// NO.25-016 | 03/2025

DISCUSSION PAPER

// OLE TEUTLOFF, ELIZA STENZHORN, AND OTTO KÄSSI

Skills, Job Application Strategies, and the Gender Wage Gap: Evidence From Online Freelancing

Leibniz Association

ZEW

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

Ole Teutloff† Eliza Stenzhorn‡ Otto Kässi§

March 3, 2025

Abstract

This paper examines how worker skills and job application behavior contribute to the gender wage gap on a major online freelancing platform. We observe significant occupational sorting by gender, with women over-represented in lower-paying project categories and tending to earn less than men even within the same categories. The unexplained gender wage gap conditional on education is initially 39.9%, but it narrows to under 2% when accounting for differences in human capital and application strategies. Our analysis shows that application behavior, including job preferences and asking wages, is the primary factor, explaining up to 90% of the wage gap. We also find that women work on longer projects and achieve higher application success rates than men, which helps offset lower hourly earnings by accumulating more work hours. While men have slightly greater platform and traditional work experience it has minimal impact on wage outcomes. These findings suggest that the gender wage gap on the platform primarily reflects distinct usage patterns between men and 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

*Author names are listed in reverse alphabetical order. All authors have contributed equally to the paper. We thank Fabian Braesemann, Estrella Gomez-Herrera, Fabian Stephany, Michael Ward, and the participants at EEA-ESEM 2024, EALE 2024, the ZEW ICT Conference 2024, and the University of Luxembourg Gender and Economics Workshop for their valuable feedback. Valeriia Khlopchyk and Olavi Kylliäinen provided excellent research assistance. All remaining errors are our own.

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

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

§Etla Economic Research Helsinki, Finland; Oxford Internet Institute, University of Oxford, United Kingdom. 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). Similarly, in 2022, the median hourly earnings for female wage and salary workers were $17.18, compared to $19.70 for male workers, revealing a $2.52 per hour difference (Statista Research Department, 2024a,b). 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., 2019). 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 platform that allows employers to divide work into distinct tasks, facilitating both precise oversight of worker outputs and a high degree of interchangeability between workers. 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. Because online freelancing offers flexible scheduling, remote work, and relatively standardized tasks, many common explanations for the gender wage gap – such as penalties for constrained schedule – should be smaller. Yet, despite these features, we observe a considerable gender gap in wages. Conditional on education, the gap in hourly wages is higher than in traditional labor markets.

We make three main contributions to the literature on gender wage gaps and digital labor markets. First, we leverage granular data on worker skills, and introduce a novel Random Forest based approach to measure the market value of workers' skills. Second, we document that workers' application behavior – including which projects they choose to pursue and the wages they ask for – drives a large share of the gender wage gap. Third, we demonstrate that even in a setting where work is fully remote and flexible, the gap is not eliminated; if anything, it can exceed that in traditional labor markets once we account for workers' education.

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, short-duration tasks. In ridesharing, customer-facing interactions may lead to customer-driven discrimination, but there is minimal scope for skill differentiation or bargaining. Similarly, in clickwork, tasks are often anonymous, and gender is often unobservable, eliminating the possibility of wage negotiations. 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.

There are a few reasons why studying online platform work is a particularly interesting setting for studying the gender wage gap. First, in addition to realized wages, the platform data include comprehensive information on job requirements, such as desired worker experience, expected project duration, weekly hours, and skill prerequisites. The data also capture 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, the data encompass 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. In summary, the data allow for an in-depth analysis of how men and women differ in what skills they possess, what jobs they apply for, and, conditional on applications, what they are paid for their work. Second, the wage structure on the platform does not have an explicit return to working long hours (such as over-time pay), or return to tenure. This structure likely reduces any job flexibility penalty that could arise from a convex hours-earnings relationship Goldin and Katz (2016). Third, platform work can be done from home, which could help even out gender discrepancies due to care responsibilities. Finally, in the platform context, realized wages become visible to all labor market participants after a contract is completed, reducing information asymmetries and increasing transparency regarding wages.

Our analysis proceeds in three main steps. First, we measure the raw gender wage gap, controlling for education level and field. Second, we employ a Random Forest (Breiman, 2001) to map workers' binary skill tags into a single measure of skill value, capturing how project skill requirements relate to wages. This allows us to estimate expected wages based on each worker's specific skill set. Finally, we analyze differences in application behavior, focusing on project attributes such as duration, workload, contract type, and required experience.

We find that women's raw hourly wages are, on average, 33.6 log-points (about 30%) lower than men's.[1] Accounting for differences in human capital reduces this gap to 12.6%, while further controlling for application strategies lowers it to 1.8%. We follow the linear decomposition method of Gelbach (2016) to decompose the gender wage gap into the proportion explained by human capital and application strategies respectively. While our decomposition rests on the usual strong assumptions, we find that the differences in application strategies drive the majority of the gender wage gap. Approximately 90% of the explained wage gap can be explained by different application strategies. The remaining 10% of the gender wage gap is attributed to differences in human capital. The difference in education levels and fields of education has minimal impact on the gender wage gap. Instead, the differences in skills account for approximately 10% of the gender wage gap.

A natural follow-up question then is *why* do men and women apply for different jobs? Our data lacks quasi-experimental variation, so it 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, our findings are consistent with compensatory behavior, where women might intentionally apply for less challenging jobs, or ask for lower wages to increase their project length or probability of being hired. Moreover, we find evidence against the hypothesis that the gender differences in wages stem from disparities in return to experience. We generally find that the return to work experience is broadly similar between men and women and that gender differences in work experience are not large enough to explain the differences in hourly wages.

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.[2] Goldin (2014) suggests that women have different preferences regarding workplace arrangements due to motherhood, which results in 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

[1]Our primary measure is the log-hourly wage. The gap in log-points refers to the difference in the natural log of hourly wages between men and women.

[2]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).

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, 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 several 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% among 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 applied-for jobs affect the gender wage gap.

In contrast to ridesharing and clickwork, jobs on online freelancing platforms are longer, more diverse, and complex and generally require higher skills. Chan and Wang (2018), Foong et al. (2018) and Gomez-Herrera and Mueller-Langer (2024) focus on gender differences in labor market outcomes in online freelancing. Across all of these studies, a consistent finding is that women and men engage differently with online labor markets. Chan and Wang (2018) focus on *hiring biases* in online labor markets, and provide evidence that gender influences the likelihood of being hired, particularly in feminine-typed jobs. However, their study focuses on hiring probabilities and does not extend to analyzing *wage outcomes after hiring* or the role that skills and job choices play in wage differences. In contrast, our study addresses wage-setting mechanisms and shows that gender wage differences in our platform are largely explained by skill sets and application strategies.

Using a global dataset, Gomez-Herrera and Mueller-Langer (2024) identify a 16.8% gender wage gap that can be entirely attributed to the bidding behavior of workers. According to their results, women tend to bid more competitively (i.e. they are willing to work for lower pay) than men, and bid for lower-paying projects. Gomez-Herrera and Mueller-Langer (2024) 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 related to differences in skills, job requirements, or other factors. In contrast, our approach allows us to disentangle how workers' skills and application strategies beyond wage bids contribute to the wage gaps.

Foong et al. (2018), on the other hand, focus 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 and results are 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 provide additional evidence and discuss 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.[3]

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) full flexibility concerning when, how much, and for whom to work. The

[3]For further details, see Ghani et al. (2014); Kässi and Lehdonvirta (2024); Lehdonvirta et al. (2019).

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.[4]

Projects on the platform are entirely virtual and 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 weekly contract hours (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 contain information on completed project transactions and worker profiles. The project-level data include 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 are a major advantage for studying gender wage disparities. Our data allow for a fine-grained analysis of how men and women differ in what skills they possess, which projects they apply for, and what they are paid for their work, conditional on project- and worker-level control variables, such as project duration, weekly contract hours, 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 were 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'

[4] It is beyond the scope of this paper to determine 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.

work histories beyond 2017. Collecting information on these projects allows us to extend our data set to transactions to years prior to 2017.[5]

The online freelancing platform is U.S.-based but global in scope. Most of the workers are based outside the United States, mainly in India, Pakistan, the Philippines, and eastern Europe (Stephany et al., 2021). However, to minimize the unobserved heterogeneity that might affect the gender wage gap and complicate the interpretation of empirical results, 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 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 number of weekly hours for the contract, employers select a broad and specific project category. There are 12 broad project categories (e.g., *writing*, *design & creative*, and *sales & marketing*) and about 90 specific project categories (e.g., *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*). 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

[5]Uncompleted projects are unlikely to systematically differ from completed projects in ways that would introduce bias to our analysis. Project completion is typically driven by the nature of the work such as its complexity or ongoing requirements, rather than characteristics directly related to gender or wages.

[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 Appendix A.1, Table A2).

more than six months) and workload (full-time vs. part-time). For completed projects, our data also contain 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 visible to all labor market participants once the project is completed.

When creating their freelancer profiles, 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. Workers list 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% 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.[10]

To measure work experience, we use information from worker profiles on offline work experience and the number of projects completed on the platform. We convert the unstructured list of offline work experience into a single numeric value representing the

[8]In particular, we cannot rule out that women might systematically under-report their skills because of, for instance, lower self-confidence (see, e.g., Croson and Gneezy, 2009; Detilleux and Deschacht, 2024).

[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]To validate the accuracy of the gender predictions from the *gender* R package, we downloaded the profile pictures of 300 random U.S.-based workers from our analysis sample. We then manually classified these profile pictures as either male or female. Comparing the package's predictions to our manually classified "ground truth" data revealed that predictions based on first names were approximately 97% accurate.

number of years since a worker's first job – in other words, their time in the labor market. Unfortunately, worker profiles do not contain age information.

Contrary to gender and age, 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). Additionally, we use GPT-4o to classify the educational description according to the Classification of Instructional Programs taxonomy of academic disciplines at institutions of higher education in the United States. If workers list multiple degrees, we consider only the highest degree.[11]

Transactions on the platform begin with employers posting a job and workers submitting applications that include wage bids. Employers can then interview applicants and invite them to work on a project. While we observe the asking wage that workers specify on their public profiles, we do not have access to project-level wage bids. Platform users, employers, and workers cannot view this information for other workers and projects. In a minority of cases, employers directly invite a specific worker to a job without a public post; these instances are not included in our data set, as our data collection starts with publicly posted jobs.

Most completed projects contain information on applicants. Matching workers to their applications allows us to build a detailed application history for each worker. However, the application data are not complete, as some projects do not contain information on applicants. There are two reasons for this. First, in cases where a project was deleted before being filled, we do not know the identity of the applicants. Second, some workers are invited to jobs directly without applying (if a worker has completed a job without applying, this project would be included when calculating their work experience, but excluded from the analysis sample). 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[12]).

In contrast to project information, worker profiles are subject to change. As they complete projects, workers accumulate experience and adapt their asking wages, 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.

[11]We validate both approaches with a random sample of 100 hand-labeled workers, achieving close to 100% accuracy. For details on the prompting of GPT-4o see Appendix Section A.2.

[12]In comparison, data from the U.S. labor market indicate that jobseekers typically submit around 26 applications per job offer (Dalton and Groen, 2020).

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.[13]

3 Summary Statistics and the Raw Gender Wage Gap

After having described our data collection process, we move to quantifying the gender wage gap and studying how worker-level characteristics affect it.

3.1 Describing the Raw Gender Wage Gap

In this section, we present summary statistics, offer background information on the workers, and describe the types of projects they engage in. We observe a considerable gender wage gap in our data set 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.-based 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.[14] [15] Although women earn lower hourly wages, their projects tend to be longer: men average 29 hours per project, while women average 37 hours. However, this difference appears driven by a small share of very long projects, as the median project length is approximately eight hours for both men and 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 *Writ-*

[13] 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. We also show, that our results remain almost identical when we limit our attention to new workers, whose profiles are less likely to have changed (see Appendix Section A.8).

[14] We include both the hourly wage and log-hourly wage for illustration, but use only the log-hourly wage as the dependent variable in our regression analyses.

[15] We plot the hourly wage distributions by gender in Appendix Section A.1.

ing, 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 education information. Workers' fields of education are summarized in Appendix Table A1. Men are more likely to have an educational background in technical fields such as Computer and Information Sciences and Engineering, while women are more commonly found in Education, Health Professions, and Psychology. This gendered distribution across fields is also reflected in the gender distribution in different project types.

