project postings via the platform's API since January 2017. Subsequently, we collected the project details in several waves between November 2019 and October 2022. Most of the completed project observations also contain information on applicants. We use the list of applicants to collect worker profiles and match them to the projects they completed. The worker profiles also include information on workers' work histories beyond 2017. Collecting information on these projects allows us to extend our data set to transactions to years prior to 2017.

The online freelancing platform is U.S.-based but global in scope. Most workers are based outside of the United States, most notably in India, Pakistan, the Philippines, and Eastern Europe (Stephany et al., 2021). However, to minimize unobserved heterogeneity that might affect the gender wage gap and to ensure an apples-to-apples comparison, we restrict our analysis to workers based in the United States.[6].

When posting a project, employers specify the requirements, characteristics and amenities of their job opening. Most importantly, a project is either remunerated on an hourly or on a fixed basis. The hourly-pay option provides employers with additional control mechanisms, such as keystroke logging and regular screenshots of workers' screens. The trade-off for the increased monitoring facilities is that employers need to pay workers for their time regardless of the quality of work they provide. In contrast, employers cannot monitor workers while they work under fixed contracts, but they can withhold payment if the workers' output is of poor quality. From a data analysis perspective, hourly-priced projects are attractive because we observe the working hours with minimal measurement error. This is in contrast to fixed contracts, where working times are not monitored. Thus, we exclude fixed-price projects from our analysis.[7] Since digital trace data can be noisy with unrealistic outliers such as negative wages or hourly wages in the thousands of USD, we remove projects with an hourly wage in the bottom 1% and the top 99% of the distribution.

After choosing the contract type, employers select a broad and specific project category. There are 12 broad project categories (for example *writing*, *design & creative*, and

[6]Roughly 6% of the total labor supply originates from the United States while over 40% of the demand originates from the United States (Stephany et al., 2021).

[7]Projects with hourly wages and those with fixed price contracts exhibit similar characteristics along other dimensions, indicating no significant differences between the two samples. See Online Appendix Table A2.

sales & marketing) and about 90 specific project categories (such as *creative writing*, *grant writing*, or *medical writing*). Then, employers specify skill requirements by choosing from a constantly updated dictionary of roughly 4,000 skills. The dictionary includes broad skills (such as *writing*, *graphic design*, or *social media management*) and specific skills (such as *Microsoft Word*, *Adobe Photoshop* or *Google Analytics*). On average, employers list a median of four skills per project.

After specifying the skill requirements for a project, employers define a set of project characteristics and expectations: the desired experience level of the worker (novice, intermediate, or expert), the expected project duration (ranging from less than one week to more than six months) and workload (full-time vs. part-time). For completed projects, our data also contains information on who applied, who was hired, how much was paid per hour, how many hours were billed, and the resulting total earnings in USD. This information is also visible to all labor market participants once the project is completed.

When creating their freelancer profile, workers provide relevant background information, including their first name, self-description, asking wage, country of origin, language expertise, formal education, skills, completed projects, employer feedback, and availability to work. Besides the free text self-description, workers specify their expertise by selecting skills from the same skill dictionary that employers use for project requirements. On average, workers show a median of nine skills on their profile. These skills are self-reported. However, we assume the likelihood of workers misrepresenting their skills to be low, as overstating one's skills can result in poor employer ratings. As ratings play a crucial role, workers are incentivized to be truthful about their skills. That said, we cannot rule out the possibility that workers might misrepresent their skills.[8]

The workers do not explicitly mention their gender on their profiles. We infer gender based on first names and the country of residence of workers in the United States using the R-package *gender* (Mullen, 2021). This package assigns a probability of being male or female to each first name based on historical U.S. Census and Social Security data sets. A probability of 0 indicates that there were only males and 1 that only females were associated with a given name in the administrative data records. However, it is not infrequent that the same name is associated with men and women. For example, the name Andrea, depending on the country and cultural context, can be female or male. In this case, the R-package *gender* provides a probability between 0 (only male) and 1 (only female). To minimize the noise in our data set, we implement a 10% to 90% cut-off: we only include workers in our analysis with a first name that received a probability of 10%

[8]In particular, we cannot rule out that women might systematically under-report their skills because of, for instance, lower self-confidence.

or less (male) or 90% or more (female).[9] Restricting our analysis to workers based in the United States helps to further reduce the share of workers with an uncertain or unknown gender.

Contrary to gender, formal education is explicitly mentioned on worker profiles. Degree and university names are not standardized, however. Instead, workers describe their educational background in a free-text field. As a result, the data are messy. For example, workers describe a bachelor's degree in various ways, including as "bachelor," "bachelor's," or "bachelor's degree." We use a string matching approach to match the free-text input to the educational levels of the U.S. Bureau of Labor Statistics (Bureau of Labor Statistics, 2022).[10]

Most completed projects contain information on the applicants. Matching workers to their applications allows us to build a detailed application history for each worker. However, the application data is not complete as some projects do not contain information on applicants. There are two reasons for this. First, our data set does not include all projects workers applied to. Second, some workers are invited to jobs directly without applying. In total, approximately 24,000 workers in our data set applied to about three million projects and landed about 46,000 of them (about 60 applications to land one job). In contrast to project information, worker profiles are subject to change. As they complete projects, workers accumulate experience and adapt their asking wage, self-descriptions, and skills. The data on workers were collected in two waves in 2020 and 2022. Having only snapshots of time-varying worker profiles represents a limitation in our data, especially concerning two essential control variables: asking wages and worker skills.

Nonetheless, based on comparing the two snapshots, workers' profiles are rather static. In particular, approximately 75% of the workers had changed less than five skill tags in their profiles. For the asking wage, we find that the workers' asking wages had remained unchanged for 56% of the workers. Given that projects usually take place within a relatively short time window (the median time between project start dates is 20 days) and most workers' online freelancing careers are relatively short (the median time between the first and last project start dates for workers is roughly half a year), measurement error due to changing profiles should not drive our results.[11]

[9]Our analysis assumes a gender binary, inferred solely from first names, an approach that inevitably overlooks the spectrum of self-experienced gender identities. We acknowledge this limitation, yet the focus remains on perceived gender as it predominantly informs societal biases and differences in labor market outcomes between genders.

[10]We validated our approach with a random sample of 100 hand-labeled workers, achieving close to 100% accuracy.

[11]To the extent that workers' skill tags and asking wages change, we see this as a source of (classical) measurement error, which will attenuate the corresponding regression coefficients toward zero.

3 Summary Statistics and the Raw Gender Wage Gap

After having described how our data is collected, we move to quantifying the gender wage gap in our data, and study how worker-level characteristics affect the wage gap.

3.1 Describing the Raw Gender Wage Gap

In this section, we present summary statistics, provide background information on the workers, and describe the types of projects they are engaged in. We observe a considerable gender wage gap in our dataset that persists across different occupations, levels of formal education and over time. Table 1 presents basic summary statistics of worker activity by gender. Our main analysis sample consists of 45,107 projects completed by 23,425 U.S. online workers between the years 2015 and 2021. Within this sample, we observe a nearly equal distribution of workers by gender, with 11,570 identified as male and 11,855 as female. Men and women completed nearly the same number of projects, 23,421 and 21,686, respectively.

While men and women have completed approximately the same number of projects, we observe that men receive hourly wages that are, on average, approximately $12 higher than women.[12] Primarily, the wage disparity on an hourly basis, leaning largely on outliers, reveals that men typically hit a median earning of $35 per project, as compared to women who reach $25.[13] Even if women earn lower hourly wages, their projects are longer. An average project done by a male is 29 hours, while the corresponding number is 37 hours for women. However, the difference in average project duration seems to be driven by outliers. The median project length is eight hours for men, and about eight hours for women.

Table 1 presents the distribution of male and female workers across the different project categories. Men and women tend to work on projects in different categories. Women are over-represented in *Admin support*, *Customer service*, *Translation*, and *Writing*, while men are over-represented in *Data science & analytics*, *Design & creative*, *Engineering & architecture*, *IT & networking*, *Legal*, *Sales & marketing*, and *Web, mobile & software development* projects. The relative under-representation of men in the lower-skill project types such as *Admin support* and *Customer service* is not reflected in workers' self-reported education. Generally, women are slightly more educated than men. Moreover, the share of women who have not disclosed their education is smaller than the share of men with missing information on education.

[12] We include both the worker's hourly wage and log-hourly wage for illustrative purposes. We solely utilize log-hourly wage as the dependent variable in our regression analyses.

[13] For the distribution of hourly wages by gender, see appendix A.1.

Table 1. Basic summary statistics

	Male		Female		Difference in means
	Mean	Median	Mean	Median	(male – female)
Hourly wage	42.160 (26.782)	35	30.575 (21.554)	25	11.585***
Hourly wage (log)	3.537 (0.666)	3.555	3.201 (0.667)	3.219	0.336***
Worker characteristics					
Project length (hours)	28.527 (63.588)	8	37.257 (95.539)	8.330	-8.730***
PhD	0.046 (0.209)	-	0.043 (0.203)	-	0.003
Master	0.188 (0.390)	-	0.214 (0.410)	-	-0.026***
Bachelor	0.453 (0.498)	-	0.467 (0.499)	-	-0.014**
Associate	0.048 (0.214)	-	0.059 (0.237)	-	-0.011***
High school	0.019 (0.138)	-	0.019 (0.137)	-	0
No degree	0.020 (0.140)	-	0.020 (0.142)	-	0
Degree unknown	0.226 (0.418)	-	0.178 (0.382)	-	0.048***
Main project categories					
Accounting & consulting	0.051 (0.221)	-	0.048 (0.214)	-	0.003*
Admin support	0.038 (0.192)	-	0.156 (0.363)	-	-0.118***
Customer service	0.008 (0.091)	-	0.022 (0.146)	-	-0.014***
Data science & analytics	0.048 (0.214)	-	0.016 (0.127)	-	0.032***
Design & creative	0.179 (0.383)	-	0.164 (0.371)	-	0.015***
Engineering & architecture	0.043 (0.202)	-	0.015 (0.121)	-	0.028***
IT & networking	0.052 (0.223)	-	0.007 (0.086)	-	0.045***
Legal	0.026 (0.159)	-	0.019 (0.136)	-	0.007***
Sales & marketing	0.154 (0.361)	-	0.140 (0.347)	-	0.014***
Translation	0.011 (0.106)	-	0.030 (0.172)	-	-0.019***
Web, mobile & software development	0.213 (0.409)	-	0.056 (0.230)	-	0.157***
Writing	0.175 (0.380)	-	0.327 (0.469)	-	-0.152***
Number of projects	23,421		21,686		
Number of workers	11,570		11,855		
Share of females			50.61%		

Note: The values presented are based on U.S. workers who completed at least one project between the years 2015 and 2021. Standard deviation in parentheses. We report both hourly wages and log-hourly wages. In our analysis, we exclusively utilize log-hourly wages as dependent variable. In Column 6, we test the statistical significance of the differences in means between female and male workers using two-sample t-tests. The significance levels are indicated by: * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Moreover, women do not only tend to complete projects in different categories. Women complete jobs in project categories that systematically pay lower hourly wages. Figure 1 illustrates this sorting into low vs. high-paying project categories by gender. We can see that the average hourly wage by project category decreases with an increasing share of females working in the given category. Figure 1 suggests that occupational segregation might be one of the key drivers of gender differences in wages. In other words, women earn less than men because they work in project categories that pay less. According to Figure 1, the average hourly wage for workers in *IT & networking* is \$56, but the share of female workers doing these projects is less than 10%. In contrast, *Admin support* has an average hourly wage of \$20 and a share of female workers of almost 80%.

However, women also earn lower wages than men when working within the same project category. Figure 2 shows the gender wage gap in each of the project categories. In most categories, women earn significantly less than men. The wage gap is particularly large in the high-paying categories *Accounting & consulting* and *Legal* but also in the low-paying categories *Admin support* and *Customer service*. In projects related to *Design & creative*, *Engineering & architecture*, and *Translation*, we do not observe a significant gender wage gap. The large wage gap in some of the categories suggests that male and female workers do different types of work within the same broad category. Besides a gender wage gap within project categories, we find a consistent wage gap across different levels of formal education and across time. Regardless of their level of formal education, women earn only about 70% of the hourly wage of their male counterparts with the same education. With minor fluctuations, the same gap persists throughout our observation period from 2015 to 2021. For details on hourly wages by level of education and over time see Appendix A2 and A3.

Result 1. *We observe strong occupational sorting by gender on the freelancing platform. Women are over-represented in lower-paying project categories. Even within the same category, women tend to earn lower wages.*

Figure 1. Average hourly wage and share of female workers in each project category

Note: The share of female workers working in each project category is plotted against the average hourly wage in USD in that category. The size of the points represents the market share of the respective project category.

