

Women (Still) Ask For Less: Gender Differences in Hourly Rate in an Online Labor Marketplace

EUREKA FOONG, Northwestern University, USA

NICHOLAS VINCENT, Northwestern University, USA

BRENT HECHT, Northwestern University, USA

ELIZABETH M. GERBER, Northwestern University, USA

In many traditional labor markets, women earn less on average compared to men. However, it is unclear whether this discrepancy persists in the online gig economy, which bears important differences from the traditional labor market (e.g., more flexible work arrangements, shorter-term engagements, reputation systems). In this study, we collected self-determined hourly bill rates from the public profiles of 48,019 workers in the United States (48.8% women) on Upwork, a popular gig work platform. The median female worker set hourly bill rates that were 74% of the median man's hourly bill rates, a gap that cannot be entirely explained by online and offline work experience, education level, and job category. However, in some job categories, we found evidence of a more complex relationship between gender and earnings: women earned more overall than men by working more hours, outpacing the effect of lower hourly bill rates. To better support equality in the rapidly growing gig economy, we encourage continual evaluation of the complex gender dynamics on these platforms and discuss whose responsibility it is to address inequalities.

CCS Concepts: • **H.5.3 [Group and Organization Interfaces]. Computer-supported cooperative work;**

Additional Key Words and Phrases: gig economy; gender; pay; equality; occupation; online labor marketplace; data transparency

ACM Reference format:

Eureka Foong, Nicholas Vincent, Brent Hecht, and Elizabeth M. Gerber. 2018. Women (Still) Ask For Less: Gender Differences in Hourly Rate in an Online Labor Marketplace. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 53 (November 2018), 21 pages.

<https://doi.org/10.1145/3274322>

1 INTRODUCTION

Across the globe, 655 million fewer women than men take part in the paid labor force [46]. In the United States (US), only 57% of the female population take part [20]. If more women participated, gross domestic product would increase by 5 to 10% [20]. Not only do women participate in the paid labor force at lower rates than men, women also earn less on average than men [72].

The US Bureau of Labor Statistics (BLS) estimated that the median female worker made only 82% of the median man's weekly earnings in 2016 [72]. Of workers who were paid on an hourly basis,

This work was generously supported by Northwestern University's Design Research Cluster and the School of Communication, and NSF Grant IIS-1707296. We would like to thank our anonymous reviewers, members of the Delta Lab at Northwestern University, Paola Sapienza, Ho Kwan Cheung, Denae Ford, Jamie Gorson, Emily Wang, Diego Gomez Zara, and Daniel Shams for their useful comments, suggestions, and support.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

2573-0142/2018/11-ART53 \$15.00

<https://doi.org/10.1145/3274322>

the median female worker made only 87% of the median man's hourly earnings [72]. Among other factors, gender differences in the desire for flexibility contribute to these gender gaps in pay and participation: women tend to hold more domestic responsibilities than men, discouraging them from participating in the paid labor force and disproportionately reducing their pay when they seek more flexible work options (e.g., [17, 19]).

The growing gig economy has the potential to mitigate both of these concerns. While the gig economy may decrease the demand for working class jobs that employed women often hold (e.g., administrative positions) [47, 72], it also has the potential to increase female participation and pay in the paid labor market by providing more flexible employment opportunities and transparent pay rates [50]. The gig economy is a labor market in which employers post millions of short-term jobs through online platforms such as Upwork [10], Amazon Mechanical Turk (MTurk) [1], and Uber [5]. In some marketplaces, workers are assigned to tasks with platform-determined pay rates (e.g., Uber [5]). On others, independent workers can view and choose jobs that match their schedules, skills, and financial needs [29, 46]. For example, a young mother with a master's degree in English can choose to edit an article for USD\$30/hour after her child goes to bed and a daughter with an associate's degree in medical transcription can ask for USD\$15 to transcribe a 20-minute audio file while supervising her elderly father's meal preparation. Online labor marketplaces that provide flexible work arrangements where workers can set their own rates after seeing others' rates [50] may allow women to earn more equal pay to men.

In this research, we evaluated this equal-pay hypothesis. Specifically, we asked: Are the wide gender pay gaps that occur in many offline labor markets still reflected in self-determined bill rates in the online gig economy? In asking this question, this paper presents the first analysis of the role of gender in rate-setting in a single online labor marketplace in which workers set hourly bill rates. We examined data for 48,019 workers in the United States (US) on Upwork [10], one of the largest online labor marketplaces in both the US and the world [50].