Table 1. Basic summary statistics

	Male		Female		Difference in means
	Mean	Median	Mean	Median	(female – male)
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***
Web, mobile & software development	0.213 (0.409)	-	0.056 (0.230)	-	-0.157***
Translation	0.011 (0.106)	-	0.030 (0.172)	-	0.019***
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. Information on the workers' field of education can be found in Table A1 in the Appendix. 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 $p < 0.01$.

Women not only tend to complete projects in different categories but also in 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, one important reason that women earn less than men is that they work in project categories that tend to pay less. According to Figure 1, the average hourly wage for workers in *IT & networking* is \$56, but the share of women who work on 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%.

Additionally, 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.*

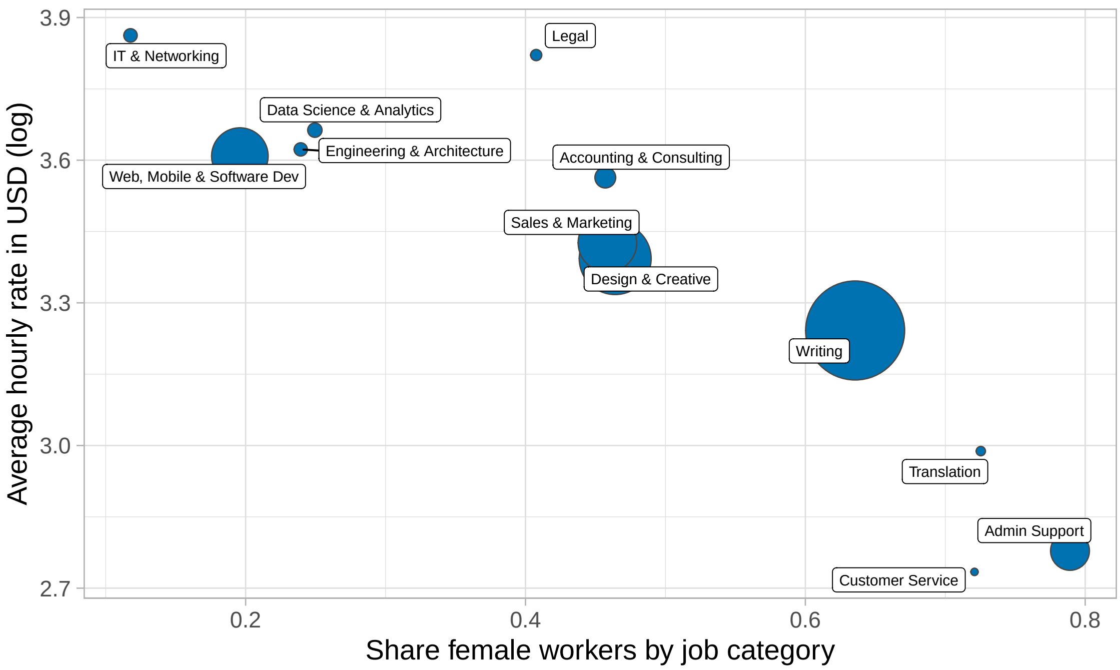

Figure 1. Average hourly wage and share of female workers in each project category

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. 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 gender norms, discrimination, or some combination of individual preferences or constraints. When project category choice itself is an outcome of gender, controlling for it effectively controls for part of the pathway through which gender influences wages, potentially leading to an underestimation of the gender wage gap.

In contrast, controlling for workers' application behavior is not subject to the "bad controls" critique. Although differences in application behavior might also arise from gendered norms, discrimination, or other external factors, application decisions clearly precede wages in the causal pathway from gender to wages. By conditioning on application behavior, we can decompose the wage gap into components that stem from differences in human capital and differences arising from application behavior, while remaining silent on the underlying reasons for these differences in application behavior. Thus, this approach allows us to distinguish wage disparities that are due to differential application strategies from those that result from workers' skills and qualifications.[16]

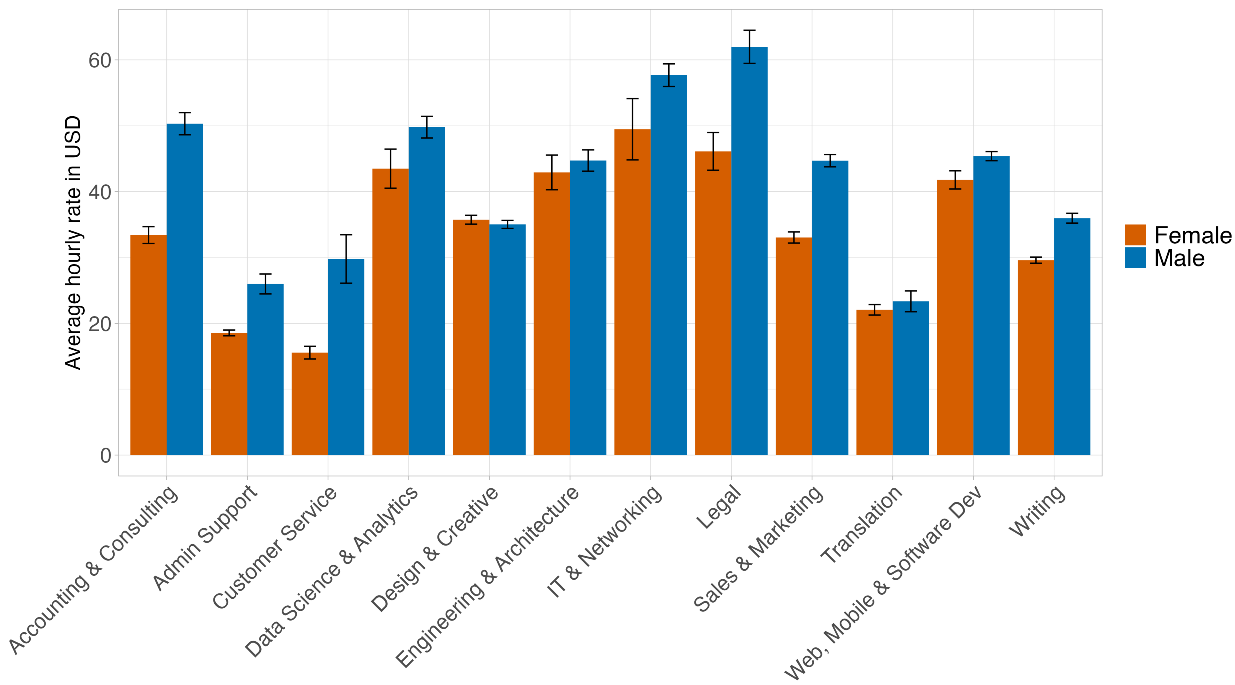

Figure 2. Average hourly wage in USD by gender and project category

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

Before moving on to the decomposition of the gender wage gap, we quantify the *raw* 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. We employ a standard wage regression model:

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

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

[16]We acknowledge the potential caveat of this approach. Specifically, one might argue that if societal norms or any form of discrimination shape the acquisition of certain skills based on gender, then controlling for these skills could inadvertently obscure specific wage disparities. While we will discuss this in detail later, it is worth noting that our primary finding suggests the gender wage gap on the platform is predominantly influenced by differences in application behavior that seemingly stem from distinct offline constraints, opportunities, or preferences.

the worker's education level and field, the project's starting year, and dummy variables for the employer's country. The coefficient attached to the gender dummy variable quantifies the difference in hourly wages between men and women in log-points, the gender wage gap. We account for potential correlations in unobservables within workers by clustering standard errors at the worker level.

Table 2. Raw gender wage gap

	Hourly wage (log)			
	(1)	(2)	(3)	(4)
Female	-0.336***	-0.290***	-0.347***	-0.300***
	(0.012)	(0.012)	(0.012)	(0.013)
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	0.174	0.065	0.173
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, the control variables include the worker's level and field of education, the year the project began, and the 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 29 log-point 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 significantly influences the gender wage gap. However, its role may be less pronounced in online freelancing markets for several reasons. First, education is self-reported and unverified by the platform, which can reduce its credibility as a selection criterion. Additionally, the inherent instability of the gig economy—with its lack of long-term employment and benefits—diminishes the importance of educational credentials in worker selection. Freelancing platforms also tend to prioritize tangible skills and proven work history over formal educational backgrounds. Nonetheless, our analysis sample is highly educated, with approximately 70% holding a college degree or higher—compared to about 40% of U.S. adults (Pew Research Center, 2022). In line with our findings, Herrmann et al. (2023) show 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 Main Results: Skills, Application Behavior and Hourly Wages

In the previous section, we demonstrate that a considerable gender wage gap persists across different occupations, levels of formal education and over time. 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 workers' human capital, or application strategies. After controlling for these, any unexplained gender gap could be due to either employer discrimination or unobservable differences in worker preferences, constraints or application strategies. This residual earning difference between genders is captured by the coefficient of the *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 skills using a machine learning-based approach. Second, we describe our approach for capturing two key aspects of application strategies: the job amenities associated with applied projects and the asking wages set by workers. Finally, we combine these components in a decomposition analysis, applying the Gelbach (2016) linear decomposition method to determine the portion of the hourly wage gap attributable to each factor.

4.1 Worker Skills

Our results so far show that men are more common in higher-paying jobs. This difference is not explained by education alone and may be due to differences in skills between workers. For instance, coding jobs require coding skills, just as translation work requires language skills. We now turn to explaining how we operationalize skill measurement within our regression framework.

Our data includes granular information on worker skills and project skill requirements, comprising over 4,000 individual skills combined in various ways. To effectively analyze this high-dimensional data, we employ a machine learning approach to *learn* the value of skills. Specifically, we use the Random Forest algorithm (Breiman, 2001) to compress the binary skill indicators into a one-dimensional continuous representation. The non-linear nature of Random Forests enables us to capture interactions between skills, allowing for a more nuanced valuation. Conceptually, this parsimonious measure represents the estimated market value of each skill combination, predicting the hourly wage a worker could expect based on their unique skill set – in short, the skill value.

In practice, we proceed in the following way. First, we train a Random Forest model on the skill requirements and log-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.[17] 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 offered by female and male workers 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).[18]

A potential concern in the approach of using realized market wages for valuing workers' skills is that if there is any gender bias in how skills are valued, this would introduce systematic bias into the estimated skill values. To probe this, we show that our results remain similar if we use only projects completed by males or projects completed by females as the training sample. Additionally, we demonstrate that the association between worker skills and realized wages does not vary by gender. These results are reported in Appendix Section A.4.

[17]See Appendix Section A.3 for details on the hyperparameter-tuning, the training process, model performance, and a comparison to other machine learning models.

[18]Additionally, we compare the predicted skill values across educational levels (see Appendix Figure A6), which shows that skill values increase with higher levels of formal education. This trend is expected and further validates our approach. We also compare predicted skill values by gender across educational levels (see Appendix Figure A7), revealing a significant gender gap in skill values across all education levels.

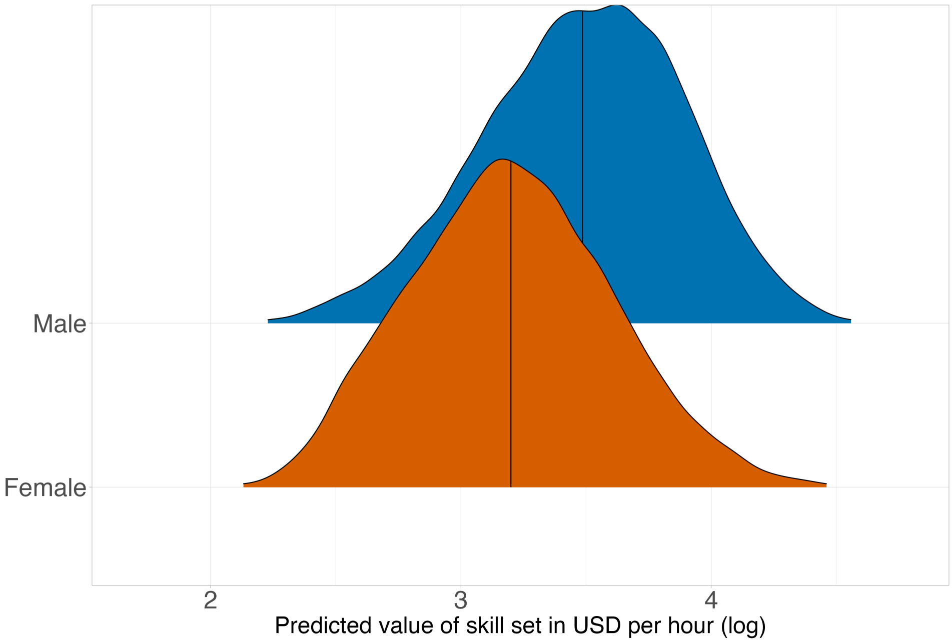

Figure 3. Predicted value of workers' skill sets by gender

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 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 utilize the one-dimensional fitted values in the next section when controlling for the skill requirements of the applications made by workers.[19]

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, Column 1) with the gender wage gap after adjusting for the predicted market value of worker skills. Moreover, we sidestep the "bad controls" issue discussed in Section 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:

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

[19]We show that our results remain similar if we use the full set of skill dummies instead of the ML based skill values (see Appendix Section A.7). While detailed results will follow in subsequent sections, we can already confirm that our core result——the decomposition of the gender wage gap into components attributable to skills and application behavior——remains consistent, regardless of how worker skills are integrated into the analysis.