The detailed data on worker skills and project skill requirements allows us to go beyond broad occupational categories and to analyze gender differences at the level of skills. In the remaining paper, we study the gender wage gap using regression analysis. Figure 2 suggests that a large share of the gender wage gap is explained by the fact that men and women work on different types of projects. However, we argue that including project categories would be subject to the "bad controls" critique because controlling for the project category inherently controls for part of the outcome (Angrist and Pischke (2009, pp. 64–68), Cunningham (2021, pp. 106–110)). To see why, assume that some women are excluded from high-paying project categories because of discrimination. Controlling for the project categories would result in controlling for a part of the labor market outcome (being hired to a high-paying project category) and consequently would result in underestimating the gender wage gap. Therefore, we do not control for project categories in our regression specifications. Instead, we focus on skills and application behavior.

Before moving on to decomposing the drivers of the gender wage gap, we quantify the *raw* gap.

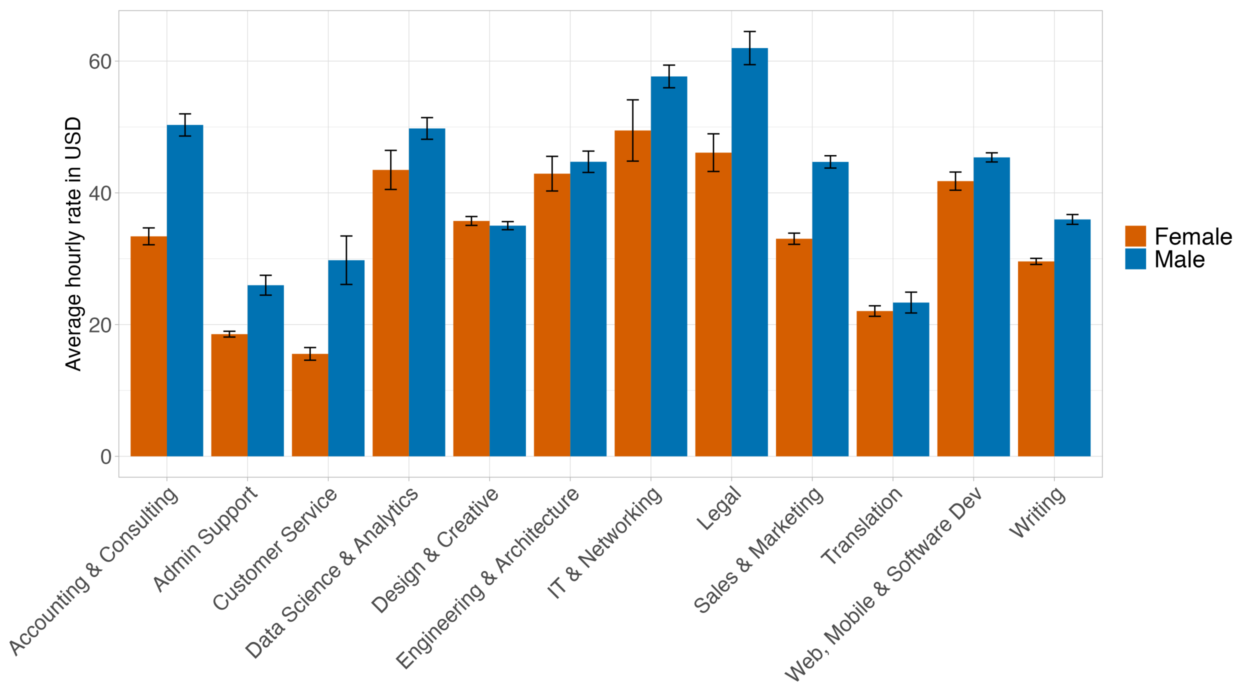

Figure 2. Average hourly wage in USD by gender and project category

Note: Average hourly wages in USD by project category and gender. Error bars represent 95% confidence intervals calculated as $+/-1.96 * \text{st. error}$.

3.2 Quantifying the Raw Gender Wage Gap

We account for formal education and time by incorporating education and year dummies, respectively. Additionally, we include dummy variables for the employers' home countries. These help account for regional or country-specific factors that could impact wages. By doing so, we can examine the wage gap without the interference of localized biases or other similar factors. We then employ a standard wage regression model, as described below.

$$\log(\text{Hourly wage})_{ijt} = \alpha + \beta\,\text{Female}_i + \rho\,X_{ijt} + \epsilon_{ijt} \tag{1}$$

where i represent a U.S. worker who completed project j in year t. The term *Hourly wage* captures the worker's log-hourly rate for each project, and *Female* is a binary variable indicating the worker's gender. The set of control variables, X_{ijt}, encompasses the worker's education level, the project's starting year, and dummy variables for the employer's country. The coefficient attached to the gender dummy variable quantifies the disparity in hourly wages between men and women in log-points, the gender wage gap. We account for potential correlations in unobservables by clustering standard errors at the worker level.

Table 2. Raw gender wage gap

	Hourly wage (log)			
	(1)	(2)	(3)	(4)
Female	-0.336***	-0.311***	-0.347***	-0.323***
	(0.012)	(0.012)	(0.012)	(0.012)
Controls		✓		✓
Employer country	All countries	All countries	U.S. only	U.S. only
Number of projects	45,107	45,107	33,045	33,045
Number of workers	23,425	23,425	18,991	18,991
Adjusted R^2	0.059	0.164	0.065	0.163
Share of females	50.61%	50.61%	50.55%	50.55%

Note: This table documents the gender wage gap in log-hourly wages. Column 1 presents the results when only considering whether the worker is *Female* or male. In Column 2, our control variables are worker's educational degree, the year the project started, and employer's country of residence. Columns 3 and 4 present the results from estimating the models using data from U.S. based employer's country of residence only. Standard errors are clustered at the worker level and are reported in parentheses. Significance of difference indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We present the findings in Table 2. The baseline estimate in Column 1 indicates that women have 33.6 log-points lower hourly wages. In Column 2, we estimate Equation 1. We show that the gender wage gap is virtually unaffected by the inclusion of education, time and employer controls. The gender wage gap decreases only marginally, resulting in a 31% difference between men's and women's earnings. This confirms our descriptive results, which indicate a substantive wage gap among platform workers.

In traditional labor markets, education is a significant factor influencing the gender wage gap. However, its role may be less pronounced in online freelancing markets for various reasons. For instance, the inherent instability of the gig economy, characterized by the absence of long-term employment and benefits, diminishes the need for employers to prioritize educational credentials during worker selection. Furthermore, freelancing platforms often value tangible skills and a proven work history over formal educational backgrounds. It's also worth highlighting that our sample is exceptionally well-educated, with approximately 70% holding a college degree or higher. In contrast, only about 40% of U.S. adults have similar qualifications (Pew Research Center, 2022). A study by Herrmann et al. (2023) examines the effects of education on success in online labor markets more closely. Their findings align with ours, indicating that higher education does not necessarily lead to higher wages in online platform labor markets.

Ghani et al. (2014) and Lehdonvirta et al. (2019) argue that costs, frictions, information asymmetries, and ethnic networks play a significant role in shaping the types of cross-border transactions that take place in platform contexts. To demonstrate that our results are not sensitive to employers' home countries, we repeat the analysis of Columns 1 and 2 using U.S. employers only as a sensitivity check in Columns 3 and 4 (Table 2). The estimation results from using data from U.S. employers only are virtually indistinguishable from the results reported in Columns 1 and 2. To maximize sample sizes and statistical power, we use data from all employer countries and corresponding employer country dummies throughout the remainder of this paper.

Result 2. *The unexplained gender wage gap conditional on education is 33.6 log-points.*

4 Skills, Application Behavior and Hourly Wages

In the previous section, we demonstrated that a considerable gender wage gap persists across different occupations, levels of formal education and over time in online freelancing. Next, we investigate how workers' skills and application behavior influence the relationship between gender and hourly wages.

We use standard linear regression models to decompose the gender wage gap into explained and unexplained parts. We proceed by gradually incorporating variables that capture factors related to either preferences that the workers have or the constraints that the workers face. After controlling for these, any unexplained gender gap could be due to either employer discrimination or unobservable differences in preferences or constraints. This residual earning difference between genders is captured by the coefficient of a female dummy variable in our regression analysis.

We approach this step by step in the following subsections. First, we outline our method for quantifying the value of using a machine learning-based method for estimating the value of workers' skills. Second, we detail our approach to capturing two key aspects of application behavior: the job amenities associated with projects to which workers apply and their asking wages. Finally, we integrate these components in a decomposition, utilizing the Gelbach (2016) method to determine the portion of the hourly wage gap each factor accounts for.

4.1 Worker Skills

Our results so far show that men are more common in higher-paying jobs. We have found that this difference is not explained by education, but it might be because of different skills between workers. For instance, one cannot work in coding without coding skills. We next proceed to discussing how we operationalize the measure of skills in regression framework.

Our data contains granular information on worker skills and the skill requirements of projects. There are more than 4,000 individual skills that are being combined in a multitude of ways. To effectively analyze such granular data, we propose a machine learning approach to *learn* the value of skills. Using the Random Forest algorithm (Breiman, 2001), we compress the high-dimensional binary skill data into a one-dimensional continuous representation. The non-linear nature of Random Forest allows us to incorporate interactions between skills. Conceptually, this parsimonious measure represents the estimated market value of each combination of skills. In other words, we predict the hourly wage a worker could expect to earn given their combination of skills.

In practice, we proceed in the following way. First, we train a Random Forest model on the skill requirements and hourly wages of project postings. Thereby, the model *learns* the relationship between skills and hourly wages. Our best model achieves an R^2-score of about 0.27 on a hold-out test set.[14] The R^2 of 0.27 implies that the best predictive model which uses project skill tags as explanatory variables explains 27% of the total variance of hourly wages. Second, we use this model to predict the market value of workers' self-reported skill sets. Third, we use this newly created variable, denoted as *Skills*, in our regression analysis to reveal what share of the gender wage gap can be attributed to differences in the skills female and male workers offer on the platform. Figure 3 displays the predicted value of workers' skill sets by gender. Our model predicts considerably lower skill values for female workers (see also Appendix Figure A5).

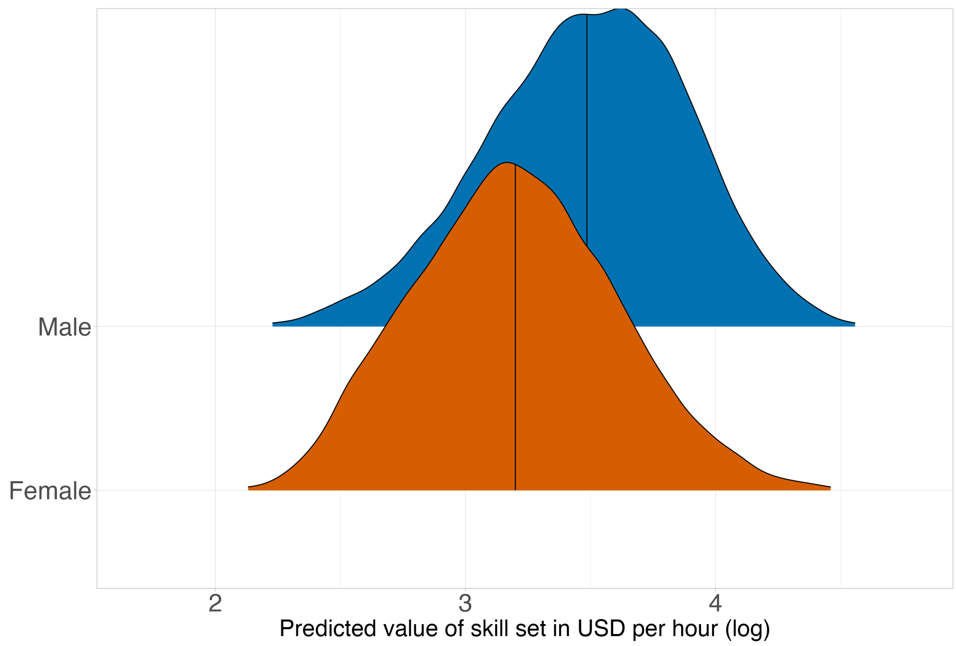

Figure 3. Predicted value of workers' skill sets by gender

Note: This figure shows the predicted value of workers' skill sets by gender in USD per hour (log). The black vertical line represents the mean.