We found that key offline inequalities in pay also exist on Upwork. The median woman on Upwork requested only 74% of what the median man requested in hourly bill rate. Controlling for job category, offline and online work experience, and highest level of education, the average treatment effect of being a woman was a \$6.28 reduction in hourly bill rate. However, despite this gap in hourly bill rate, when taking into account the total number of hours worked, we found that the total earnings of women and men on Upwork were roughly equal. Women worked more hours in total (median = 48.8) than men (median = 32.5), indicating that some women may have successfully attracted employers with their lower bill rates to earn more revenue. It also means that the cost of the bill rate gap across genders is borne in time, not in money. We discuss future qualitative work to explore these dynamics, as well as important sociotechnical interventions that leverage the transparency of online work platforms to improve gender equality in expected online bill rates.

2 BACKGROUND

2.1 Factors Influencing the Gender Pay Gap in the Offline Labor Market

In the US, women and men have not historically participated nor been treated equally in the paid labor market [80]. As noted in the introduction, female workers in the US earned just 82% of male workers' median weekly earnings and 87% of median hourly earnings in 2016 [72]. Three major factors that contribute to this gender pay gap include women working fewer hours in a day than men, employers penalizing flexible work arrangements, and differences in pay expectations (e.g., [19, 38, 61]).

With only 24 hours in day, women can work for pay during those hours when they are not working (without pay) for their families. Because women tend to do more of this unpaid labor than men [72], men have the capability to work more hours for pay in a day. Furthermore, given the often sporadic and unpredictable demands of domestic responsibilities on their time and location, women tend to need more flexibility about when and where they work for pay (e.g., [17, 19, 38, 61, 79]).

Unfortunately, workers who work fewer hours and/or have more flexible work arrangements are disproportionately penalized by their employers. For example, part-time workers can expect to receive lower compensation per hour than full-time workers in the same job [56]. Researchers speculate that this may be due in part to differences in productivity; working full-time may be related to higher productivity [56]. Within certain job categories, workers who desire flexibility may receive significantly less pay than their colleagues because workers are less easily substituted with one another in these professions [38]. For example, business and finance jobs disproportionately reward workers who are able to work more than 50 hours per week on average throughout the year [38]. This is because such workers are costly to replace. With each hour worked, they accrue institution-specific work experience and increase in value to their employer. If women initially work such long hours, but take temporary leave from full-time work before and after childbirth, they can still fall behind their male peers. Indeed, before and immediately after childbirth, the pay gap between a woman and her spouse can double owing to this reduction in total work experience [27]. As workers become older, the gap in earnings between women and men increases, partially owing to these differences in family responsibilities [38].

However for jobs in other industries such as pharmacy, where workers are more easily substituted with one another, those who work more hours and continuously over time are not disproportionately rewarded [38]. Consequently, the pay between men and women is almost equal [39] and workers can work part-time without leaving the workforce and reducing their earning ability [38]. In short, employers in certain job categories that place a disproportionately higher value on work experience reward workers who have the ability to continuously work longer hours at a full-time job, putting women at a disadvantage for earning equal pay. Factors such as number of hours worked, job category, total work experience, and age are covariates with these differences in pay.

2.1.1 Differences in Pay Expectations. The gender pay gap may also be explained by differences in pay expectations that are influenced by job category, negotiation behaviors, and perceptions of ability. Women are more likely to choose undergraduate majors that lead to roles in lower-paying job categories such as education [24], while men are more likely to choose majors that lead to roles in higher-paying job categories, such as computing and engineering [47, 72]. Aside from differences in major selection and resulting job category, women are less likely to ask for as much pay as men because they fear such asks will result in negative workplace evaluations. And their fear is real; experiments have shown that women who negotiate their salaries are seen as less hireable by men [15, 16, 21]. Women also often have lower expectancies than men in their ability to complete tasks (e.g., [28]) and hence undervalue their work. An online marketplace for technology jobs, Hired.com, showed professionals on its website the ranges of earnings for various job categories, women were more likely to set their own expected salary at the lower end of the range on the site, increasing the gender pay gap [43]. In short, gender differences in job category, negotiation behaviors, and perceptions of ability may also influence women's pay expectations and the gender pay gap.

2.2 Women's Participation and Pay in the Online Gig Economy

2.2.1 The Growing Online Gig Economy. Tens of millions of men and women take part in the growing online gig economy [50]. The online gig economy supports online marketplaces for short-term labor. Platforms act as marketplaces for transactions of short-term labor, such as remote staffing

(e.g., Freelancer [11]) and on-demand location-based jobs (e.g., Uber [5]) [53, 54, 71]. The World Bank estimated that workers' earnings aggregated across online labor marketplaces would reach USD\$4.8 billion worldwide in 2016 [53]. Moreover, job growth in the six largest English-speaking online labor marketplaces increased by 25.5% from July 2016 to July 2017 [50].