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, it means the measurement error in *Skills* may cause us to overestimate the gender wage gap when accounting for skills. However, we demonstrate in Appendix Section A.8 that concentrating on new workers, who likely have not adjusted their reported skills does not change the empirical results.

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.

Third, it is well understood that regularization methods, such as Random Forest, improve model forecasting ability but introduce bias into the regularized coefficients. Chernozhukov et al. (2018) emphasize that including a regularized term in a regression model with a binary dummy variable can transmit this regularization bias into the parameter of interest, similar to omitted variable bias. This bias can be substantial, even in moderately sized samples. It is crucial 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	-0.126***	-0.130***
	(0.011)	(0.012)
Skills	0.668***	0.673***
	(0.013)	(0.014)
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	0.319	0.322
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 level and field of education, 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 12.6 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 include data on projects, where the employer can be based in any country. To ensure that 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 narrows from 33.6 to 12.6 log points.*

4.2 Workers' Application Strategies

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. Indeed, 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 number of expected weekly contract hours. 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 practice, 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 not one clear-cut way 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 before 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.[20]

[20] See Appendix Table A6.

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 the platform API, the list of applicants is not always populated, 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 in Appendix A.1.) 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 is missing at random.

Table 4 shows the descriptive statistics for this sample. The average hourly wage and log-hourly wage are largely consistent across both analysis samples. Turning to application behavior, the differences between men and women are substantial. First, men predominantly apply to projects expected to be completed in less than 10 hours, whereas women apply to projects lasting longer. Additionally, men are more likely to apply to jobs, where the employers have declared they prefer expert-level knowledge. 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 jobs being available to workers with different skills, or gender differences in preferences concerning job amenities.

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

	Male		Female		Difference in means
	Mean	Median	Mean	Median	(female – male)
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: Weekly contract hours					
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)		
Weekly contract hours 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	0.000	0.114	0.000	-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	0.000	-0.006**
	(0.171)		(0.176)		
3 to 6 months	0.073	0.000	0.082	0.000	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)		
Application success rate ◇	0.355	0.250	0.384	0.286	0.029***
	(0.284)		(0.290)		

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)		
Expected value of past applications					
Expected 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. Information on worker characteristics and main project categories can be found in Table A3 in the Appendix. 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 $p < 0.01$.

⋄ Due to missing information, the variable *Application success rate* is available for 24,876 observations (5,664 men and 5,681 women).

To disentangle these channels, we extend Equation 2 to account for the application strategies of workers by:

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

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). *Weekly contract hours* 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 Expected value of past applications$_{ijt}$. We create this variable by predicting the value of applied-for projects based on skill requirements using the same machine learning model as described in Section 4.1.

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 not statistically different from the one reported in Table 3. We proceed by gradually introducing additional variables related to workers' application behavior in subsequent columns. We find that *Desired worker experience*, *Weekly contract hours*, *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 reduces the gender wage gap to 1.8 log points.

The fact that both individual worker skills (variables Skills$_i$) and the expected value of past applications have independent predictive power on wages, indicates that men and women apply for jobs with different skill requirements, even when skill differences between workers are controlled. Furthermore, the expected value of past applications 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 have lower skill demands, even conditional on other covariates.

One concern in the results presented in Table 5 is that the *Expected value of past applications* and *Asking wage* variables might capture some degree of sorting across project categories. If this were the case, they might inadvertently introduce the "bad controls" issue they are meant to avoid. As a sensitivity check, we include project category dummies as additional controls in Column 10. Intuitively, this is equivalent to comparing workers within the same project category. A comparison between Columns 9 and 10 shows that the regression coefficients remain virtually unchanged, providing no evidence of bias due to sorting. Moreover, we note, that the adjusted R^2 increases only marginally after the inclusion of job category dummies.

Overall, our findings indicate, that the differences in human capital, and application strategies between men and women reduce the unexplained gender wage gap to less than 2%. This operates via four channels. First, women have skills that are less valued in the labor market. Second, conditional on skills, women apply for projects that pay, on average, less. Thirdly, 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 four differences explain the gender wage gap in online freelancing.

Result 4. *Accounting for workers' differences in human capital and application strategies reduces the unexplained gender wage gap to under 2 log-points.*

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

	Hourly wage (log)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female	-0.355***	-0.310***	-0.140***	-0.101***	-0.099***	-0.093***	-0.094***	-0.042***	-0.018**	-0.010
	(0.016)	(0.016)	(0.015)	(0.013)	(0.013)	(0.013)	(0.012)	(0.009)	(0.009)	(0.009)
Skills			0.683***	0.524***	0.519***	0.501***	0.502***	0.178***	0.110***	0.095***
			(0.018)	(0.015)	(0.015)	(0.015)	(0.015)	(0.014)	(0.014)	(0.014)
Asking wage (log)								0.521***	0.497***	0.494***
								(0.011)	(0.011)	(0.011)
Expected value of past applications									0.250***	0.213***
									(0.013)	(0.012)
Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓	✓
Number of past applications							✓	✓	✓	✓
Project categories										✓
Number of projects	27,698	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	13,267
Adjusted R^2	0.067	0.189	0.342	0.472	0.473	0.478	0.479	0.663	0.672	0.675
Share of females	50.24%	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 level and field of education). Column 3 presents the results from the regression specification where we further account for the market value of workers' skills. In Columns 4 to 9, 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. In Column 10, we include project category dummies as additional controls. 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 explanatory power of workers' human capital, and job-application strategies on 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 Gelbach decomposition technique aims to break down the aggregate explanatory power of covariates on the gender wage gap in a way that is unaffected by the sequence in which covariates are introduced. This method enables us to isolate the contribution of each covariate to shifts in the *Female* coefficient. It is particularly useful when factors are correlated with one another, as with worker skills, the expected value of past applications, and asking wages. The decomposition disentangles these contributions, while the estimated standard error accounts for uncertainty arising from potential multicollinearity among regressors. More concretely, consider a model where the dependent variable Y, log-hourly wage, is a function of X_1 and X_2:

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

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 $(X_1'X_1)^{-1}X_1'\mathbf{X_2}\boldsymbol{\beta_2}$. 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 into the relative contribution of each covariate to the gender wage gap.

Figure 4 presents the computed $\tilde{\pi}_k$ parameters and their 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, estimated using the estimating equation 3).

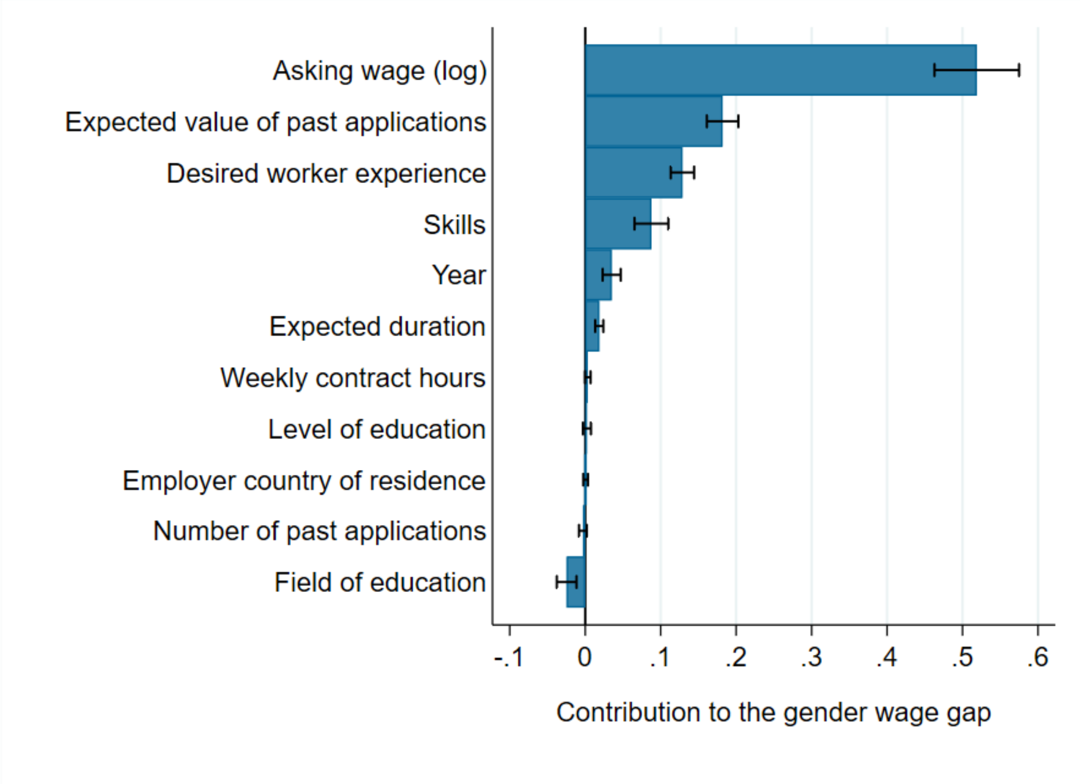

Figure 4. Gelbach decomposition

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

The main takeaway of Figure 4 is that application strategy can account for the majority of the gender wage gap. *Asking wage* alone accounts for 52% of the explained gender wage gap. *Expected value of past applications* accounts for the second-largest share, contributing approximately 20% to the explained wage disparity. *Desired worker experience* contributes about 13% of the gap. Taken together, application strategy accounts for 85% of the gender wage gap. At the same time, human capital (*Skills*, education level, and field) accounts for approximately 7% of the gender wage gap. The remainder of the explained share of the gender wage gap is explained by year dummies.

As with any decomposition method, the ultimate results depend on which covariates are included in the regression. For instance, if we excluded the covariates related to application strategies from the model, the decomposition would correspond to Equation 2, and would attribute the majority of the gender wage gap to differences in human capital.

More generally, the decomposition is susceptible to omitted variables bias. In particular, if there are unobserved gender differences in confidence, negotiation skills, preferences or other omitted variables that are correlated with both application strategies and hourly wages, the decomposition will attribute this difference to gender differences in application strategies. Indeed, one possible omitted variable is employer discrimination: if women anticipate lower wages or fewer opportunities in certain types of projects, they may strategically avoid applying to those projects, choosing instead to focus on areas where they expect higher success. If these unobservable factors are correlated with workers' application strategy, then the decomposition would overestimate the portion of the gender wage gap accounted for by differences in job application strategy, making this an upper-bound estimate. With this caveat, we summarize the decomposition analysis as follows:

Result 5. *Application strategy – the difference in the types of jobs workers apply to – accounts for up to 90% of the gender wage gap.*

5 Additional Results: Gender Gap in Hours Worked, Application Success and Return to Experience

Our analysis in the previous section shows that while women earn substantially less than men per hour, their application behavior and human capital almost completely account for this difference. Women tend to ask for lower wages and apply for less demanding projects (measured by expected hourly wages). According to our decomposition analysis, this difference in "ambition" can account for over 70% of the gender wage gap.

However, the mechanisms underlying these differences remain unclear. While platform data provide a unique opportunity to document gender differences in application behavior across the skill distribution, it is less well suited for exploring the reasons behind these patterns. In this section, we document – within the limitations of our data – additional findings that provide suggestive evidence on the underlying mechanisms.

5.1 Gender Differences in Asking Wages and Applied-For Jobs

Figure 4 demonstrates that women tend to target lower-paying jobs and ask for lower wages. This behavior is consistent with findings from the research literature, which shows that women ask for smaller wages and are less likely to negotiate, contributing to the gender wage gap Roussille (2021); Dreber et al. (2022); Recalde and Vesterlund (2023); Gomez-Herrera and Mueller-Langer (2024); Foong et al. (2018); Munoz et al. (2024). Additionally, women are more likely to target jobs with lower skill requirements, even when accounting for their skills and asking wages. This aligns with Fluchtmann et al. (2021), who found that men are more likely to apply for jobs that offer wages above the industry average. Moreover, our findings confirm the results from Gomez-Herrera and Mueller-Langer (2024) who find, that the difference in asking wages is a major determinant of the gender wage gap in the platform economy context.