While it might be statistically feasible to use worker skill dummies directly in a regression analysis to account for workers' skills, we choose a machine learning approach

[14]See Appendix A.2 for details on the hyperparameter-tuning, the training process, model performance, and a comparison to other machine learning models.

for several compelling reasons. First and foremost, from a theoretical perspective, the Random Forest model has the advantage of recognizing potential interactions between variables. Recent research demonstrated that the complementarity between skills plays an important role for workers' earnings (Stephany and Teutloff, 2024). For instance, possessing skills in both Python and Javascript might be more valuable than just being adept at one of them individually. Moreover, we will utilise the one-dimensional fitted values in the next section when controlling for the skill requirements of the applications made by workers.[15]

We introduce the variable *Skills* into the model to capture the market value of workers' self-declared skills. This allows us to compare the raw gender wage gap (Table 2) with the gender wage gap after adjusting for the predicted market value of worker skills. Moreover, we sidestep the "bad controls" issue discussed in 3.2, which would arise if we directly controlled for project categories. We use ordinary least-squares regressions to estimate the gender wage gap. The estimating equation is:

$$\log(\text{Hourly wage})_{ijt} = \alpha + \beta\,\text{Female}_i + \gamma\,\text{Skills}_i + \rho\,X_{ijt} + \epsilon_{ijt} \tag{2}$$

where i refers to a U.S. worker who completed project j in year t. *Hourly wage* denotes the worker's log-hourly rate per project, while *Female* is a binary variable indicating the gender of the worker. *Skills* represents the machine learning-based prediction of the market value of workers' skills. The set of control variables, denoted by X_{ijt}, includes the same controls as outlined in Equation 1. We cluster the standard errors at the worker level.

A few remarks on Equation 2 are worth making. First, it is important to note that a worker's skill set at the beginning of a project might not match their skill set at the time of our data collection. This difference introduces a measurement error in the *Skills* variable. Such classical measurement error can bias regression coefficients toward zero. If we further assume this error is not correlated with gender — a logical assumption — this means the measurement error in *Skills* may cause us to overestimate the gender wage gap when accounting for skills.

Second, constructing the variable *Skills* captures the average market wages for different skill combinations. Thus, it does not take into account the possibility that men and women working on projects with the same skill requirements might be paid different wages. Nonetheless, when we include the variable *Skills* in the regression model, the differences in the average market wages conditional on skills will be reflected in the coefficient of the *Female* dummy variable.

[15] The robustness of our findings has been tested through regression analyses with skill dummies (see Online Appendix A.7). Detailed results will be presented in subsequent sections, but we can already indicate that our ultimate result — the decomposition of the gender wage gap into skills and application behavior — remains unaffected, irrespective of how we integrate worker skills into our analyses.

Third, it is well understood that regularization methods such as Random Forest result in better forecasting ability of the model, but at the same time introduce bias into the coefficients that are being regularized. Chernozhukov et al. (2018) emphasize that including a regularized term into a regression model with a binary dummy variable can result in a situation, where the bias due to regularization also transmits into the parameter of interest. This regularization bias issue is similar to omitted variable bias. This bias can be substantial even in moderately sized samples. It is essential to clarify our approach in this context: we utilize data on project skill requirements to estimate the machine learning model without incorporating a gender dummy. Only the predictions from this model are subsequently used as a control variable in a regression that includes a gender dummy. In other words, we do not estimate a model where the regularized term and a binary indicator variable would be estimated simultaneously using the same data. As such, the coefficient on the gender dummy in this latter regression remains unaffected by regularization bias.

Table 3. Gender wage gap conditional on workers' skills

	Hourly wage (log)	
	(1)	(2)
Female	-0.124***	-0.130***
	(0.011)	(0.012)
Skills	0.675***	0.680***
Controls	✓	✓
Employer country	All countries	U.S. only
Number of projects	45,107	33,045
Number of workers	23,425	18,991
Adjusted R^2	0.317	0.320
Share of females	50.61%	50.55%

Note: This table documents the gender wage gap in log-hourly wages. The control variables in our analysis comprise the worker's educational degree, the year in which the project commenced, and employer country dummies. In addition, the regression model includes a measure of the market value of workers' skills derived from a machine learning model. In Column 2, we report the results from estimating the model using data from U.S. based employer's country of residence only. Standard errors are clustered at the worker level and are reported in parentheses. Significance of difference levels indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In Table 3, we present our estimation results. After accounting for worker skills, the gender-based difference in hourly wages amounts to 12.4 log points. When comparing this to the findings in Table 2, it is evident that differences in worker skills explain a significant portion of the gender wage gap. In Column 1 of Table 3, we analyze data from all workers. To ensure our results are not influenced by the employer's country of residence, we conduct a sensitivity analysis in Column 2, focusing exclusively on U.S.-based employers. The point estimate remains virtually unchanged between Columns 1 and 2.

Result 3. *After accounting for skills, the unexplained gender wage gap decreases from 33.6 to 12.4 log points.*

4.2 Workers' Past Application Behavior

A potential factor contributing to the gender wage gap is the systematic sorting of men and women into projects with differing characteristics. Even when they have the required skills, women might shy away from demanding, higher-paying jobs, if they come with less appealing job characteristics. Additionally, there exists laboratory and behavioral evidence suggesting that men tend to display greater confidence during job applications than women, further motivating this analysis. (Niederle and Vesterlund, 2011).

Our data capture various measures of project characteristics at the job posting level. When employers post a project, they use standardized terms to detail its specifics, including the expected engagement duration and contract type. In addition, we observe the projects' preferred worker tier, as expressed by the employer, and presented to workers considering bidding for projects. When creating a project, the employer can choose what tier of worker they are looking for, choosing from three options ranging from "Looking for someone relatively new to this field" to "Looking for comprehensive and deep expertise in this field." While projects vary in numerous characteristics (some not captured in our data set), we expect these to be the most salient to the workers because the platform user interface allows workers to filter the projects along these dimensions in the search dialogue. We also have information on workers' expected wages declared at the time of the data collection and the number of applications submitted over time for both hourly and fixed-price projects. In addition, we control for the average skill value of the projects workers applied to in the past, obtained by using machine learning as outlined in Section 4.1. This variable captures the estimated expected hourly wage of the applied-for projects.

As in Section 4.1, we want to avoid controlling for the job characteristics of the project they are currently doing (project indexed as j). On the other hand, when measuring *past* applications, we want to encompass all applications made by the worker, irrespective of the hiring outcome. In practise, we operationalize the project characteristics in the regression models as shares of projects with a certain characteristic in the last 30 days. For instance, if the worker has applied to n (where $n > 1$) projects over the past 30 days, and half of the applied projects had an expected duration of *Under 1 week*, and half of the *1 to 3 months*, then the value of these two variables would be 0.5, while the share of other contract lengths would be zero.

There are no obvious theoretical approaches for choosing the optimal length for the time window over which to average the job applications. We have to strike a balance between minimizing unobserved changes and maximizing the sample size. The job application time window should be short enough so that we can be relatively confident that the workers' offline circumstances (such as employment status, living arrangements, or education level) have not changed. On the other hand, the time window should be long enough for the sample sizes to not become too small. Therefore, our primary analysis consists of the applications made by workers during the 30 days prior to starting a given project. However, we demonstrate that our results remain qualitatively and quantitatively similar if, instead of 30 days, we look at alternative time window lengths of up to 365 days.[16]

Including application behavior information reduces the sample sizes compared to those reported in Table 1. There are two reasons for this. First, some workers get the first project they apply to. Additionally, due to limitations of our data collection we do not always capture the complete list of job applicants in projects, resulting in missing data. Consequently, the application behavior data consist of 27,698 projects carried out by 13,267 workers (down from 45,107 projects completed by 23,425 workers). When comparing worker and project characteristics (see Table A3) to the primary analysis sample presented in Table 1, we find that educational qualifications, as well as project length are largely unchanged between the two analysis samples. This supports the assumption that information on applications are missing at random from the data.

Table 4 shows the descriptive statistics for this particular sample. We observe that the average hourly wage and average log-hourly wage remain largely consistent between the two analysis samples. When examining job-search behavior, the data reveals that men are more inclined to apply to projects seeking experts. In contrast, women seem to favor projects offering longer-term contracts. More specifically, men predominantly apply to projects expected to be completed in less than 10 hours, whereas women show a preference for extended durations. Furthermore, the hourly wages women typically ask for are about $21 less than their male counterparts. Given the considerable differences in worker skills, it is unclear whether these differences in asking wages result from different types of contracts being available to workers with different skills or gender differences in preferences concerning job amenities.

[16] See Online Appendix Table A4.

Table 4. Basic summary statistics with information on job-search behavior

	Male		Female		Difference in means
	Mean	Median	Mean	Median	(male – female)
Hourly wage	43.481	37	31.244	25	12.237***
	(27.065)		(21.893)		
Hourly wage (log)	3.577	3.611	3.222	3.219	0.355***
	(0.650)		(0.672)		
Application behavior					
Share of applications: Desired worker experience					
Novice	0.117	0	0.169	0.067	-0.052***
	(0.190)		(0.237)		
Intermediate	0.470	0.500	0.505	0.500	-0.035***
	(0.269)		(0.282)		
Expert	0.404	0.385	0.314	0.278	0.090***
	(0.286)		(0.284)		
Desired worker experience unknown	0.009	0	0.012	0	-0.003***
	(0.065)		(0.077)		
Share of applications: Contract type					
Less than 10 hours	0.268	0.231	0.240	0.200	0.028***
	(0.242)		(0.244)		
Part-time	0.305	0.262	0.347	0.312	-0.042***
	(0.255)		(0.276)		
Full-time	0.161	0.036	0.169	0.026	-0.008**
	(0.240)		(0.249)		
Contract type unknown	0.266	0.250	0.244	0.182	0.022***
	(0.255)		(0.260)		
Share of applications: Expected duration					
Under 1 week	0.126	0	0.114	0	0.012***
	(0.192)		(0.192)		
Less than 1 month	0.177	0.100	0.135	0.037	0.042***
	(0.232)		(0.212)		
1 to 3 months	0.107	0.036	0.101	0	0.006**
	(0.171)		(0.176)		
3 to 6 months	0.073	0	0.082	0	-0.009***
	(0.152)		(0.169)		
More than 6 months	0.180	0.100	0.241	0.150	-0.061***
	(0.233)		(0.280)		
Expected duration unknown	0.338	0.333	0.327	0.333	0.011***
	(0.249)		(0.263)		
Number of past applications					
Number of applications	12.867	7	10.716	6	2.151***
	(19.199)		(15.326)		
Number of applications for fixed projects	5.230	3	4.348	2	0.882***
	(8.619)		(7.140)		

Table 4 continued	Male Mean	Male Median	Female Mean	Female Median	Difference in means (male – female)
Worker's declared wage					
Asking wage	70.171	55.780	48.799	38	21.372***
	(61.036)		(40.833)		
Asking wage (log)	3.999	4.021	3.638	3.638	0.361***
	(0.715)		(0.703)		
Mean value of past applications					
Mean value of past applications	3.476	3.490	3.217	3.228	0.259***
	(0.333)		(0.370)		
Number of projects	14,621		13,077		
Number of workers	6,602		6,665		
Share of females			50.24%		

Note: The values presented are based on U.S. online workers who completed at least one project between the years 2015 and 2021. Standard deviation in parentheses. We report both hourly wages and log-hourly wages, as well as asking wage and the log of the asking wage. In our analysis, we exclusively utilize log-hourly wages as dependent variable and the the log of the asking wage as independent variable. Besides *Asking wage*, *Asking wage (log)*, and *Mean value of past applications*, information related to application behavior is coded as shares. In Column 4, we test the statistical significance of the differences in means between female and male workers using two-sample t-tests. The significance levels are indicated by: * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

We extend Equation 2 to account for the application behavior as follows:

$$\begin{aligned}\log(\text{Hourly wage})_{ijt} = \alpha &+ \beta\,\text{Female}_i + \delta_1\,\text{Desired worker experience}_{ijt} \\ &+ \delta_2\,\text{Contract type}_{ijt} + \delta_3\,\text{Expected duration}_{ijt} \\ &+ \delta_4\,\text{Number of past applications}_{ijt} + \delta_5\,\log(\text{Asking wage})_i \\ &+ \gamma\,\text{Skills}_i + \zeta\,\widetilde{\text{Skills}}_{ijt} + \rho\,X_{ijt} + \epsilon_{ijt},\end{aligned} \tag{3}$$

where t refers to the year when worker i started working on project j. *Desired worker experience* captures the share of applications made to projects with different desired worker experience (entry level, intermediate, and expert). *Contract type* captures the share of applications by contract type, namely less than 10 hours, part-time, full-time, and unknown. *Expected duration* captures the distribution of expected durations for the applied projects, ranging from *Under 1 week* to *More than 6 months*. The *Number of past applications* vector includes the total number of applications made by the worker, including both hourly and fixed payment projects, as well as only the number of applications for projects with fixed payment. *Asking wage* is the worker's asking log-wage declared in their profile. Finally, as in Equation 2, we control for the workers' skills by including the estimated value of worker skills (*Skills*). Moreover, we include a measure of the expected value of the applied-for projects during the 30-day time window, denoted by $\widetilde{\text{Skills}}$. We create this variable by predicting the value of a project based on its skill requirements using the same machine learning model as described in Section 4.1.