While the online gig economy cannot solve the challenge that women engage in more unpaid labor than men, it offers women the opportunity to arrange when and where they work, and earn more equal pay to men. Online platforms allow workers to seek short-term gigs when they are able to, providing more flexible paid labor opportunities [50]. In the US, most women who participate in the gig economy report that the flexible gigs allow them to more easily attend to domestic responsibilities than conventional jobs [44]. Because workers can be hired quickly and on demand (e.g., [73]), online gigs may also make workers more substitutable with one another over time [38], which may reduce gaps in pay rates between those who take leave from short-term labor and those who do not. Online gigs also draw on a range of knowledge and skills, making them accessible to workers regardless of gender. Gigs can range from routine tasks, such as doing laundry on TaskRabbit [9] and transcribing an audio recording on MTurk [1], to creative tasks, such as developing a new product on InnoCentive [4] or creating a business plan on Upwork [10].

2.2.2 Women's Participation and Treatment in the Online Gig Economy. Despite the initial promise of allowing more women to participate in the paid labor force, emerging evidence of equal female-to-male participation in the online gig economy is mixed [44, 66]. On MTurk, a popular gig work platform specializing in routine tasks for small pay, more than 50% of American workers are female [66]. In contrast, a recent survey of all US freelancers, including those who may not have used gig work platforms, estimated that only 41% of freelancers were female, compared to 47% of workers in the entire US economy [48].

Moreover, gender inequalities in pay and workplace evaluation in the offline labor market may persist in online labor marketplaces. Researchers have found gender inequalities in both: 1) marketplaces that offer a fixed occupational context with platform-determined rates (e.g., driving for Uber [5]) and 2) marketplaces that offer a range of occupational contexts and varied rates determined by individual clients or workers (e.g., web development fees negotiated by workers on Upwork [10]). In marketplaces with platform-determined rates (i.e., Uber [5]), women still earn less than men overall due to behavioral differences in completing the work, such as how quickly they drive [29]. In some marketplaces with varied rates, such as Fiverr [3] and TaskRabbit [9], women receive fewer client reviews, influencing their position in search rankings and their employability [3, 29, 41].

While we are beginning to understand the dynamics of female participation in the gig economy, we still lack a key understanding about gender differences in pay outcomes, specifically in marketplaces where workers are able to determine their pay rates. Such marketplaces are increasingly used in social computing research and applications to crowdsource creative tasks (e.g., [73]), making it critical for us to better understand gender dynamics on these gig work platforms.

2.2.3 Platform Experience and Performance Affect Pay Rates in the Online Gig Economy. In contrast to the offline labor market, pay rates in the online gig economy may also be subject to additional factors, such as specific platform experience and performance. While workers on platforms such as iStockphoto [12] and MTurk are primarily motivated to earn discretionary income [22, 23, 49, 64], workers may be willing to work for lower pay to build platform work experience or be given more time to complete a task [58, 81]. Highly rated reviews from past clients on platforms such as Fiverr [3] and TaskRabbit [9] can also influence workers' visibility to future clients and their earning abilities [41]. On Upwork [10], workers can take skill tests that may improve their prospects of winning clients [7]. The platform also displays workers' feedback ratings; the difference between

positive and negative feedback ratings is used to calculate a worker's Upwork Job Success Score [6]. These additional factors increase the complexity of predicting online gender-pay dynamics based on our current limited understanding of the offline gender pay gap.

2.3 Research Questions

The above related work informed two specific research questions that we sought to explore:

- **RQ1: How do the hourly bill rates that women in the US set for themselves in online labor marketplaces compare to men's hourly bill rates?** This question seeks to examine hourly bill rates (and resulting overall revenue) at a high level, allowing us to understand if the gender-pay dynamics in online labor marketplaces are similar or different to those in the traditional US economy.
- **RQ2: How do gender representation and discrepancy in hourly bill rate compare across job categories in online labor marketplaces?** As noted above, prior work has found that job categories are an important consideration in gender-pay dynamics [19, 38, 47, 61]. Hence, we sought to repeat RQ1, but focusing individually on each job category in online labor marketplaces. For instance, are there different discrepancies in hourly bill rates in Information Technology (IT) and Networking jobs versus Administrative Support jobs? We additionally seek to put these numbers in the context of the equivalent statistics in the traditional US economy as defined by job categories surveyed by the US BLS [72].

3 METHODS

3.1 Studying Upwork as an Online Labor Marketplace

Upwork [10], formerly oDesk, is a useful platform for studying gender differences in hourly bill rates and earnings across diverse job categories and experience levels. Upwork is one of the largest online labor marketplaces by earnings in the world [14]. It attracts workers across the US with a variety of skills in areas such as writing, web and software development, and law [52, 73], allowing us to observe gender-pay dynamics across many job categories. In comparison to other gig economy platforms, such as Uber [5], workers have significant control of their hourly rates, which allows us to observe differences in rate-setting behaviors. Given the popularity of Upwork and that the majority of demand for online gig work originates from the US [50], we focus our study on US workers on the Upwork platform.