These findings suggest that the platform could boost both female earnings and its own revenue (since it takes a cut of platform earnings) by encouraging women to aim for higher-paid jobs and ask for higher wages. However, it is not clear that women would always benefit from doing so. In fact, Exley et al. (2020) show in a lab setting that women who are pushed to ask for higher wages often end up with worse outcomes than those who are less aggressive in negotiations. The study finds that pushing for higher wages can increase the risk of accepting a bad offer, suggesting that women may be aware of the risks related to asking for more. This suggests that women's lower asking wages could be a strategic choice to avoid the risk of not being hired. Therefore, the platform should weigh the risks of encouraging women to ask for higher wages, as pushing them to be more aggressive may actually lower their earnings.

In the next section, we show that women tend to work longer hours and have a higher probability of being hired than men. This supports the idea that men, by asking for higher wages, may reduce their chances of being hired and reduce their project hours, which in turn makes the benefits of asking for higher pay unclear.

5.2 Gender Gap in Hours Worked and Application Success Rate

Using the regression model from Equation 3, we set log-project hours as the dependent variable and report results in Table 6. The gender dummy in all specifications is positive, suggesting that, conditional on worker skills and asking wages, women work on longer projects than men. However, when we control for the expected project length of applied-for projects, the gender difference is no longer statistically significant. Table 6 also shows that asking wages, worker skills, and expected wage in past applications are negatively associated with project length. These results suggest that women offset lower hourly wages by working longer hours.

We investigate gender differences in project length and application success rates. To begin, we repeat the regression analysis from Equation 3, setting log-project hours as the dependent variable. The results, presented in Table 6, show that the gender dummy is positive across all specifications. This indicates that, conditional on worker skills and asking wages, women work on longer projects than men. However, the difference becomes statistically insignificant, though it remains positive, when we control for the expected project length of applied-for projects. This suggests that the primary reason women work on longer projects is their tendency to apply for longer ones. Table 6 also reveals that asking wages, worker skills, and the expected wage in past applications are negatively associated with project length. These findings suggest that women may partially offset their lower hourly wages by accumulating more work hours on longer projects.

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.[21] According to results in Table 7, women are more likely to win the projects they apply to, even after conditioning on job application strategies.[22] Our findings on women being more successful applicants confirm the results of Chan and Wang (2018), who provide evidence of hiring bias in favor of females in platform freelancing. Furthermore, we observe that the gender dummy coefficient remains virtually unchanged when conditioning on application behavior and asking wages. This suggests that the employer bias in favor of females is not mediated by differences in skills or application behavior between genders. However, we note that this observation is conditional on the applications workers have submitted, and it is unclear whether women apply for less selective jobs than men. Chan and Wang (2018) also provide evidence that employers perceive women as more trustworthy. To the extent that perceived trustworthiness is independent of workers' application behavior, it could further explain our results.

To summarize, we find, that women work longer hours and are more likely to win projects they bid on compared to men, even if at lower hourly wages. Our finding is in contrast to the findings of Bertrand et al. (2010) and Goldin (2014), who study the incomes of MBAs working in the financial and corporate sectors. They find that men earn more per hour than women because they work longer hours ("convex hours-pay relationship"). Our results point in the opposite direction: women in platform labor earn less per hour while working longer hours and winning more projects.

In the absence of information on workers' off-line work opportunities, it is challenging to draw conclusions on why women seem to exchange lower hourly pay for longer hours, but one explanation is that women are more dependent on platform work. 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. Table 4 also indicates that men are likelier to go for part-time and short-term projects requiring expert skills. Considering all this information, the evidence suggests that male workers might have better job chances off-platform, making full-time platform work less attractive for them. We summarize the results regarding differences in project duration and success rate as follows:

Result 6. *Women work on longer projects than men. Conditional on asking wages and skills, women engage in longer projects, suggesting that they partially offset lower hourly wages on the platform by accumulating more work hours.*

[21]To avoid a mechanical relationship between a regressor and the dependent variable we exclude the number of past applications from the set of regressors.

[22]We note, however, that we do not observe the effort cost of each application submitted. It is not clear, whether men and women spend the same time when preparing applications.

Table 6. Gender gap in log-hours

	Project duration (log-hours)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female	0.090***	0.121***	0.058**	0.058**	0.047*	0.026	0.027	0.018	0.012
	(0.025)	(0.025)	(0.026)	(0.026)	(0.025)	(0.024)	(0.024)	(0.024)	(0.024)
Skills			-0.252***	-0.245***	-0.241***	-0.177***	-0.180***	-0.126***	-0.108***
			(0.031)	(0.032)	(0.032)	(0.031)	(0.030)	(0.033)	(0.034)
Asking wage (log)								-0.086***	-0.080***
								(0.021)	(0.022)
Expected value of past applications									-0.065*
									(0.037)
Controls		✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓
Number of past applications							✓	✓	✓
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	0.001	0.018	0.022	0.024	0.031	0.046	0.048	0.049	0.049
Share of females	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%	50.24%

Note: This table documents the gender wage gap in log-hours. Column 1 reports the raw wage gap without controls. In Column 2 we include control variables for worker's highest degree and field of education, the year in which the project started, and employer country dummies. In Columns 3 to 9, we progressively incorporate controls for workers' skills and job application behavior. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels are indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Result 7. *Women have a higher application success rate than men. Conditional on worker skills and other observables, women are more likely to secure projects they apply to.*

5.3 No Gender Differences in Return to Experience

The literature on platform labor consistently highlights substantial returns to experience (Pallais, 2014; Agrawal et al., 2016; Cook et al., 2021). For example, Cook et al. (2021) show that men's earnings advantage in ride-hailing arises partly from a strong learning-by-doing effect. Another explanation is statistical discrimination: workers with less experience may face disadvantages due to information frictions (Pallais, 2014; Agrawal et al., 2016; Lehdonvirta et al., 2019; Kässi and Lehdonvirta, 2024). Both the learning-by-doing and statistical discrimination hypotheses share the same empirical implication: more experienced workers should earn higher hourly rates, either because of decreased employer uncertainty, or increased productivity.

To study these hypotheses, we first plot platform work experience and offline work experience by gender. We define platform work experience as the number of completed projects at the start of each new project, which is visible to prospective employers. Traditional work experience is measured as the time between the start date of the first (self-)reported job in the worker's profile and the start date of the project, measured in full years.

We plot the distribution of both in Figures 5 and 6. Since both platform work experience and total work experience are heavily right-skewed, we plot the distribution of

Table 7. Gender gap in application success rate

	Application Success Rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female	0.019***	0.021***	0.017***	0.015***	0.013***	0.014***	0.014***	0.013***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Skills			-0.019***	-0.013**	-0.013**	-0.014***	-0.014**	-0.011*
			(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)
Asking wage (log)							-0.001	0.0001
							(0.004)	(0.004)
Expected value of past applications								-0.011
								(0.007)
Controls		✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓
Expected duration						✓	✓	✓
Number of projects	22,056	22,056	22,056	22,056	22,056	22,056	22,056	22,056
Number of workers	11,679	11,679	11,679	11,679	11,679	11,679	11,679	11,679
Adjusted R^2	0.002	0.102	0.104	0.107	0.124	0.127	0.127	0.127
Share of females	46.84%	46.84%	46.84%	46.84%	46.84%	46.84%	46.84%	46.84%

Note: This table documents the gender gap in application success rate. Column 1 reports the raw relationship between gender and application success rate. In Column 2 we include control variables for worker's highest degree, field of education, and employer country dummies. In Columns 3 to 8, we progressively incorporate controls for workers' skills and job application behavior. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels are indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

both in four quartiles. Figures 5 and 6 both show, that men on the platform have slightly more experience than women.

To study whether this difference in experience translates to differences in hourly wages, we estimate models where we have interacted worker experience levels (using both measures for experience) with the gender dummy. To facilitate a comparison between the two measures of experience that are on different scale, the experience measures are normalized by demeaning and dividing by the standard deviation. These are reported in Table 8. We find that while experience and hourly wages are slightly positively correlated, there is no difference in return to experience between genders. Thus, while our results provide suggestive evidence in favor of both statistical discrimination and learning-by-doing hypotheses, the returns to experience are far too small to account for the gender wage gap.

Result 8. *Men have slightly more platform and traditional work experience than women, but gender differences in returns to experience are too small to explain the gender wage gap.*

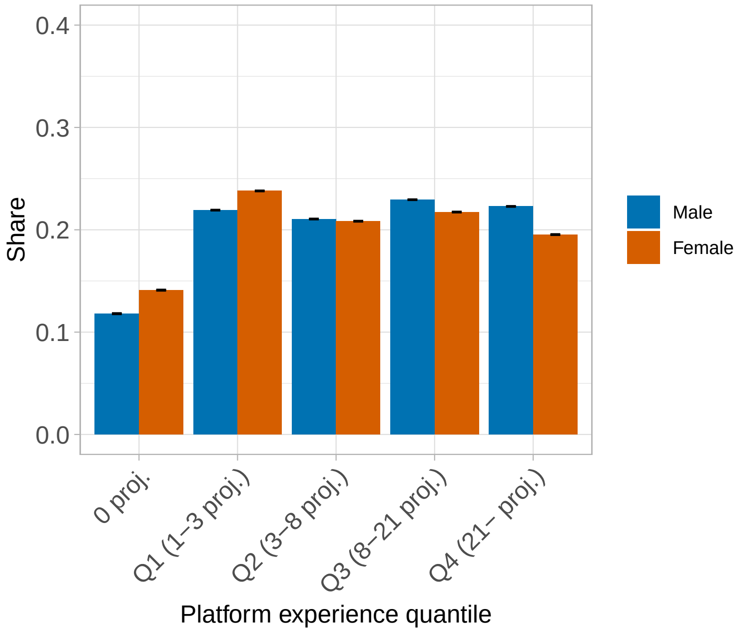

Figure 5. Gender distribution within experience quantiles

Note: This figure plots the distribution of completed projects by quartile. The number of completed projects is measured at the time of project start, and each worker is observed more than once.

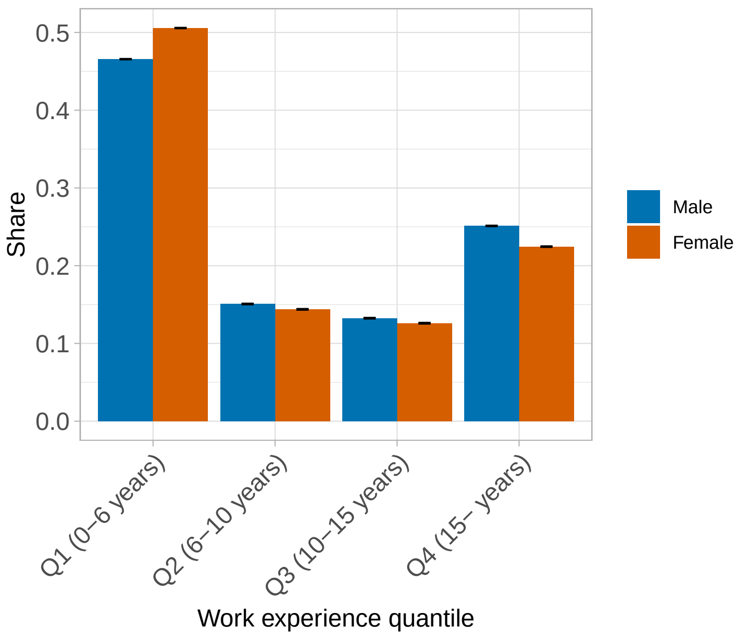

Figure 6. Gender distribution within experience quantiles

Note: This figure plots the gender distribution of self-reported work experience measured by the difference between earliest reported work experience and the start of project measured in years.