The *Asking wage* variable is measured with some error because it is based on a snapshot collected at the time of data collection, which might differ from the asking wage at the time of applying for a project. Following a similar line of argumentation to the *Skills* variable, we expect that this measurement error is uncorrelated with gender, which will lead to attenuation of the regression coefficient on the *Asking wage*, and, consequently, will lead to overestimation of the gender wage gap conditional on asking wages.

We present the estimation results in Table 5. In Column 1, we examine whether our previous findings from Table 3 hold in the smaller sample. The coefficient on the *Female* dummy is statistically indistinguishable from the one reported in Table 3. This suggests that the attrition from the sample is not correlated with either gender or the labor market outcomes of workers. We proceed by gradually introducing additional variables related to workers' application behavior in subsequent columns. We find that *Desired worker experience*, *Contract type*, *Expected duration*, and *Number of past applications* have minimal impact on the gender wage gap. However, controlling for *Asking wage* leads to a significant decrease of 5.7 log-points in the gender wage gap, from 9.3 to 3.6. Finally, controlling for the expected hourly wages of the applied-for projects based on their skill requirements eliminates the gender wage gap entirely.

The fact that both individual worker skills (variables Skills_i) and the average skill level of past applications ($\widetilde{\text{Skills}}_{ijt}$) independently affect wages suggests that men and women apply for jobs with different skill requirements, even when skill differences between them are controlled. Furthermore, $\widetilde{\text{Skills}}_{ijt}$ continues to be a strong predictor of wages, as shown in Column 9 of Table 5, even after accounting for other job characteristics of applied-for jobs. This indicates that women apply for jobs that are have looser skill demands, even when all other observable job characteristics and worker skills are taken into account. A likely explanation for this might be that, as laboratory evidence suggests (see (Niederle and Vesterlund, 2011)), men typically exhibit greater confidence than women in job applications.

Overall, our findings indicate that workers' application behavior can account for 13.7 log-points of the gender wage gap, conditional on the market value of workers' skills. This effect operates via three channels. First, conditional on skills, women apply for projects with different job characteristics and skill requirements. Moreover, women ask for lower hourly wages even when their job application behavior is held constant. Lastly, women apply for projects with lower expected hourly wages. Taken together, these three differences are enough to explain the gender wage gap.

Result 4. *Workers' skills and application behavior fully explain the gender wage gap.*

Table 5. Gender wage gap conditional on skills and application behavior

	Hourly wage (log)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female	-0.355***	-0.327***	-0.137***	-0.100***	-0.098***	-0.091***	-0.093***	-0.036***	-0.011
	(0.016)	(0.016)	(0.015)	(0.013)	(0.012)	(0.012)	(0.013)	(0.009)	(0.009)
Skills			0.689***	0.530***	0.525***	0.507***	0.508***	0.177***	0.107***
			(0.017)	(0.015)	(0.015)	(0.015)	(0.015)	(0.014)	(0.014)
Mean value of past applications									0.247***
									(0.013)
Controls		✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓
Contract type					✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓
Number of past applications							✓	✓	✓
Asking wage (log)								✓	✓
Number of projects	27,698	27,698	27,698	27,698	27,698	27,698	27,698	27,698	27,698
Number of workers	13,267	13,267	13,267	13,267	13,267	13,267	13,267	13,267	13,267
Adjusted R^2	0.067	0.179	0.339	0.470	0.471	0.476	0.477	0.662	0.671
Share of females	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%

Note: This table presents the gender wage gap conditional on application behavior. Column 1 shows the raw gender wage gap. In Column 2, we control for conventional controls (project start year, employer country dummies, and worker education). Column 3 presents the results from the regression specification where we account for the market value of workers' skills and control variables (project start year, employer country dummies, and worker education). In Columns 2 to 7, we progressively incorporate controls for workers' job application behavior. *Number of past applications* encompasses the number of all applications and the number of applications to fixed price projects. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels are indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.3 Decomposition Analysis

In the previous section, we used ordinary least squares regressions to examine the impact of traditional control variables, workers' skills, and job-application behavior to explain the gender gap in hourly wages. For a more nuanced understanding of the factors driving the gender wage gap, we estimate the specific contribution of each variable to the observed wage discrepancy using the decomposition methodology developed in Gelbach (2016).

The core idea of the Gelbach decomposition technique is to break down the aggregate influence of covariates on the gender wage gap in a manner that remains unaffected by the sequence in which additional covariates are incorporated. Through this approach, we are able to discern the specific contribution of every omitted variable, as articulated in Equation 3, to the shift in the coefficient of the *Female* variable. More concretely, consider a model where the dependent variable Y, log-hourly wage, is a function of X_1 and X_2:

$$Y = \alpha \, \text{Female} + \beta_1 X_1 + \beta_2 X_2 + \epsilon \tag{4}$$

Here, X_1 denotes a single variable, and X_2 encompasses all other covariates from Equation 3. Now, suppose that we exclude the matrix X_2 from our model. We can quantify the resulting bias in β_1 as $(X_1'X_1)^{-1}X_1'X_2\beta_2$. The contribution of each element k in X_2 contributes to this bias which can be expressed as $(X_1'X_1)^{-1}X_1'X_{2k}\hat{\beta}_{2k} = \hat{\Gamma}_k\hat{\beta}_{2k}$. $\hat{\Gamma}$ is an estimate we get using an auxiliary regression of gender on each k. By dividing the estimate of this bias from the omitted variable by $\hat{\alpha}$ — the raw gender pay gap — we get an estimate of k's contribution as a fraction of the baseline unconditional wage gap:

$$\tilde{\pi}_k = \frac{\hat{\Gamma}_k\hat{\beta}_{2k}}{\hat{\alpha}} \tag{5}$$

Aggregating these relative contributions across all omitted variables illustrates their joint influence on the baseline unconditional gender wage gap. While the results from any decomposition method are — by their nature — correlational, they provide useful insights on the relative contribution of each covariate to the gender wage gap.

Figure 4 presents the computed $\tilde{\pi}_k$ parameters, paired with their respective 95% confidence intervals. These parameters represent the decomposition of the shift in point estimates between the baseline model (Table 5, Column 1) and the full model as specified in Equation 3 (Table 5, Column 9).

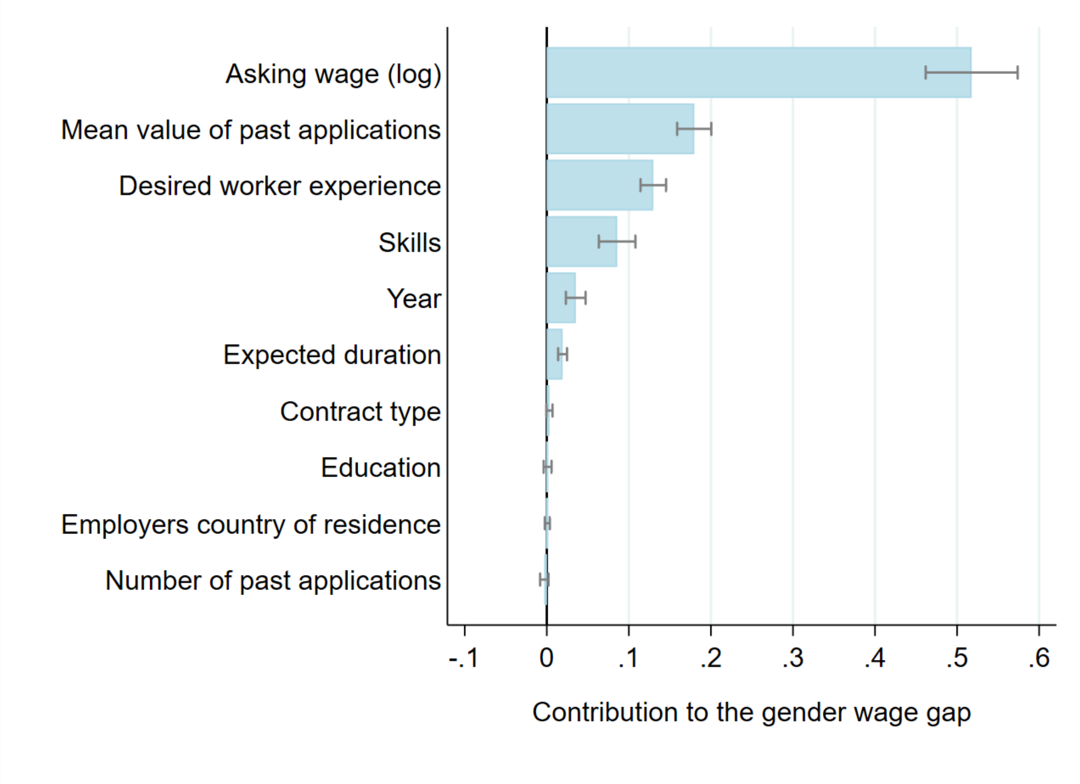

Figure 4. Gelbach decomposition

Note: The figure applies the method from Gelbach (2016) to show how much each factor contributes to the gender wage gap. These factors are education, preferred worker experience, type of contract, expected duration, number of past applications, the wage workers ask for, skills, average skill value of projects workers applied to over 30 days, employer's country, and year dummies. Error bars represent 95% confidence intervals calculated as $+/-1.96 * \text{st. dev}$.

Our findings highlight that *Asking wage*, *Mean value of past applications*, *Desired worker experience*, and *Skills* contribute most to the wage gap. The *Asking wage* alone accounts for nearly half of the link between gender and hourly wage. This suggests that increased wage requests by females could potentially mitigate wage disparities. *Mean value of past applications* accounts for the second-largest share, approximately one-fifth of wage disparities, with women tending to apply for lower wage projects compared to men, even when *Skills* are held constant.

Desired worker experience contributes to around 13% of the gap. Factors such as occupational segregation, career interruptions due to family responsibilities, or the prevalence of part-time work among women may affect these aspects. *Skills* explain about 9% of the wage gap, reflecting the differential skill sets and the varying market value of these skills among men and women. The remaining 6% of the gender wage gap can be attributed to the project's *Expected duration*, *Contract type*, worker's *Education*, *Employer country of residence*, and *Year* the project started.

Result 5. *The asking wage is the most important factor contributing to the gender wage gap — accounting for nearly half of the hourly wage differences between men and women.*

5 Discussion of Potential Mechanisms

Our analysis presented in the previous section shows that while women earn substantially less than men, their application behavior and skills can fully account for the difference. Moreover, we find that, on average, women apply to a larger number of projects and for projects with a longer duration, which, at least partly, might mitigate the gap in hourly wages. So far, our analysis has been silent on the potential underlying mechanisms that result in these differences. In this section, we proceed to explore potential explanations for the underlying patterns we observe.

The potential mechanisms we test for include the so-called self-fulfilling discrimination hypothesis, according to which women expect to be discriminated against, which leads them to not apply to certain projects or adjust their wage bids down to counteract this discrimination (Coate and Loury, 1993; Fluchtmann et al., 2021; Lundberg and Startz, 1983). Another potential mechanism that has been identified in the literature is the so-called "job-flexibility penalty." Goldin (2014) presents evidence for imperfect substitution between workers that can result in a convex hours-earnings relationship. In other words, women earn less per hour than men because they tend to work fewer hours. In addition, Cook et al. (2021) demonstrate that a major factor why male Uber drivers earn more than their female counterparts is the return to experience. Male drivers work more hours per week, which increases their productivity per hour.

Our data are inconsistent with all of these hypotheses. Instead, our data suggest that women work more hours and are more likely to get hired, even if they earn less per hour of work. This could be because men, on average, have better off-platform job market opportunities, which increase their on-platform wages while resulting in fewer platform work hours. On the other hand, women might have fewer opportunities offline, making online work — even with a smaller pay compared to comparable males — lucrative for them. Since we do not observe the family structure of workers, it is difficult to know if the gender gaps are more pronounced for workers with young children. Nonetheless, the flexibility of online freelancing may be particularly appealing to women, particularly those managing both work and family commitments. These family responsibilities could be one potential explanation for the differences in offline work opportunities. Kleven et al. (2019a), among others, has convincingly shown that child-rearing is strongly associated with the emergence of the gender wage gap. Unfortunately, a limitation in our analysis is that our data do not include information about whether the workers have children or are of child-bearing age.