Table 8. Return to experience and the gender wage gap

	Hourly wage (log)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Platform experience and log-hourly wage									
Female	-0.348***	-0.307***	-0.136***	-0.100***	-0.099***	-0.092***	-0.093***	-0.039***	-0.016*
	(0.017)	(0.017)	(0.016)	(0.013)	(0.013)	(0.013)	(0.013)	(0.009)	(0.009)
Platform experience	0.049***	0.016	0.024**	0.012	0.012	0.010	0.014	0.024***	0.024***
	(0.017)	(0.014)	(0.012)	(0.010)	(0.010)	(0.010)	(0.010)	(0.006)	(0.006)
Female × Platform experience	0.035	0.027	0.008	-0.012	-0.012	-0.014	-0.013	0.005	0.001
	(0.035)	(0.029)	(0.026)	(0.022)	(0.022)	(0.022)	(0.020)	(0.011)	(0.011)
Controls		✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓
Number of past applications							✓	✓	✓
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	0.075	0.189	0.342	0.471	0.473	0.478	0.478	0.663	0.673
Share of females	47.21%	47.21%	47.21%	47.21%	47.21%	47.21%	47.21%	47.21%	47.21%
Panel B: Total work experience and log-hourly wage									
Female	-0.349***	-0.300***	-0.131***	-0.097***	-0.095***	-0.088***	-0.087***	-0.032***	-0.007
	(0.019)	(0.019)	(0.018)	(0.015)	(0.015)	(0.015)	(0.015)	(0.010)	(0.010)
Total experience	0.088***	0.064***	0.069***	0.042***	0.042***	0.040***	0.040***	0.058***	0.057***
	(0.016)	(0.016)	(0.013)	(0.011)	(0.011)	(0.011)	(0.011)	(0.008)	(0.008)
Female × Total experience	0.012	0.031	0.015	0.008	0.008	0.009	0.009	-0.002	-0.002
	(0.021)	(0.020)	(0.017)	(0.014)	(0.014)	(0.014)	(0.014)	(0.010)	(0.010)
Controls		✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓
Number of past applications							✓	✓	✓
Number of projects	20,229	20,229	20,229	20,229	20,229	20,229	20,229	20,229	20,229
Number of workers	9,653	9,653	9,653	9,653	9,653	9,653	9,653	9,653	9,653
Adjusted R^2	0.087	0.211	0.369	0.483	0.484	0.489	0.490	0.671	0.681
Share of females	49.70%	49.70%	49.70%	49.70%	49.70%	49.70%	49.70%	49.70%	49.70%

Note: Panel A presents the effect of platform work experience on gender wage gap, where platform work experience is defined as number of completed projects at the time of project start. Panel B repeats the analysis using years between the start of the project and earliest reported work experience as the measure of experience. Both experience measures are standardized by demeaning and dividing by standard deviation. Column 1 shows the raw gender wage gap. In Column 2, we include controls for project start year, employer country, and worker level and field of education. Column 3 presents the results from the regression specification where we account for the market value of workers' skills and control variables. In Columns 4 to 9, 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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 of 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 disparity typically seen in traditional labor markets (Statista Research Department, 2024a,b). 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 virtually disappears. Differences in application strategies account for the majority 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 occasional gigs. The differences in application strategies 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 would imply that women, on average, are more dependent on income from platform labor than men.

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. Nonetheless, our results align well with Fluchtmann et al. (2021), who report considerable gaps in applied-for jobs between genders that closely reflect the wage gap between genders. This, along with findings platform labor in general, bolsters the external validity of our results and suggests that similar dynamics may be at play across different labor market contexts.

Our findings highlight how offline constraints (such as other available opportunities) and preferences influence online labor market outcomes. Addressing these issues requires interventions beyond the platform rather than regulating the contractual working arrangements of individual platform providers. Strategies to enhance women's participation in higher-paying sectors include targeted, industry-specific training and mentoring for women, as well as providing low-cost childcare options for families. These measures may contribute to a more equitable distribution of opportunities and rewards across both traditional and online labor markets. Future research could explore how platforms might use behavioral nudges, such as informing female freelancers of potentially higher-earning opportunities in male-dominated project categories, to reduce gender segregation in platform work.

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.

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

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.

Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons, "Does standardized information in online markets disproportionately benefit job applicants from less developed countries?," *Journal of international Economics*, 2016, *103*, 1–12.

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.

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.

Chan, Jason and Jing Wang, "Hiring preferences in online labor markets: Evidence of a female hiring bias," *Management Science*, 2018, *64* (7), 2973–2994.

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.

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.

Croson, Rachel and Uri Gneezy, "Gender differences in preferences," *Journal of Economic Literature*, 2009, *47* (2), 448–474.

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.

Dalton, Michael R. and Jeffrey A. Groen, "How do jobseekers search for jobs? New data on applications, interviews, and job offers," *Beyond the Numbers: Employment and Unemployment, U.S. Bureau of Labor Statistics*, 2020, *9* (14).

Detilleux, Céline and Nick Deschacht, "The gender gap in the wage sensitivity of job transitions: a decomposition analysis," *Journal for Labour Market Research*, 2024, *58* (1), 1–12.

Dreber, Anna, Emma Heikensten, and Jenny Säve-Söderbergh, "Why do women ask for less?," *Labour Economics*, 2022, *78*, 102204.

Exley, Christine L, Muriel Niederle, and Lise Vesterlund, "Knowing when to ask: The cost of leaning in," *Journal of Political Economy*, 2020, *128* (3), 816–854.

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 Mueller-Langer, "Does Information Disclosure Affect the Gender Gap in Bidding Behavior? Empirical Evidence from a Natural Experiment on a Large Online Labor Platform," *Labour Economics*, 2024, *90*, 102585.

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*, 2024, *59* (4), 1284–1318.

Kleven, Henrik, Camille Landais, 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.

Le Barbanchon, Thomas, 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.

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.

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.

Munoz, Isabel, Pyeonghwa Kim, Clea O'Neil, Michael Dunn, and Steve Sawyer, "Platformization of Inequality: Gender and Race in Digital Labor Platforms," *Proceedings of the ACM on Human-Computer Interaction*, 2024, *8* (CSCW1), 1–22.

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.

Pallais, Amanda, "Inefficient hiring in entry-level labor markets," *American Economic Review*, 2014, *104* (11), 3565–3599.

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.

Recalde, Maria P and Lise Vesterlund, "Gender differences in negotiation: can interventions reduce the gap?," *Annual Review of Economics*, 2023, *15* (1), 633–657.

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.

Statista Research Department, "Median hourly earnings of female wage and salary workers in the United States from 1979 to 2022," https://www.statista.com/statistics/185345/median-hourly-earnings-of-female-wage-and-salary-workers/ 2024. Last accessed: November 1, 2024.

_ , "Median hourly earnings of male wage and salary workers in the United States from 1979 to 2022," https://www.statista.com/statistics/185350/median-hourly-earnings-of-male-wage-and-salary-workers/ 2024. Last accessed: November 1, 2024.

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 Appendix

A.1 Descriptive Statistics

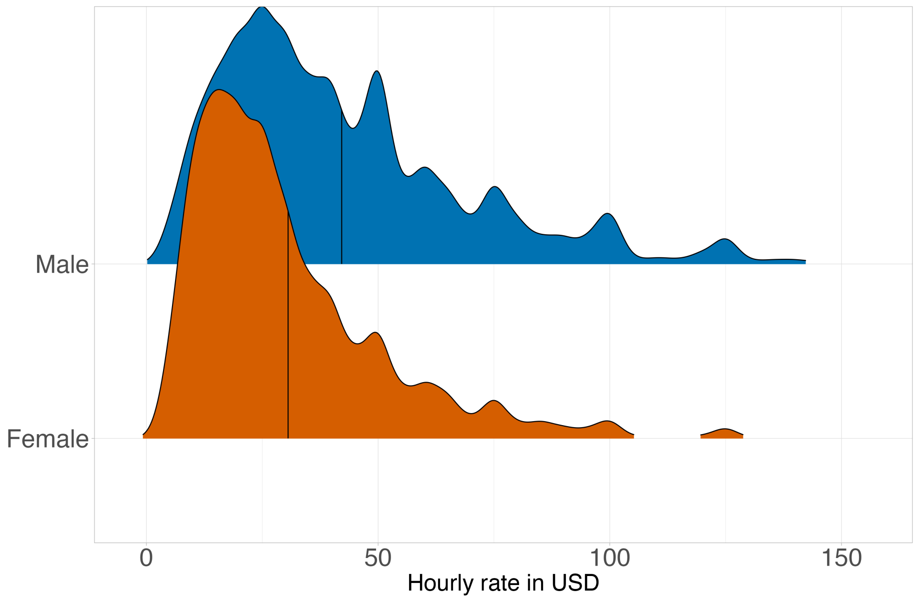

Appendix Figure A1. Distribution of hourly wage by gender

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

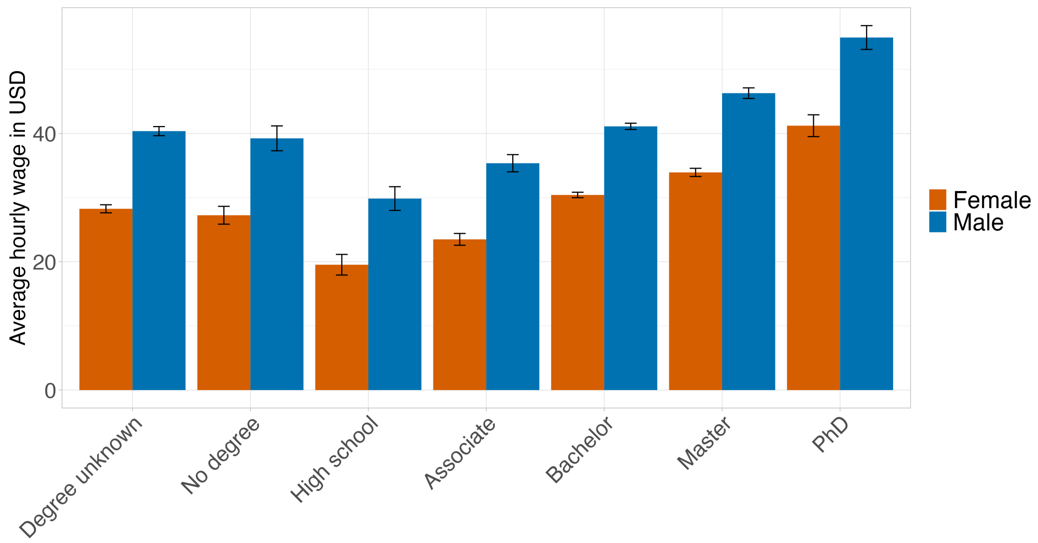

Appendix Figure A2. Hourly wage by education and gender

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

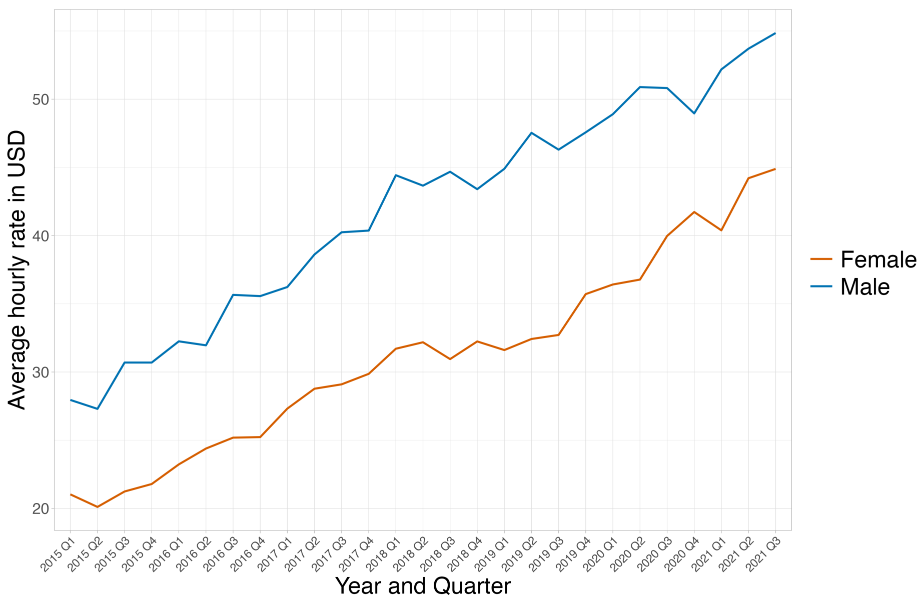

Appendix Figure A3. Hourly wage by gender over time

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

Appendix Table A1. Basic summary statistics for field of education

	Male		Female		Difference in means
	Mean	Median	Mean	Median	(female – male)
Agriculture	0.001	-	0.003	-	0.002***
	(0.036)		(0.059)		
Environmental studies	0.003	-	0.004	-	0.001
	(0.053)		(0.063)		
Architecture	0.005	-	0.003	-	-0.002
	(0.068)		(0.057)		
Cultural studies	0.001	-	0.003	-	0.002***
	(0.038)		(0.059)		
Media studies	0.044	-	0.074	-	0.030***
	(0.206)		(0.262)		
Communication services	0.004	-	0.003	-	-0.001*
	(0.065)		(0.053)		
Computer and information sciences	0.104	-	0.028	-	-0.076***
	(0.305)		(0.165)		
Personal services	0.001	-	0.003	-	0.002***
	(0.025)		(0.053)		
Education	0.011	-	0.030	-	0.019***
	(0.104)		(0.171)		
Engineering	0.052	-	0.011	-	-0.041***
	(0.223)		(0.103)		
Engineering technologies	0.006	-	0.002	-	-0.004***
	(0.075)		(0.044)		
Foreign languages	0.007	-	0.017	-	0.010***
	(0.083)		(0.129)		
Human sciences	0	-	0.002	-	0.002***
	(0.013)		(0.047)		
Legal studies	0.014	-	0.016	-	0.002
	(0.119)		(0.124)		
Literature	0.025	-	0.060	-	0.035***
	(0.156)		(0.238)		
Liberal arts & sciences	0.004	-	0.007	-	0.003***
	(0.063)		(0.085)		
Library science	0.001	-	0.002	-	0.001***
	(0.025)		(0.047)		
Biological & biomedical sciences	0.008	-	0.014	-	0.006***
	(0.090)		(0.117)		
Mathematics & statistics	0.009	-	0.004	-	-0.005***
	(0.096)		(0.066)		
Military technologies	0	-	0	-	0
	(0.019)		(0.013)		
Interdisciplinary studies	0.005	-	0.005	-	0
	(0.069)		(0.070)		
Recreation studies	0.003	-	0.003	-	0
	(0.051)		(0.056)		