5.1 Self-Fulfilling Discrimination

The finding that workers' asking wages are an extremely strong predictor of the gender wage gap leads to a hypothesis that the asking wages could result from workers' equilibrium wage-setting and application behavior in the presence of discrimination.[17]

This effect has been characterized by the term "self-fulfilling discrimination" in other contexts (Coate and Loury, 1993; Fluchtmann et al., 2021; Lundberg and Startz, 1983). In the context of job application behavior, this could emerge if women know that employers will discriminate against them; in response, they may adjust their wages down or apply to less demanding jobs. Without experimental data, it is difficult to gauge whether this effect is present in our data conclusively. Nonetheless, we can check if the gender gap exists in other labor market outcomes that we can measure, such as project length and the share of successful applications.

We show in Table 6 that the gender wage gap in project length is positive. That is, conditional on worker skills, women work on longer projects than men. Perhaps unsurprisingly, this effect is eliminated when we account for the desired worker experience and contract type. This suggests that women can — at least partly — offset their lower hourly wages by working more hours.

[17] The gap in asking wages is also documented in the platform labor context in Foong et al. (2018). Moreover, Roussille (2021) documents the existence of asking wage gaps in recruitment in traditional labor markets.

Table 6. Gender wage gap in project length

	Hours (log)			
	(1)	(2)	(3)	(4)
Female	0.100***	0.097***	0.047**	0.014
	(0.019)	(0.019)	(0.020)	(0.025)
Controls		✓	✓	✓
Skills			✓	✓
Desired worker experience				✓
Contract type				✓
Number of projects	45,107	45,107	45,107	27,698
Number of workers	23,425	23,425	23,425	13,267
Adjusted R^2	0.001	0.012	0.014	0.025
Share of females	50.61%	50.61%	50.61%	50.24%

Note: This table documents the gender wage gap in log-hours. The control variables in our analysis comprise the worker's educational degree, the year in which the project commenced, and employer country dummies. In addition, the regression model in Column 3 includes a measure of the market value of workers' skills derived from a machine learning model. In Column 4, we include additional controls for share of applications made to projects of into different contract types and the proportion of projects with different expected worker experience levels. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We repeat a similar exercise in Table 7 using application success rate as the dependent variable. We take the number of applications used as a control variable in Equation 3 and divide the number of successful applications by the total number of applications. According to results in Table 7, women are more likely to win the projects they apply to.

In summary, the results from regressions using either working hours or success rates as the dependent variable indicate that women, on average, work more hours and are more likely to get hired, even after we have accounted for skill differences. To test if women have to forego more in terms of wages to get hired is to study if wages have a higher predictive power on the application success for men than for women. We implement this test by calculating the correlation between the regression residuals calculated from specifications 1 to 4 in Table 7, and the hourly wages. Here, the idea is the following: if women have to adjust their wage bids down to counteract discrimination, we expect the correlation between wages and success rates to be negative and larger in absolute value for women than men. We report the results of this exercise in Figure 5. In general, we find that the correlation between wages and success rate — conditional on observable characteristics — in applications is small and non-distinguishable between men and women at conventional risk levels. The sole exception, where a statistically significant difference emerges, is specification 1, where no control variables are included. Even in this model, the correlation between wages and success rates is smaller for women than for men.

To summarize, in general, women work longer hours and are more likely to win projects they bid on compared to men, even if at lower hourly wages. At the same time, the correlation between success rates and hourly wages is of similar magnitude for men and women. In other words, women seem more successful than men along other dimensions than hourly pay. While our findings suggest that women are not necessarily at a disadvantage compared to men in all aspects of labor market outcomes, it does not definitively eliminate the possibility of self-fulfilling discrimination. To conclusively draw such a conclusion, we would require information on the success rates of women and men in counterfactual projects they might have pursued but chose not to. Unfortunately, our current data set does not allow for such analysis.

Table 7. Gender wage gap in application success rate

	Share of successful applications			
	(1)	(2)	(3)	(4)
Female	0.019***	0.019***	0.014***	0.011**
	(0.004)	(0.004)	(0.004)	(0.004)
Controls		✓	✓	✓
Skills			✓	✓
Desired worker experience				✓
Contract type				✓
Number of projects	22,056	22,056	22,056	22,056
Number of workers	10,326	10,326	10,326	10,326
Adjusted R^2	0.002	0.097	0.098	0.119
Share of females	49.81%	49.81%	49.81%	49.81%

Note: This table documents the gender wage gap in application success rates. The control variables in our analysis comprise the worker's educational degree, the year in which the project commenced, and employer country dummies. In addition, the regression model in Column 3 includes a measure of the market value of workers' skills derived from a machine learning model. In Column 4, we include additional controls for share of applications made to projects of into different contract types and the proportion of projects with different expected worker experience levels. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

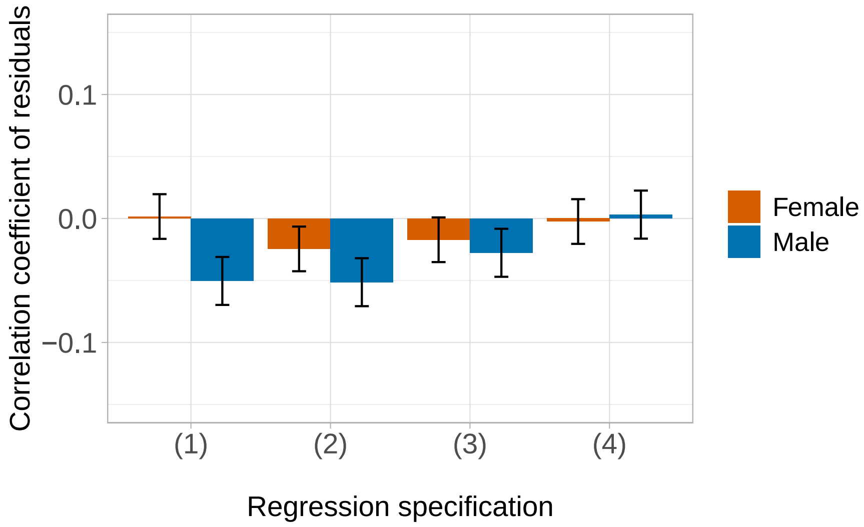

Figure 5. Correlation between wage and application success regression residuals

Note: This figure plots the correlation coefficients and their 95% confidence intervals of regression residuals from two regression models: with success rate and wage as the dependent variable, calculated separately for males and females. The numbers (1), (2), (3), and (4) correspond to specifications reported in columns of Table 7. See text for more details. Error bars represent 95% confidence intervals calculated as $+/-1.96 *$ st. dev..

5.2 Return to Work Experience and Learning by Doing

In the context of drivers on a ride-hailing platform, Cook et al. (2021) show that a major reason why men earn more than women is that men have accumulated more work experience than women. They further argue that this finding indicates that drivers who have completed more rides have become more productive due to the learning-by-doing effect. Even if our data are from a non-rideshare context, a similar mechanism could be at play in other gig economy platforms. We look at this by first plotting the gender distribution by experience quantiles in Figure 6.[18] We find that in general, women are slightly over-represented among workers with higher platform experience.

To further show that the gender gap is not driven by experience, we report the results from regression models where we have interacted worker experience levels — measured by completed projects at project start — with the gender dummy. If learning by doing drives the gender gap in hourly wages, we would expect wages to diverge as workers gain more work experience. We report the results in Table 8. In general, we find that while experience and hourly wages are slightly positively correlated, the difference in return to experience is small compared to the wage gap. Thus, we find no effect that learning by doing would tilt the gender wage gap in favor of men.

[18] Our measure of work experience does not account for projects that employers have marked as private. This could lead to an underestimation of actual experience. As a result, we might be overestimating the return to (observed) work experience.

Our finding is in contrast to the findings of Bertrand et al. (2010) and Goldin (2014), who study the incomes of MBAs working in financial and corporate sectors. They find that men earn more per hour than women because they work longer hours than women ("convex hours-pay relationship"). If anything, our results point in the opposite direction: women earn less per hour while working longer hours and more projects. This result aligns with Gomez-Herrera and Müller-Langer (2019), who also study data from an online freelancing platform. They find that men earn higher wages than women but do not find a gender wage gap in total earnings. Their interpretation of this is that women are able to compensate for their lower hourly wages by working more hours.

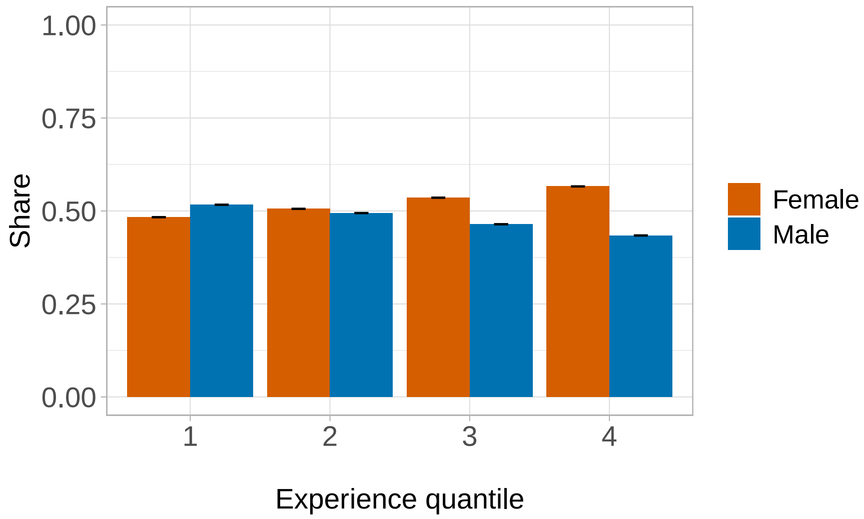

Figure 6. Gender distribution within experience quantiles

Note: This figure plots the gender distribution in different experience quantiles.

Table 8. Gender wage gap by different levels of experience

	Hourly wage (log)			
	(1)	(2)	(3)	(4)
Female	-0.358***	-0.328***	-0.134***	-0.100***
	(0.012)	(0.011)	(0.011)	(0.013)
Experience	0.003***	0.001*	0.002***	0.0002
	(0.001)	(0.001)	(0.001)	(0.0006)
Female × Experience	0.005***	0.003*	0.002	0.0002
	(0.002)	(0.002)	(0.001)	(0.0013)
Controls		✓	✓	✓
Skills			✓	✓
Desired worker experience				✓
Contract type				✓
Number of projects	45,107	45,107	45,107	27,698
Number of workers	23,425	23,425	23,425	13,267
Adjusted R^2	0.070	0.167	0.320	0.471
Share of females	50.61%	50.61%	50.61%	50.24%

Note: This table documents the interaction between gender wage gap and experience. Experience measured by number of initiated projects. The control variables in our analysis comprise the worker's educational degree, the year in which the project commenced, and employer country dummies. In addition, the regression model in Column 3 includes a measure of the market value of workers' skills derived from a machine learning model. In Column 4, we include additional controls for number of applications made by worker in 30 days, share of applications made to projects of into different contract types, and the proportion of projects with different expected worker experience levels. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Differences in Offline Opportunities, Risk-Aversion, and Overconfidence

Although our data on workers' outside options off-platform are limited, they allow for some speculation. As shown in Figure 1, men are more commonly found in high-skilled and high-paying projects than women. They also tend to ask for higher wages. This might be why they choose certain projects on the platform. Table 4 also points out that men are likelier to go for part-time and short-term projects requiring expert skills. Considering all this information, male workers on the platform seem to have better job chances off-platform as well. In simple terms, if men have the required skills and experience to be hired in expert jobs offline, they will likely find expert jobs online, too.

It is also worth noting that higher wage expectations and a higher prevalence of men in job applications demanding expert skills could suggest that men have systematically higher wage expectations, or less risk-aversion compared to women (see, e.g., Cortés et al., 2021; Roussille, 2021). If men consistently ask for higher wages and are more likely to apply for positions that require advanced skills, it may reflect a stronger belief in their abilities or, conversely, a higher tolerance for risk. This hypothesis gains support from the fact that women exhibit a higher success rate in their job applications compared to men, as illustrated in Table 7.

Given the observational nature of our data, it is challenging to disentangle the three explanations. However, it is worth noting that the gender gap persists, and its magnitude remains unchanged even among workers with long platform work histories. If we assume that workers with longer work experience possess a more realistic understanding of market demand for their skills and potential wages, it lends support to the hypothesis that men have better job opportunities outside the platform.[19] This external advantage for men might ultimately contribute to the gender wage gap observed on the platform.

6 Conclusion

In this paper, we use transaction-level data from a prominent U.S. online freelance labor market to investigate the gender wage gap. While some analysts posit that the gig economy's flexibility might help reduce the gender wage gap (see, e.g., Cook et al., 2021), our findings challenge this view. We find a substantial gender wage gap, likely primarily driven by choices reflecting the preferences and constraints faced by the workers. This suggests that transitioning to a more flexible, platform-mediated labor market would not necessarily narrow the gender wage gap in the broader economy.

We find that the gender wage gap between men and women is approximately 30%, a more significant disparity than the 20% to 1% typically seen in traditional labor markets (Eurostat, 2022). However, when we account for three factors: workers' skills, the types of projects they apply to, and their asking wages, we find that the unexplained gender wage gap disappears. The differences in application behavior are a central driver of the gender wage gap: Women predominantly apply for longer, full-time projects with lower skill and experience requirements, while men lean towards short-term expert gigs. The differences in application behavior suggest that platform work is more likely to be a full-time occupation for women. In contrast, men are more likely to use platform work to supplement their primary income. This also implies that women, on average, are more dependent on platform labor income than men. A decomposition of the gender wage gap further reveals that gender differences in asking wages can account for almost half of the gender wage gap, even among men and women with similar skills who apply for similar projects. It remains unclear, however, whether lower asking wages of women stem from

[19] We also highlight that the realized wages are visible to all labor market participants after a contract has been completed. Thus, it does not seem likely that workers would have systematically biased views of their earnings potential.

biased beliefs of market wages or result from an informed (rational) choice. Crucially, we find no evidence of gender discrimination. Regardless of gender, workers with identical skills applying for the same job and requesting the same wages receive equal pay.

Our results indicate that a labor market with no discrimination, high flexibility, and little uncertainty about worker productivity can still give rise to gender wage gaps. These differences likely are a result of gender differences in preferences or constraints. In particular, women seem to pay a price for striving for full-time predictable work arrangements. Instead of a preference for flexibility, the women in our data seem to give a higher weight for "schedule predictability" (Bolotnyy and Emanuel, 2022), which results in a wage penalty.

Our analysis focuses on a specific labor market, prioritizing one-off transactions over careers and organizational dynamics. Online freelancing represents a fraction of the labor market as a whole (Garin et al., 2022). Research suggests that gig work often serves as supplemental income (Farrell and Greig, 2016). Online freelancing also lacks many elements common in traditional, offline markets, such as team interactions, hierarchical management, or other social and work-related structures. While our methodology is transparent and can easily be applied in other contexts where data on workers' skills and application behavior are available, our results might have limited external validity beyond the online freelancing context. Additionally, although our dataset encompasses a host of project characteristics such as contract length, hourly demands, and employer expectations, it omits several aspects that might influence gender differences in hourly wages. These unobserved factors include deadline stringency, frequency of night work, work intensity, and the likelihood of re-hiring, among other things.

Our findings highlight how offline constraints and preferences influence online labor market outcomes. Addressing these issues requires interventions outside of the platform rather than regulating the contractual working arrangements of individual platform providers. Strategies to enhance women's participation in higher-paying sectors include promoting STEM education for females, and providing low-cost childcare for families. These measures may contribute to a more equitable distribution of opportunities and rewards from both traditional and online labor. More speculatively, our decomposition analysis demonstrates that the differences in asking wages can explain a significant portion of the wage differences between men and women. To the extent that biased beliefs about market wages drive differences in asking wages, informing workers about suitable pay may reduce gender wage gaps (Roussille, 2021).

Our research indicates that online freelancing markets provide equal pay for equal work. Yet, there are significant differences in how men and women use these markets and the jobs they choose. While advancements in technology have the potential to be transformative, they alone are not enough to close the gender wage gap.

CRediT Authorship Contribution Statement

Otto Kässi: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing – original draft, Writing – review & editing, Visualization, Supervision, Project administration.

Eliza Stenzhorn: Conceptualization, Methodology, Software, Validation, Formal analysis, Writing – original draft, Writing – review & editing, Visualization, Project administration.

Ole Teutloff: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing – original draft, Writing – review & editing, Visualization, Project administration.

Acknowledgements

We thank Fabian Stephany, Fabian Braesemann and Estrella Gomez-Herrera for useful comments. Valeriia Khlopchyk and Olavi Kylliäinen provided excellent research assistance. All remaining errors are our own.

Funding

This work was initiated when Otto Kässi was funded by the European Research Council (grant number 639652).

Competing interests

The authors declare no competing interests.

Declaration of the use of AI and AI-assisted technologies

Ole Teutloff has used ChatGPT to debug code. Otto Kässi has used ChatGPT for debugging code and generating BibTeX references, and Grammarly for proofreading and language checks. Eliza Stenzhorn has used ChatGPT to debug code and to improve readability and language.

References

Adams-Prassl, Abi, Kotaro Hara, Kristy Milland, and Chris Callison-Burch, "The gender wage gap in an online labor market: The cost of interruptions," *The Review of Economics and Statistics*, 02 2023, pp. 1–23.

Adda, Jérôme, Christian Dustmann, and Katrien Stevens, "The career costs of children," *Journal of Political Economy*, 2017, *125* (2), 293–337.

Angelov, Nikolay, Per Johansson, and Erica Lindahl, "Parenthood and the gender gap in pay," *Journal of labor economics*, 2016, *34* (3), 545–579.

Angrist, Joshua D and Jörn-Steffen Pischke, *Mostly harmless econometrics: An empiricist's companion*, Princeton University Press, 2009.

Azmat, Ghazala and Rosa Ferrer, "Gender gaps in performance: Evidence from young lawyers," *Journal of Political Economy*, 2017, *125* (5), 1306–1355.

Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet, "Gender differences in job search: Trading off commute against wage," *The Quarterly Journal of Economics*, 2021, *136* (1), 381–426.

Barth, Erling, Sari Pekkala Kerr, and Claudia Olivetti, "The dynamics of gender earnings differentials: Evidence from establishment data," *European Economic Review*, 2021, *134*, 103713.

Bavel, Jan Van, Christine R Schwartz, and Albert Esteve, "The reversal of the gender gap in education and its consequences for family life," *Annual Review of Sociology*, 2018, *44*, 341–360.

Bertrand, Marianne, Claudia Goldin, and Lawrence F Katz, "Dynamics of the gender gap for young professionals in the financial and corporate sectors," *American Economic Journal: Applied Economics*, 2010, *2* (3), 228–55.

Blau, Francine D and Lawrence M Kahn, "The feasibility and importance of adding measures of actual experience to cross-sectional data collection," *Journal of Labor Economics*, 2013, *31* (S1), S17–S58.

_ **and** _ , "The gender wage gap: Extent, trends, and explanations," *Journal of Economic Literature*, 2017, *55* (3), 789–865.

Bolotnyy, Valentin and Natalia Emanuel, "Why do women earn less than men? Evidence from bus and train operators," *Journal of Labor Economics*, 2022, *40* (2), 283–323.

Breiman, Leo, "Random forests," *Machine Learning*, 2001, *45*, 5–32.

Bureau of Labor Statistics, "*Employment projections data definitions*," https://www.bls.gov/emp/documentation/definitions.htm#education 2022. Last accessed on June 27, 2023.

Card, David and A Abigail Payne, "High school choices and the gender gap in STEM," *Economic Inquiry*, 2021, *59* (1), 9–28.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins, "Double/debiased machine learning for treatment and structural parameters," *The Econometrics Journal*, 01 2018, *21* (1), C1–C68.

Coate, Stephen and Glenn C Loury, "Will affirmative-action policies eliminate negative stereotypes?," *American Economic Review*, 1993, pp. 1220–1240.

Cook, Cody, Rebecca Diamond, Jonathan V Hall, John A List, and Paul Oyer, "The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers," *The Review of Economic Studies*, 2021, *88* (5), 2210–2238.

Corporaal, Greetje F and Vili Lehdonvirta, "Platform sourcing: How Fortune 500 firms are adopting online freelancing platforms," *University of Oxford*, 2017.

Cortés, Patricia, Jessica Pan, Laura Pilossoph, and Basit Zafar, "Gender differences in job search and the earnings gap: Evidence from business majors," Technical Report, National Bureau of Economic Research 2021.

Cullen, Zoe B, John Eric Humphries, and Bobak Pakzad-Hurson, "Gender and sorting in the on-demand economy," in "ASSA Annual Meeting, Atlanta, GA" 2018.

Cunningham, Scott, *Causal inference: The mixtape*, Yale University Press, 2021.

Eurostat, "*Gender pay gap statistics*," `https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Gender_pay_gap_statistics` 2022. Last accessed on January 27, 2023.

Farrell, Diana and Fiona Greig, "Paychecks, paydays, and the online platform economy," in "Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association," Vol. 109 JSTOR 2016, pp. 1–40.

Fluchtmann, Jonas, Anita Glenny, Nikolaj A Harmon, and Jonas Maibom, "The Gender Application Gap: Do men and women apply for the same jobs?," *IZA Discussion Paper No. 14906*, 2021.

Foong, Eureka, Nicholas Vincent, Brent Hecht, and Elizabeth M. Gerber, "Women (still) ask for less: Gender differences in hourly rate in an online labor marketplace," *Proceedings of the ACM on Human-Computer Interaction*, November 2018, *2* (CSCW).

Gallen, Yana, "Motherhood and the gender productivity gap," *Becker Friedman Institute for Research in Economics Working Paper*, 2018, (2018-41).

Garin, Andrew, Emilie Jackson, and Dmitri Koustas, "Is gig work changing the labor market? Key lessons from tax data," *National Tax Journal*, 2022, *75* (4), 791–816.

Gelbach, Jonah B, "When do covariates matter? And which ones, and how much?," *Journal of Labor Economics*, 2016, *34* (2), 509–543.

Ghani, Ejaz, William R Kerr, and Christopher Stanton, "Diasporas and outsourcing: Evidence from oDesk and India," *Management Science*, 2014, *60* (7), 1677–1697.

Goldin, Claudia, *Understanding the gender gap: An economic history of American women*, Oxford University Press, 1990.

_ , "A grand gender convergence: Its last chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

_ , *Career and family: Women's century-long journey toward equity*, Princeton University Press, 2021.

_ **and Lawrence F Katz**, "A most egalitarian profession: pharmacy and the evolution of a family-friendly occupation," *Journal of Labor Economics*, 2016, *34* (3), 705–746.

Gomez-Herrera, Estrella and Frank Müller-Langer, "Is there a gender wage gap in online labor markets? Evidence from over 250,000 projects and 2.5 million wage bill proposals," *CESifo Working Paper No. 7779*, 2019, *19-07.*

Herrmann, Andrea M, Petra M Zaal, Maryse MH Chappin, Brita Schemmann, and Amelie Lühmann, " 'We don't need no (higher) education' – How the gig economy challenges the education-income paradigm," *Technological Forecasting and Social Change*, 2023, *186*, 122136.

Kahn, Shulamit and Donna Ginther, "Women and STEM," Technical Report, National Bureau of Economic Research 2017.

Kässi, Otto and Vili Lehdonvirta, "Online labour index: Measuring the online gig economy for policy and research," *Technological Forecasting and Social Change*, 2018, *137*, 241–248.

_ **and** _ , "Do microcredentials help new workers enter the market? Evidence from an online labor platform," *Journal of Human Resources*, 2022, pp. 0519–10226R3.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard, "Children and gender inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 2019, *11* (4), 181–209.

_ **,** _ **, Johanna Posch, Andreas Steinhauer, and Josef Zweimüller**, "Child penalties across countries: Evidence and explanations," in "AEA Papers and Proceedings," Vol. 109 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2019, pp. 122–126.

Lehdonvirta, Vili, Otto Kässi, Isis Hjorth, Helena Barnard, and Mark Graham, "The global platform economy: A new offshoring institution enabling emerging-economy microproviders," *Journal of Management*, 2019, *45* (2), 567–599.

Lundberg, Shelly J and Richard Startz, "Private discrimination and social intervention in competitive labor market," *American Economic Review*, 1983, *73* (3), 340–347.

Mullen, Lincoln, "*Gender: Predict gender from names using historical data.* R-package version 0.6.0.," https://github.com/lmullen/gender 2021. Last accessed January 27, 2023.

Niederle, Muriel and Lise Vesterlund, "Gender and competition," *Annual Review of Economics*, 2011, *3* (1), 601–630.

Olivetti, Claudia and Barbara Petrongolo, "The evolution of gender gaps in industrialized countries," *Annual Review of Economics*, 2016, *8*, 405–434.

Pew Research Center, "*10 facts about today's college graduates*," https://www.pewresearch.org/short-reads/2022/04/12/10-facts-about-todays-college-graduates/ 2022. Last accessed on September 5, 2023.

Roussille, Nina, *The central role of the ask gap in gender pay inequality*, University of California, Berkeley, 2021.

Schofer, Evan and John W Meyer, "The worldwide expansion of higher education in the twentieth century," *American Sociological Review*, 2005, *70* (6), 898–920.