Table A1 continued	Male Mean	Male Median	Female Mean	Female Median	Difference in means (female – male)
Basic skills & developmental education	0	-	0	-	0
	(0.000)		(0.013)		
Philosophy & religious studies	0.004	-	0.002	-	-0.002***
	(0.065)		(0.045)		
Theology	0.005	-	0.003	-	-0.002**
	(0.068)		(0.053)		
Physical sciences	0.010	-	0.004	-	-0.006***
	(0.097)		(0.061)		
Psychology	0.012	-	0.027	-	0.015***
	(0.110)		(0.162)		
Homeland security	0.003	-	0.004	-	0.001
	(0.052)		(0.062)		
Public administration	0.003	-	0.010	-	0.007***
	(0.054)		(0.101)		
Social sciences	0.028	-	0.029	-	0.001
	(0.165)		(0.168)		
Construction trades	0.001	-	0.	-	-0.001**
	(0.031)		(0.016)		
Mechanic trades	0	-	0	-	0
	(0.016)		(0.009)		
Precision production	0	-	0	-	0
	(0.016)		(0.000)		
Transportation	0	-	0	-	0
	(0.021)		(0.021)		
Visual & performing arts	0.081	-	0.100	-	0.019***
	(0.272)		(0.300)		
Health professions	0.006	-	0.027	-	0.021***
	(0.076)		(0.163)		
Business management	0.119	-	0.122	-	0.003
	(0.324)		(0.328)		
History	0.005	-	0.006	-	0.001
	(0.070)		(0.077)		
Health professions residency	0	-	0	-	0
	(0.009)		(0.009)		
Field of education unknown	0.375	-	0.329	-	-0.046***
	(0.484)		(0.470)		
No field of education	0.039	-	0.040	-	0.001
	(0.195)		(0.195)		
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. 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 $p < 0.01$.

Appendix Table A2. Summary statistics: 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
Project categories										
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.

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

	Male		Female		Difference in means
	Mean	Median	Mean	Median	(female – male)
Worker characteristics					
Project length (hours)	27.256	7.830	35.642	8.170	8.386***
	(59.368)		(93.631)		
Worker skill prediction	3.522	3.541	3.239	3.237	-0.283***
	(0.399)		(0.407)		
PhD	0.045	-	0.042	-	-0.003
	(0.208)		(0.200)		
Master	0.193	-	0.207	-	0.014**
	(0.394)		(0.405)		
Bachelor	0.456	-	0.473	-	0.017*
	(0.498)		(0.499)		
Associate	0.050	-	0.062	-	0.012***
	(0.217)		(0.240)		
High school	0.018	-	0.020	-	0.002
	(0.131)		(0.141)		
No degree	0.022	-	0.026	-	0.004
	(0.147)		(0.158)		
Degree unknown	0.217	-	0.170	-	-0.047***
	(0.412)		(0.376)		
Field of education					
Agriculture	0.001	-	0.002	-	0.001**
	(0.030)		(0.049)		
Environmental studies	0.003	-	0.004	-	0.001
	(0.054)		(0.061)		
Architecture	0.005	-	0.004	-	-0.001
	(0.067)		(0.060)		
Cultural studies	0.002	-	0.003	-	0.001*
	(0.041)		(0.057)		
Media studies	0.047	-	0.071	-	0.024***
	(0.212)		(0.258)		
Communication services	0.004	-	0.003	-	-0.001
	(0.065)		(0.055)		
Computer and information sciences	0.104	-	0.029	-	-0.075***
	(0.306)		(0.169)		
Personal services	0.001	-	0.003	-	0.002***
	(0.025)		(0.052)		
Education	0.010	-	0.030	-	0.020***
	(0.097)		(0.171)		
Engineering	0.051	-	0.010	-	-0.041***
	(0.220)		(0.098)		
Engineering technologies	0.005	-	0.002	-	-0.003***
	(0.074)		(0.039)		
Foreign languages	0.006	-	0.017	-	0.011***
	(0.079)		(0.131)		
Human sciences	0.000	-	0.002	-	0.002***
	(0.017)		(0.044)		

Table A3 continued	Male		Female		Difference in means
	Mean	Median	Mean	Median	(female – male)
Legal studies	0.017	-	0.017	-	0
	(0.127)		(0.130)		
Literature	0.025	-	0.056	-	0.031***
	(0.157)		(0.230)		
Liberal arts & sciences	0.005	-	0.007	-	0.002
	(0.074)		(0.083)		
Library science	0.001	-	0.002	-	0.001*
	(0.030)		(0.046)		
Biological & biomedical sciences	0.008	-	0.012	-	0.004**
	(0.089)		(0.108)		
Mathematics & statistics	0.009	-	0.004	-	-0.005***
	(0.096)		(0.062)		
Military technologies	0.000	-	0.000	-	0
	(0.017)		(0.017)		
Interdisciplinary studies	0.005	-	0.005	-	0
	(0.073)		(0.071)		
Recreation studies	0.003	-	0.002	-	-0.001
	(0.054)		(0.047)		
Basic skills & developmental education	0.000	-	0.000	-	0
	(0.000)		(0.012)		
Philosophy & religious studies	0.004	-	0.002	-	-0.002***
	(0.066)		(0.041)		
Theology	0.004	-	0.003	-	-0.001
	(0.061)		(0.053)		
Physical sciences	0.008	-	0.004	-	-0.004***
	(0.088)		(0.064)		
Psychology	0.012	-	0.029	-	0.017***
	(0.110)		(0.167)		
Homeland security	0.003	-	0.003	-	0
	(0.051)		(0.059)		
Public administration	0.003	-	0.011	-	0.008***
	(0.056)		(0.102)		
Social sciences	0.029	-	0.029	-	0
	(0.168)		(0.169)		
Construction trades	0.000	-	0.000	-	0
	(0.021)		(0.000)		
Mechanic trades	0.000	-	0.000	-	0
	(0.021)		(0.000)		
Precision production	0.000	-	0.000	-	0
	(0.017)		(0.000)		
Transportation	0.000	-	0.000	-	0
	(0.021)		(0.012)		
Visual & performing arts	0.077	-	0.103	-	0.026***
	(0.266)		(0.304)		
Health professions	0.005	-	0.026	-	0.021***
	(0.068)		(0.158)		
Business management	0.130	-	0.128	-	-0.002
	(0.336)		(0.334)		

Table A3 continued	Male Mean	Male Median	Female Mean	Female Median	Difference in means (female – male)
History	0.005	-	0.006	-	0.001
	(0.069)		(0.076)		
Health professions residency	0.000	-	0.000	-	0
	(0.000)		(0.012)		
Field of education unknown	0.367	-	0.325	-	-0.042***
	(0.482)		(0.469)		
No field of education	0.039	-	0.046	-	0.007*
	(0.195)		(0.209)		
Main project categories					
Accounting & consulting	0.055	-	0.049	-	-0.006**
	(0.228)		(0.216)		
Admin support	0.035	-	0.155	-	0.120***
	(0.184)		(0.362)		
Customer service	0.006	-	0.021	-	0.015***
	(0.080)		(0.143)		
Data science & analytics	0.046	-	0.014	-	-0.032***
	(0.210)		(0.117)		
Design & creative	0.168	-	0.168	-	0
	(0.374)		(0.373)		
Engineering & architecture	0.037	-	0.012	-	-0.025***
	(0.188)		(0.110)		
IT & networking	0.053	-	0.008	-	-0.045***
	(0.224)		(0.087)		
Legal	0.032	-	0.020	-	-0.012***
	(0.175)		(0.140)		
Sales & marketing	0.174	-	0.153	-	-0.021***
	(0.380)		(0.360)		
Web, mobile & software development	0.212	-	0.058	-	-0.154***
	(0.409)		(0.233)		
Translation	0.009	-	0.027	-	0.018***
	(0.096)		(0.161)		
Writing	0.172	-	0.316	-	0.144***
	(0.377)		(0.465)		
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.2 Identification of Freelancers' Field of Study

We use GPT-4o to classify the free text educational degree descriptions on freelancer profiles according to the Classification of Instructional Programs (CIP) taxonomy of academic disciplines at institutions of higher education in the United States. Examples of the CIP taxonomy are "Communication, Journalism, and related programs" (code 09), "Computer and Information Sciences and Support Services" (code 11) or "Visual and Performing Arts" (code 50). Examples of workers' free text educational degree descriptions are: "Bachelor of Fine Arts (B.F.A.) in Painting & Sculpture", "Bachelor of Arts (B.A.) in Psychology" or "Master of Arts (M.A.) in International Relations." If workers list multiple degrees, we consider only the highest degree. In total, we find that freelancers in our data set have degrees from 40 different fields of study. The five most common fields of study are: 1) "Business, Management, Marketing, and Related Support Services," 2) "Visual and Performing Arts," 3) "Computer and Information Sciences and Support Services," 4) "Communication, Journalism, and Related Programs," and 5) "English Language and Literature/Letters."

For each worker, we run the following prompt through the GPT-4o API:

```
# Task: I will give you information about the education of a
   worker.

# Your task is to extract different variables from the raw text
   data.

Note: 'missing' indicates a missing value. If a value is missing,
    return 'missing' as the response.

### First variable: Please classify the area of study according
   to the Classification of Instructional Programs (CIP). Use the
    CIP Title and CIP code.

- study area: {ed_area}

# Response Format: Please provide your response in JSON format
   following this structure:

{{'category_title': <CIP Title>, 'category_code': <CIP code>, '
   ranking': <low, medium, high, excellent>}}
```

A.3 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.[23] 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 A4 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 A4. 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.

[23]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.

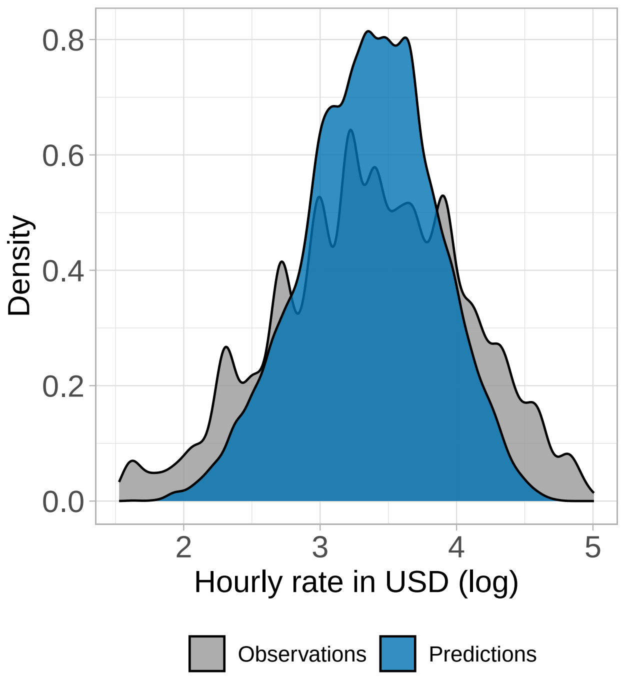

Appendix Figure A4. Comparing predicted and observed hourly rates of projects

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.