Stephany, Fabian and Ole Teutloff, "What is the price of a skill? The value of complementarity," *Research Policy*, 2024, *53* (1), 104898.

_, Otto Kässi, Uma Rani, and Vili Lehdonvirta, "Online labour index 2020: New ways to measure the world's remote freelancing market," *Big Data & Society*, 2021, *8* (2), 20539517211043240.

A Online Appendix

A.1 Descriptive Statistics

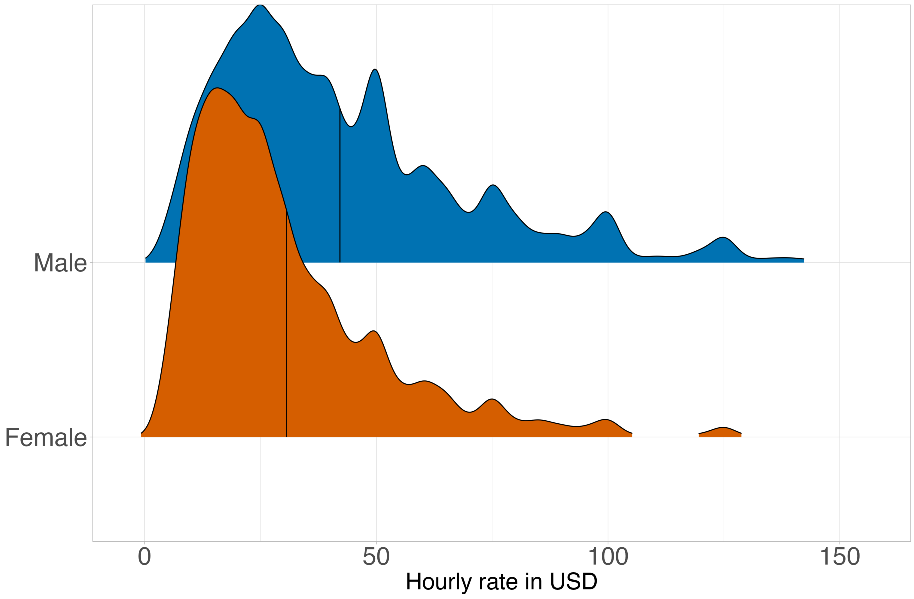

Appendix Figure A1. Distribution of hourly wage by gender

Note: The distribution of hourly wages in USD by gender. The black vertical line represents the mean.

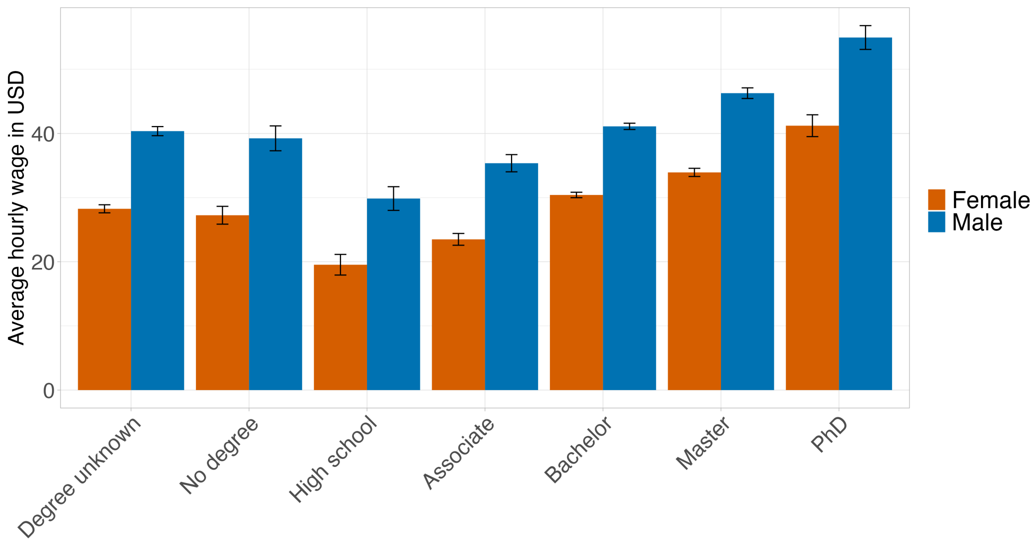

Appendix Figure A2. Hourly wage by education and gender

Note: The average hourly wage in USD by level of formal education and by gender. Error bars represent 95% confidence intervals calculated as $+/-1.96 *$ st. error.

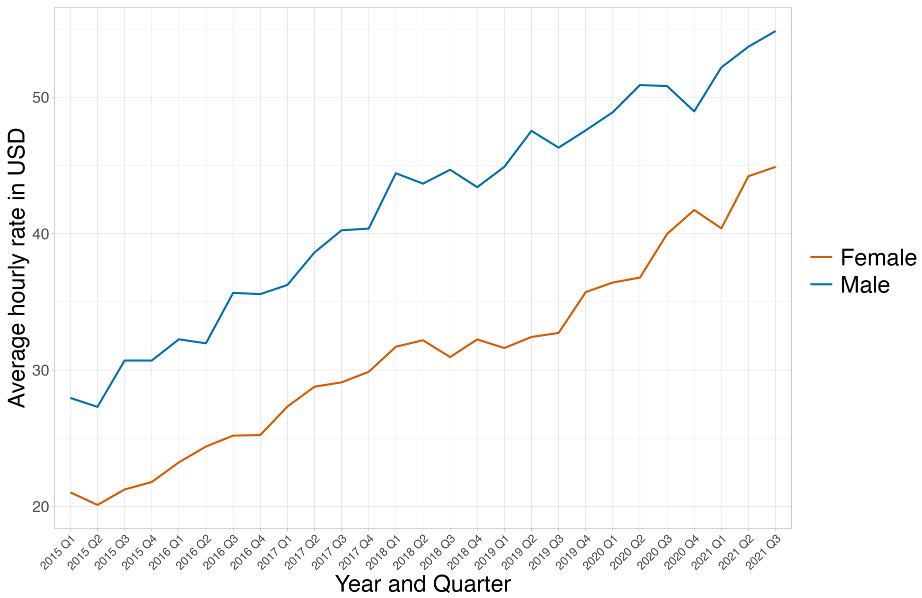

Appendix Figure A3. Hourly wage by gender over time

Note: The Hourly wage in USD by gender over time as quarterly averages.

A.2 Learning the Value of Worker Skills

In the following, we provide details on how we use machine learning to learn the value of skills. To begin, we randomly partition the data, encompassing 45,581 projects with skill requirements, into 80% training and 20% test set.[20] We train and compare three different shallow machine-learning algorithms: Elastic Net, XGBoost and Random Forest. For each, we perform hyper-parameter tuning using 10-fold cross-validation on the training set. Table A1 reports the prediction performance on the test set for the models with the respective best hyper-parameters. Random Forest shows the best performance. Therefore, we use Random Forest in our subsequent analysis.

Appendix Table A1. ML-Model comparison

Model	R-Squared on test set
Elastic Net	0.22
XGBoost	0.24
Random Forest	0.27

We proceed by predicting the value in hourly USD (log) of all projects based solely on their skill requirements. Figure A4 shows how our predictions compare to the observed hourly rates. As we can see, Random Forest performs well for values in the center of the distribution. However, our model has less predictive power in the tails of the distribution. One reason might be that the model does not have access to information on other project characteristics that matter for hourly wages such as for example desired worker experience.

[20]Note that the data set used for our analyses comprises 45,107 observations. This reduction stems from our data merging process; 12,053 observations lacked either skill or project information and were therefore excluded from the final data set.

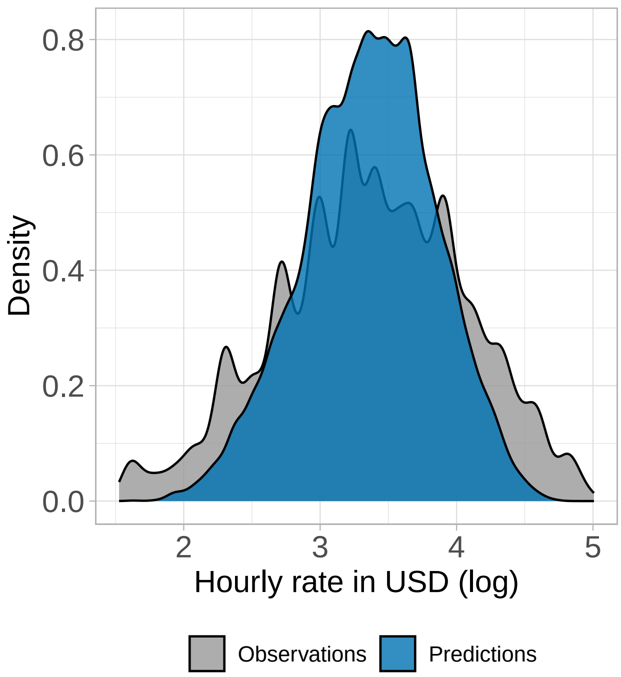

Appendix Figure A4. Comparing predicted and observed hourly rates of projects

Note: This figure compares the predictions of the Random Forest model to the actually observed hourly rates of all projects with skill requirements. Hourly rates are in USD (log).

Using the same Random Forest model, we predict the value of workers' skill sets. There exists no ground truth for the value of workers' skill sets against which we could compare our predictions. Figure 3 displays the predicted value of workers' skill sets by gender. As we can see, our model predicts significantly lower skill values for female workers.

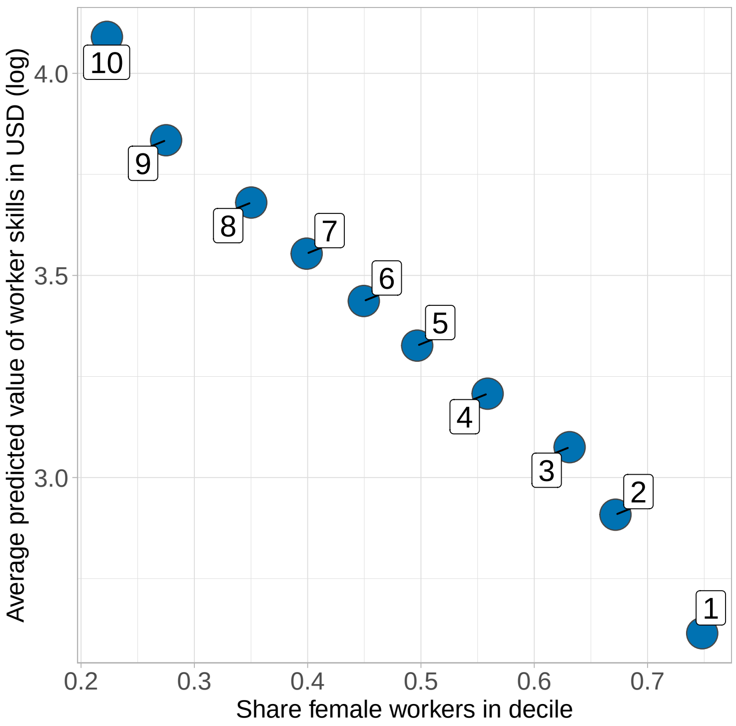

Appendix Figure A5. Average hourly wage and share of female workers in each skill value decile

Note: The share of female workers in each decile of the skill value predictions is plotted gaianst the average hourly wage in USD (log) in that decile.

Female workers offer skill sets of systematically lower value than their male counterparts. Figure A5 illustrates the share of female workers in each decile of the predicted skill value variable. We can observe a clear negative relationship between the average skill value by decile and the share of female workers.

A.3 Descriptive Statistics for Hourly and Fixed Projects

Appendix Table A2. Summary statistics for hourly and fixed price projects

	Hourly projects					Fixed price projects				
	Mean	St. Dev.	Min	Median	Max	Mean	St. Dev.	Min	Median	Max
Wage	36.590	25.086	4.600	30	149.500	332.726	848.777	5.010	100	13,568
Worker characteristics										
PhD	0.044	0.206	0	-	1	0.050	0.219	0	-	1
Master	0.201	0.401	0	-	1	0.200	0.400	0	-	1
Bachelor	0.460	0.498	0	-	1	0.459	0.498	0	-	1
Associate	0.054	0.226	0	-	1	0.052	0.221	0	-	1
High school	0.019	0.138	0	-	1	0.017	0.130	0	-	1
No degree	0.020	0.141	0	-	1	0.021	0.143	0	-	1
Degree unknown	0.202	0.401	0	-	1	0.201	0.401	0	-	1
Main project characteristics										
Accounting & consulting	0.050	0.217	0	-	1	0.019	0.137	0	-	1
Admin support	0.095	0.293	0	-	1	0.048	0.213	0	-	1
Customer service	0.015	0.121	0	-	1	0.003	0.055	0	-	1
Data science & analytics	0.033	0.178	0	-	1	0.020	0.139	0	-	1
Design & creative	0.172	0.377	0	-	1	0.240	0.427	0	-	1
Engineering & architecture	0.029	0.168	0	-	1	0.020	0.139	0	-	1
IT & networking	0.031	0.172	0	-	1	0.012	0.110	0	-	1
Legal	0.023	0.148	0	-	1	0.023	0.150	0	-	1
Sales & marketing	0.147	0.354	0	-	1	0.062	0.242	0	-	1
Translation	0.021	0.142	0	-	1	0.037	0.189	0	-	1
Web, mobile & software development	0.137	0.344	0	-	1	0.093	0.291	0	-	1
Writing	0.248	0.432	0	-	1	0.423	0.494	0	-	1
Number of projects	45,107					127,463				
Number of workers	23,425					31,353				
Share of females	50.61%					50.74%				

Note: The value presented are based on U.S. online workers who completed at least one project between the years 2015 to 2021. We report the descriptive statistics for hourly-priced projects and lump sum priced projects. *Wage* signifies the hourly wage in projects with hourly pricing; for projects under lump sum contracts, it refers to the total lump sum earnings.

A.4 Descriptive Statistics: Job-Applicant Sample

Appendix Table A3. Basic summary statistics: Sample on job-search behavior

	Male		Female		Difference in means
	Mean	Median	Mean	Median	(male – female)
Worker characteristics					
Project length (hours)	27.256 (59.368)	7.830	35.642 (93.631)	8.170	-8.386***
PhD	0.045 (0.208)	-	0.042 (0.200)	-	0.003
Master	0.193 (0.394)	-	0.207 (0.405)	-	-0.014**
Bachelor	0.456 (0.498)	-	0.473 (0.499)	-	-0.017*
Associate	0.050 (0.217)	-	0.062 (0.240)	-	-0.012***
High school	0.018 (0.131)	-	0.020 (0.141)	-	-0.002
No degree	0.022 (0.147)	-	0.026 (0.158)	-	-0.004
Degree unknown	0.217 (0.412)	-	0.170 (0.376)	-	0.047***
Main project categories					
Accounting & consulting	0.055 (0.228)	-	0.049 (0.216)	-	0.006**
Admin support	0.035 (0.184)	-	0.155 (0.362)	-	-0.120***
Customer service	0.006 (0.080)	-	0.021 (0.143)	-	-0.015***
Data science & analytics	0.046 (0.210)	-	0.014 (0.117)	-	0.032***
Design & creative	0.168 (0.374)	-	0.168 (0.373)	-	-
Engineering & architecture	0.037 (0.188)	-	0.012 (0.110)	-	0.025***
IT & networking	0.053 (0.224)	-	0.008 (0.087)	-	0.045***
Legal	0.032 (0.175)	-	0.020 (0.140)	-	0.012***
Sales & marketing	0.174 (0.380)	-	0.153 (0.360)	-	0.021***
Translation	0.009 (0.096)	-	0.027 (0.161)	-	-0.018***
Web, mobile, & software development	0.212 (0.409)	-	0.058 (0.233)	-	0.154***
Writing	0.172 (0.377)	-	0.316 (0.465)	-	-0.144***
Number of projects	14,621		13,077		
Number of workers	6,602		6,665		
Share of females			50.24%		

Note: The values presented are based on U.S. online workers who completed at least one project between the years 2015 and 2021. Standard deviation in parentheses. In Column 6, we test the statistical significance of the differences in means between female and male workers using two-sample t-tests. The significance levels are indicated by: * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

A.5 Past Application Behavior: Different Time Windows

Appendix Table A4. Gender wage gap conditional on skills and application behavior

	Hourly wage (log)								
Panel A	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female	-0.341***	-0.313***	-0.123***	-0.078***	-0.076***	-0.068***	-0.069***	-0.025***	0.002
	(0.015)	(0.015)	(0.014)	(0.011)	(0.011)	(0.011)	(0.011)	(0.008)	(0.008)
Skills			0.692***	0.504***	0.499***	0.475***	0.476***	0.168***	0.092***
			(0.016)	(0.014)	(0.014)	(0.013)	(0.014)	(0.012)	(0.012)
Mean value of past applications									0.278***
									(0.014)
Controls		✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓
Contract type					✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓
Number of past applications							✓	✓	✓
Asking wage (log)								✓	✓
Number of projects	33,160	33,160	33,160	33,160	33,160	33,160	33,160	33,160	33,160
Number of workers	15,598	15,598	15,598	15,598	15,598	15,837	15,598	15,598	15,598
Adjusted R^2	0.062	0.178	0.340	0.495	0.497	0.503	0.503	0.671	0.681
Share of females	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%

	Hourly wage (log)								
Panel B (Table A4 continued)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female	-0.338***	-0.310***	-0.122***	-0.068***	-0.065***	-0.057***	-0.058***	-0.019***	0.012
	(0.014)	(0.014)	(0.013)	(0.010)	(0.010)	(0.010)	(0.010)	(0.007)	(0.007)
Skills			0.685***	0.471***	0.464***	0.438***	0.438***	0.148***	0.064***
			(0.015)	(0.013)	(0.013)	(0.012)	(0.013)	(0.011)	(0.011)
Mean value of past applications									0.312***
									(0.015)
Controls		✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓
Contract type					✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓
Number of past applications							✓	✓	✓
Asking wage (log)								✓	✓
Number of projects	37,419	37,419	37,419	37,419	37,419	37,419	37,419	37,419	37,419
Number of workers	17,681	17,681	17,681	17,681	17,681	17,681	17,681	17,681	17,681
Adjusted R^2	0.061	0.178	0.336	0.510	0.512	0.518	0.518	0.678	0.688
Share of females	50.12%	50.12%	50.12%	50.12%	50.12%	50.12%	50.12%	50.12%	50.12%

Note: This table presents the gender wage gap conditional on application behavior. Panel A provides the results when considering the workers' application behaviors over the past 90 days at the time of data collection, and Panel B provides the results for the past 365 days. Column 1 shows the raw gender wage gap. In Column 2, we control for conventional controls (project start year, employer country dummies, and worker education). Column 3 presents the results from the regression specification where we account for the market value of workers' skills and control variables (project start year, employer country dummies, and worker education). In Columns 2 to 7, we progressively incorporate controls for workers' job application behavior. *Number of past applications* encompasses the number of all applications and the number of applications to fixed price projects. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels are indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.6 Sample Restriction: U.S. Employers Only

Appendix Table A5. Gender wage gap conditional on skills and application behavior

	Hourly wage (log)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female	-0.360***	-0.333***	-0.139***	-0.099***	-0.098***	-0.091***	-0.092***	-0.038***	-0.010
	(0.017)	(0.017)	(0.015)	(0.013)	(0.013)	(0.013)	(0.013)	(0.009)	(0.009)
Skills			0.695***	0.536***	0.530***	0.512***	0.513***	0.184***	0.110***
			(0.018)	(0.016)	(0.016)	(0.016)	(0.016)	(0.014)	(0.014)
Mean value of past applications									0.258***
									(0.014)
Controls		✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓
Contract type					✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓
Number of past applications							✓	✓	✓
Asking wage (log)								✓	✓
Number of projects	20,614	20,614	20,614	20,614	20,614	20,614	20,614	20,614	20,614
Number of workers	10,925	10,925	10,925	10,925	10,925	10,925	10,925	10,925	10,925
Adjusted R^2	0.070	0.176	0.342	0.474	0.476	0.481	0.481	0.663	0.673
Share of females	49.89%	49.89%	49.89%	49.89%	49.89%	49.89%	49.89%	49.89%	49.89%

Note: This table presents the gender wage gap conditional on application behavior. Column 1 shows the raw gender wage gap. In Column 2, we control for conventional controls (project start year, employer country dummies, and worker education). Column 3 presents the results from the regression specification where we account for the market value of workers' skills and control variables (project start year, employer country dummies, and worker education). In Columns 2 to 7, we progressively incorporate controls for workers' job application behavior. *Number of past applications* encompasses the number of all applications and the number of applications to fixed price projects. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels are indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.7 Skill Dummies

Appendix Table A6. Gender wage gap conditional on skills and application behavior

	Hourly wage (log)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female	-0.355***	-0.327***	-0.081***	-0.059***	-0.059***	-0.056***	-0.057***	-0.030**	-0.019*
	(0.016)	(0.016)	(0.014)	(0.012)	(0.012)	(0.012)	(0.012)	(0.009)	(0.009)
Mean value of past applications									0.194***
									(0.012)
Controls		✓	✓	✓	✓	✓	✓	✓	✓
Skill dummies			✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓
Contract type					✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓
Number of past applications							✓	✓	✓
Asking wage (log)								✓	✓
Number of projects	27,698	27,698	27,698	27,698	27,698	27,698	27,698	27,698	27,698
Number of workers	13,267	13,267	13,267	13,267	13,267	13,267	13,267	13,267	13,267
Adjusted R^2	0.067	0.179	0.506	0.581	0.582	0.584	0.584	0.711	0.715
Share of females	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%

Note: This table presents the gender wage gap conditional on application behavior. Column 1 shows the raw gender wage gap. In Column 2, we control for conventional controls (project start year, employer country dummies, and worker education). Column 3 presents the results from the regression specification where we account for the market value of workers' skills and control variables (project start year, employer country dummies, and worker education). In Columns 2 to 7, we progressively incorporate controls for workers' job application behavior. *Number of past applications* encompasses the number of all applications and the number of applications to fixed price projects. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels are indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

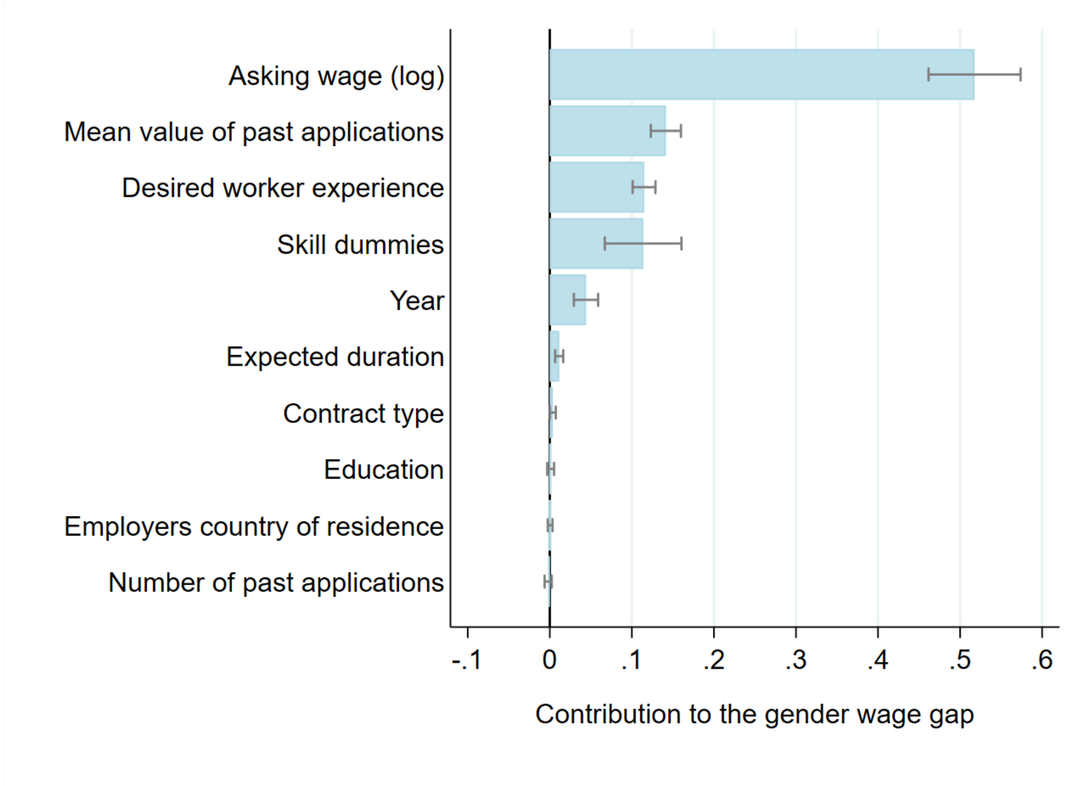

Appendix Figure A6. Gelbach decomposition

Note: We apply the method from Gelbach (2016) to show how much each factor contributes to the gender wage gap. These factors are education, preferred worker experience, type of contract, expected duration, number of past applications, the wage workers ask for, skills, average skill value of projects workers applied to over 30 days, employer's country, and year dummies. Error bars represent 95% confidence intervals calculated as $+/-1.96 * \text{st. dev.}$