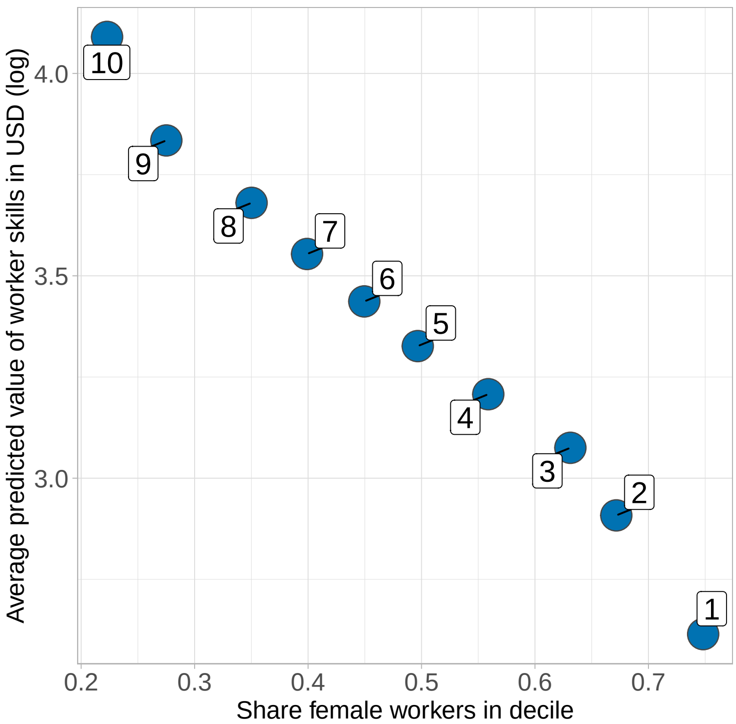

Appendix Figure A5. Average hourly wage and share of female workers in each skill value decile

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.

Predicted worker skill values increase with higher levels of formal education for both female and male freelancers (see Figure A7). However, a significant gender gap persists across all educational levels, with women having skills of notably lower market value despite holding the same level of formal education.

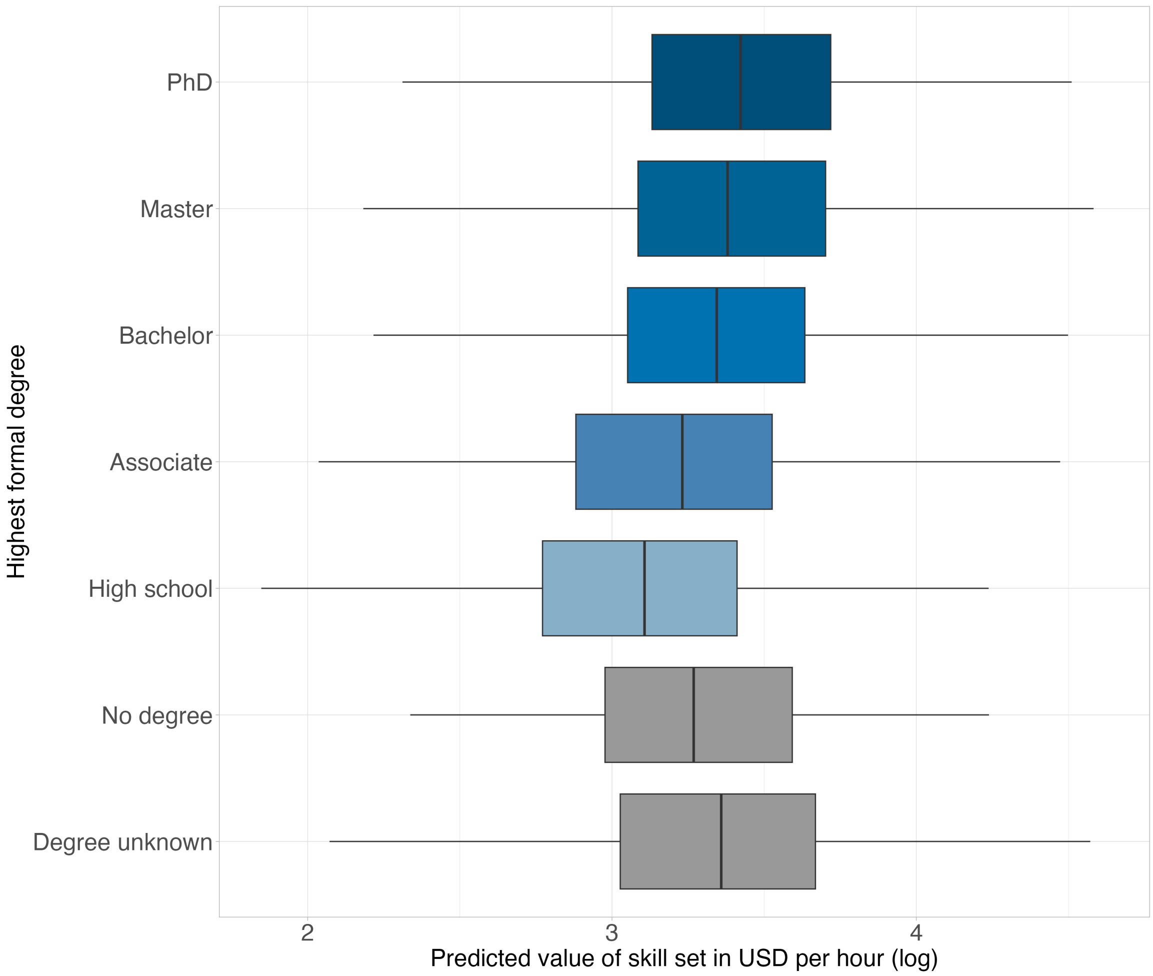

Appendix Figure A6. Predicted skill value of workers by level of formal education

Note: This plot illustrates the predicted skill set values of workers across educational levels. The categories range from "Degree unknown" and "No degree" to "PhD," with colors indicating educational progression. Each box plot represents the distribution of predicted skill values within each education level, highlighting median and interquartile ranges.

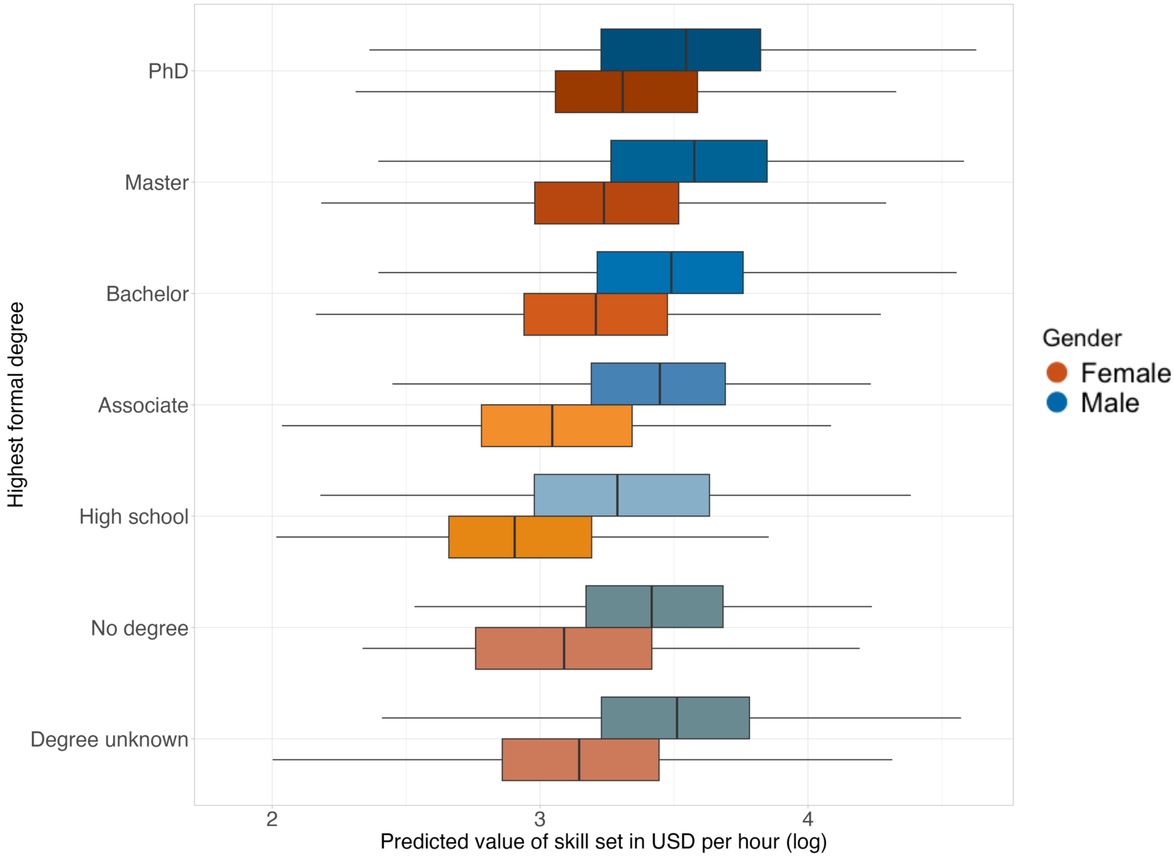

Appendix Figure A7. Predicted skill value of workers by gender and level of formal education

Note: This plot illustrates the predicted skill set values of workers by gender across educational levels. The categories range from "Degree unknown" and "No degree" to "PhD," with colors indicating educational progression. Each box plot represents the distribution of predicted skill values within each education level, highlighting median and interquartile ranges.

A.4 Addressing the Potential Bias in Skill Value Prediction

A potential concern with using skill values estimated from the same wage data as the outcome variable is the risk of capturing confounders, such as gender, that influence both skill reporting and wages. This could result in an underestimation of the unexplained portion of the gender wage gap.

To address this concern, we fit the Random Forest model using three subsets of the projects: the full sample, male workers only, and female workers only. This approach yields three sets of estimated skill values (*Skills*) and three sets of expected values for past applications (*Expected value of past applications*). We then repeat the analysis from Table 5 with each of these three sets of skill values. Figure A8 demonstrates that the results are largely consistent regardless of the training sample used for skill value estimation. The gender dummy coefficients are slightly smaller when only projects completed by women are used as the training data, but this difference becomes virtually zero after controlling for asking wages. The differences between the estimates are never statistically significant, suggesting that the skill value estimation is not biased by omitted variables.

Additionally, we present the results in Table A5 from a regression model using the job-search behavior sample, where we have interacted the worker skill value prediction with the *Female* dummy. We find that the coefficients on the gender dummy are virtually identical to those reported in Table 5. Moreover, after adding *Desired worker experience* as control variables, the coefficient on the interaction term (*Female* $\times$ *Skills*) becomes zero. Together, these findings support the assumption that the skill value estimation method does not introduce significant bias into the regression models we use.

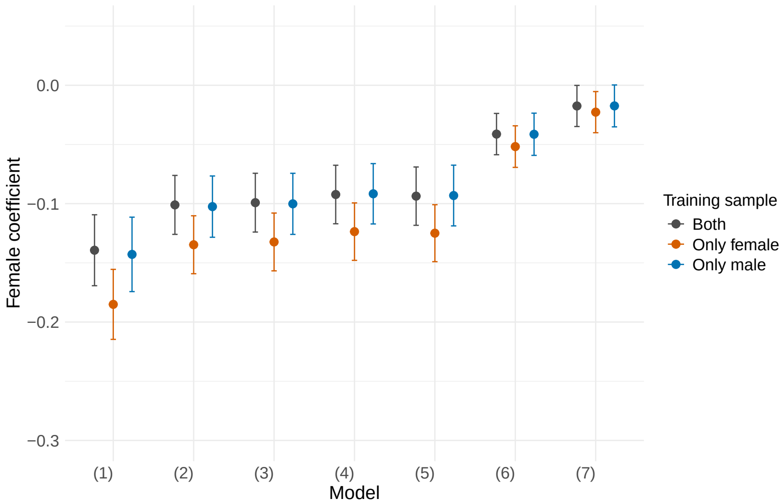

Appendix Figure A8. Skill value estimation: different subsamples

Note: This figure displays the coefficient of the *Female* dummy variable across different regression specifications (full sample, male workers only, female workers only), using subsamples of the data as the training sample for skill value estimation. The specifications are labeled (1)–(7), with each successive model incorporating additional controls:

- (1) Includes controls for education level, field, and the estimated skill values.
- (2) Adds application shares for different expected worker experience levels.
- (3) Further includes application shares for different weekly contract hours.
- (4) Incorporates expected project duration.
- (5) Includes the number of applications.
- (6) Adds the asking wage.
- (7) Further includes the expected value of applied-for projects.

Appendix Table A5. Gender wage gap: gender difference in return to skills

	Hourly wage (log)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female	-0.138***	-0.100***	-0.099***	-0.092***	-0.093***	-0.041***	-0.017**
	(0.015)	(0.013)	(0.013)	(0.013)	(0.013)	(0.009)	(0.009)
Skills	0.272***	0.216***	0.215***	0.208***	0.208***	0.073***	0.045***
	(0.011)	(0.009)	(0.009)	(0.009)	(0.009)	(0.007)	(0.007)
Asking wage (log)						0.521***	0.497***
						(0.011)	(0.011)
Expected value of past applications							0.250***
							(0.013)
Female × Skills	0.040***	0.015	0.014	0.011	0.012	0.006	0.003
	(0.014)	(0.012)	(0.012)	(0.012)	(0.012)	(0.009)	(0.009)
Controls	✓	✓	✓	✓	✓	✓	✓
Desired worker experience		✓	✓	✓	✓	✓	✓
Weekly contract hours			✓	✓	✓	✓	✓
Expected duration				✓	✓	✓	✓
Number of past applications					✓	✓	✓
Number of projects	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
Adjusted R^2	0.189	0.342	0.472	0.473	0.478	0.479	0.663
Share of females	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, and the predicted skill value interacted with the gender dummy. Controls refers to year, employer country, education level and field dummies. 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.5 Past Application Behavior: Different Time Windows

Appendix Table A6. Gender wage gap conditional on skills and application behavior: different time windows

	Hourly wage (log)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: 90 days										
Female	-0.341***	-0.294***	-0.125***	-0.077***	-0.075***	-0.068***	-0.068***	-0.029***	-0.004	0.003
	(0.015)	(0.015)	(0.014)	(0.011)	(0.011)	(0.011)	(0.011)	(0.008)	(0.008)	(0.008)
Skills			0.686***	0.497***	0.492***	0.468***	0.468***	0.169***	0.094***	0.083***
			(0.016)	(0.014)	(0.014)	(0.014)	(0.014)	(0.012)	(0.012)	(0.012)
Asking wage (log)								0.506***	0.480***	0.478***
								(0.010)	(0.010)	(0.010)
Expected value of past application									0.282***	0.243***
									(0.014)	(0.013)
Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓	✓
Number of past applications							✓	✓	✓	✓
Project categories										✓
Number of projects	33,160	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	15,598
Adjusted R^2	0.062	0.188	0.342	0.497	0.499	0.505	0.505	0.672	0.682	0.684
Share of females	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%	50.10%
Panel B: 365 days										
Female	-0.338***	-0.291***	-0.125***	-0.068***	-0.065***	-0.057***	-0.057***	-0.022***	0.006	0.011
	(0.014)	(0.014)	(0.013)	(0.010)	(0.010)	(0.010)	(0.010)	(0.007)	(0.007)	(0.007)
Skills			0.679***	0.465***	0.458***	0.431***	0.431***	0.149***	0.067***	0.060***
			(0.015)	(0.013)	(0.013)	(0.013)	(0.013)	(0.011)	(0.011)	(0.011)
Asking wage (log)								0.501***	0.473***	0.472***
								(0.009)	(0.010)	(0.009)
Expected value of past applications									0.319***	0.277***
									(0.015)	(0.015)
Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓	✓
Number of past applications							✓	✓	✓	✓
Project categories										✓
Number of projects	37,419	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	17,681
Adjusted R^2	0.061	0.187	0.339	0.512	0.513	0.520	0.520	0.678	0.688	0.690
Share of females	50.12%	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 across different time windows for 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 level and field of education). Column 3 presents the results from the regression specification where we further account for the market value of workers' skills. In Columns 4 to 9, 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. In Column 10, we add the main project categories as controls to assess the robustness of our results. 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 A7. Gender wage gap conditional on skills and application behavior: U.S. employers

	Hourly wage (log)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female	-0.360***	-0.315***	-0.139***	-0.098***	-0.097***	-0.089***	-0.091***	-0.041***	-0.016*	-0.008
	(0.017)	(0.017)	(0.016)	(0.013)	(0.013)	(0.013)	(0.013)	(0.009)	(0.009)	(0.009)
Skills			0.688***	0.530***	0.523***	0.505***	0.506***	0.184***	0.112***	0.098***
			(0.018)	(0.016)	(0.016)	(0.016)	(0.016)	(0.014)	(0.014)	(0.014)
Asking wage (log)								0.512***	0.487***	0.484***
								(0.011)	(0.011)	(0.011)
Expected value of past applications									0.260***	0.220***
									(0.014)	(0.014)
Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓	✓
Number of past applications							✓	✓	✓	✓
Project categories										✓
Number of projects	20,614	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	10,925
Adjusted R^2	0.070	0.186	0.345	0.476	0.478	0.483	0.483	0.664	0.674	0.677
Share of females	49.89%	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, considering only U.S. employers. Column 1 shows the raw gender wage gap. In Column 2, we control for conventional controls (project start year, employer country dummies, and worker level and field of education). Column 3 presents the results from the regression specification where we further account for the market value of workers' skills. In Columns 4 to 9, 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. In Column 10, we add the main project categories as controls to assess the robustness of our results. 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 A8. Gender wage gap conditional on skills and application behavior: skill dummies

	Hourly wage (log)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female	-0.355***	-0.310***	-0.081***	-0.060***	-0.059***	-0.056***	-0.057***	-0.030**	-0.020*	-0.014
	(0.016)	(0.016)	(0.014)	(0.012)	(0.012)	(0.012)	(0.012)	(0.009)	(0.009)	(0.009)
Asking wage (log)								0.509***	0.495***	0.489***
								(0.010)	(0.010)	(0.010)
Expected value of past applications									0.194***	0.165***
									(0.012)	(0.011)
Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓
Skill dummies			✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓	✓
Number of past applications							✓	✓	✓	✓
Project categories										✓
Number of projects	27,698	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	13,267
Adjusted R^2	0.067	0.189	0.508	0.582	0.583	0.585	0.585	0.711	0.716	0.719
Share of females	50.24%	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, considering only U.S. employers. Column 1 shows the raw gender wage gap. In Column 2, we control for conventional controls (project start year, employer country dummies, and worker level and field of education). Column 3 presents the results from the regression specification where we further account for the full set of skill dummies instead of the ML based skill values. In Columns 4 to 9, 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. In Column 10, we add the main project categories as controls to assess the robustness of our results. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels are indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

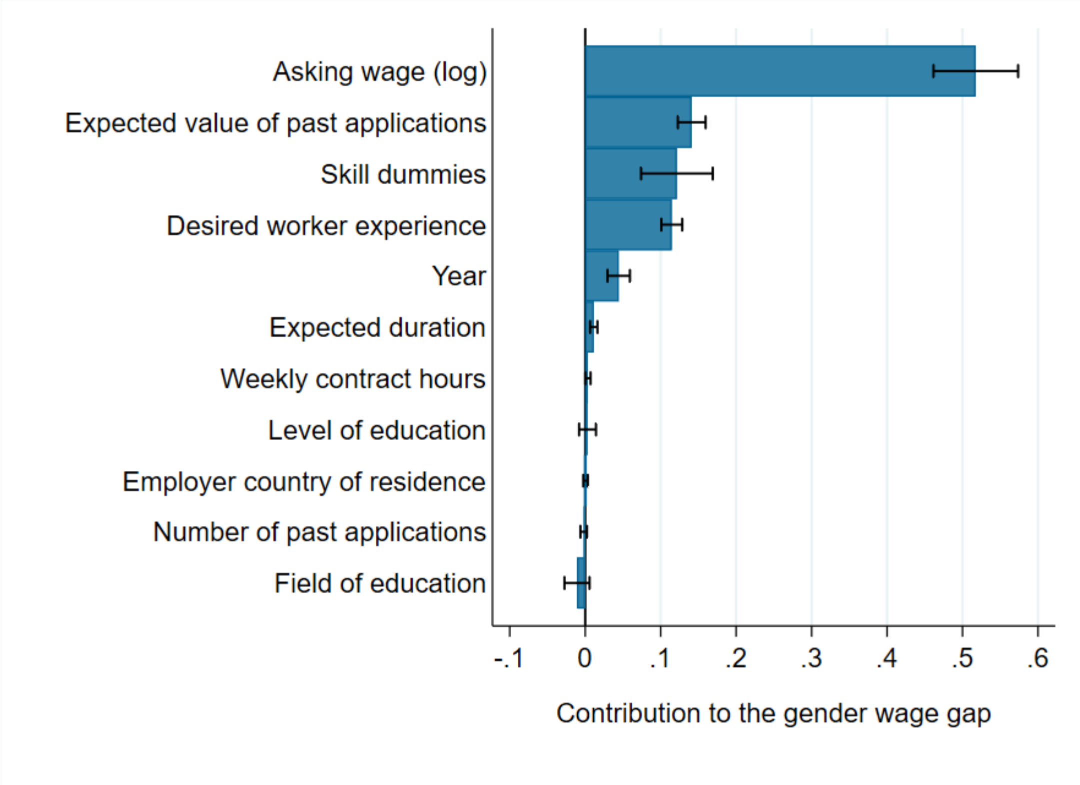

Appendix Figure A9. Gelbach decomposition

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

A.8 Results For New Workers Only

Our two central explanatory variables, asking wage and worker skills are time-invariant in our data while in reality, they vary in time. We demonstrate, that the results stay similar when we concentrate on workers who are at the start of their careers. Since the workers are at the start of their careers, it is less likely that they would have acquired new skills or changed their asking wages. Comparison between Tables 5 and A9 shows that the coefficients on both the gender dummy and skill and asking variables are almost unchanged, even if the sample sizes are much smaller.

Appendix Table A9. Gender wage gap conditional on skills and application behavior: new workers

	Hourly wage (log)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female	-0.342***	-0.304***	-0.132***	-0.099***	-0.100***	-0.093***	-0.093***	-0.047***	-0.018	-0.005
	(0.015)	(0.015)	(0.015)	(0.013)	(0.013)	(0.013)	(0.013)	(0.011)	(0.011)	(0.011)
Skills			0.633***	0.478***	0.472***	0.462***	0.460***	0.169***	0.091***	0.070***
			(0.017)	(0.016)	(0.016)	(0.016)	(0.016)	(0.015)	(0.016)	(0.016)
Asking wage (log)								0.462***	0.436***	0.432***
								(0.011)	(0.011)	(0.011)
Expected value of past applications									0.274***	0.226***
									(0.017)	(0.017)
Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓
Desired worker experience				✓	✓	✓	✓	✓	✓	✓
Weekly contract hours					✓	✓	✓	✓	✓	✓
Expected duration						✓	✓	✓	✓	✓
Number of past applications							✓	✓	✓	✓
Project categories										✓
Number of projects	9,887	9,887	9,887	9,887	9,887	9,887	9,887	9,887	9,887	9,887
Number of workers	8,072	8,072	8,072	8,072	8,072	8,072	8,072	8,072	8,072	8,072
Adjusted R^2	0.063	0.165	0.298	0.421	0.424	0.426	0.427	0.587	0.600	0.606
Share of females	50.13%	50.13%	50.13%	50.13%	50.13%	50.13%	50.13%	50.13%	50.13%	50.13%

Note: This table presents the gender wage gap conditional on application behavior for workers who have completed a maximum of three projects. Column 1 shows the raw gender wage gap. In Column 2, we control for conventional controls (project start year, employer country dummies, and worker level and field of education). Column 3 presents the results from the regression specification where we further account for the market value of workers' skills. In Columns 4 to 9, 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. Column 10 includes project category dummies as additional controls. Standard errors are clustered at the worker level and are reported in parentheses. Significance levels are indicated by: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Download ZEW Discussion Papers:

https://www.zew.de/en/publications/zew-discussion-papers

or see:

https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html
https://ideas.repec.org/s/zbw/zewdip.html

//
IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de