3.1.1 Expected Pay and Actual Pay on Upwork. While workers have control over parts of their profile, workers' actual pay on Upwork may differ from their expected hourly bill rate. On Upwork, workers either bid on individual job assignments that pay on a fixed or hourly basis, or are solicited directly by clients who find them through search. During the bidding and interview process with clients, workers may negotiate the terms of their contract. However, past research suggests that fierce competition in online labor marketplaces encourages workers to underbid and accept rates without bargaining [40]. Hence, it is likely that workers are being paid less than their expected hourly bill rate. However, due to the limitations of the data we are able to collect publicly from Upwork [10], we are only able to examine workers' expected bill rates, rather than their actual pay. Therefore, we focused our current study to examine only workers' rate-setting behaviors on Upwork.

3.2 Analyses Conducted on Upwork Data

To address RQ1, we used a combination of descriptive and causal analysis techniques to compare male and female workers' hourly bill rates, controlling for potential covariates. The covariates we considered were Upwork job category, work experience (including work experience outside of

Upwork), and highest level of education. We also attempted to detect workers' age range and use this as a potential covariate. However, age range was not used in our analyses due to low precision, which we explain below. To address RQ2, we compared median bill rate and gender representation in Upwork job categories to equivalent BLS job categories [72].

3.3 Data Collection

Adopting the approach of key prior work on the gig economy discussed above (e.g., [41, 76]), we used information gathered from publicly available user profiles to form our primary dataset. Specifically, we collected data from 55,518 public profiles of US workers on Upwork from December 13-31, 2017. This data includes information such as first name, requested hourly bill rate in US Dollars, employment history on and off Upwork, and education history. We also estimated how much workers have earned in total hourly revenue given the total number of hours they have worked on Upwork and their stated hourly bill rate. Our final sample included 48,019 workers (48.8% whom we identified as women) after removing 13.5% of the sample (7,499 workers) for whom we could not identify either gender or at least one of the key covariates listed in the next section.

3.4 Approach to Gender and Covariate Detection

Below, we describe our approach to detecting gender and its potential covariates with expected bill rate. We also report the number of workers for whom we could not identify either gender or a key covariate and excluded from the causal analysis.

3.4.1 Gender. Because Upwork does not collect information on gender, we inferred gender using the name-based technique implemented in the `genderComputer` Python package [74]. This is a technique and a software package that is widely used in social computing, computational social science, and beyond (e.g., [26, 37, 41, 57, 78]), with a reported precision of 0.95 (recall = 0.88) [8, 74]. `genderComputer` compares an input first name against lists of first names that are highly likely to belong to men and to women, and does so on a country-by-country basis. In our case, `genderComputer` compared workers' names to a list of baby names from the 1990 US Census. If a name was given to a woman in at least twice as many cases as it was given to a man, `genderComputer` assumes the name belongs to a woman, and vice versa. Names that fell in between these two thresholds were assigned the label of `unisex`, whereas names not in the database were not assigned a gender (see below for more details). As in prior work (e.g., [37]), we did not include in our primary analysis workers whose gender was labelled as `unisex` (1.48% of sample, 821 workers) or could not be identified by `genderComputer` (5.04% of sample, 2,800 workers).

To verify that the `genderComputer` package did not have specific issues with the names in our dataset, we did an additional back-of-the-napkin verification of its results in which two independent raters labeled the perceived gender of 200 randomly selected workers in the dataset using profile pictures ($\text{Kappa} = 0.99$) [30]. Due to the high agreement of the raters, we chose one of the sets of labels as our ground truth dataset. The `genderComputer` classifier achieved similar average precision (0.92) and recall (0.89) against our ground truth dataset as it did for Vasilescu and colleagues [74], giving us additional confidence in its results.

While this name-based approach implemented in `genderComputer` is well established in the literature, gender identification is a challenging problem and this approach presents a necessarily limited view. Key limitations of studies that rely on this approach include an ignorance of non-binary genders (a particularly important limitation) and the inability to incorporate people with more ambiguous names. We mitigated some of these concerns by comparing the bill rates of workers who were identified as women on Upwork with those whose names were labeled as `unisex` or could not be identified (see below). We discuss these issues further in the Limitations section.

3.4.2 Job Category. To capture workers' job category, we used their primary job category as indicated on Upwork. We removed from further analysis the 507 (0.91%) of workers who did not indicate a primary job category.

3.4.3 Work Experience. In our analyses, we accounted for both the offline and online work experience of Upwork users using different measures. We estimated offline work experience using years of offline work experience and number of skills listed on a worker's profile. We calculated years of offline work experience by measuring the total, non-overlapping years of work completed as reported on the worker's employment history. Because workers are not required to fill out this section of their profile, we were careful not to classify workers who did not report work experience as having no work experience at all. Therefore, we removed workers who did not report any work experience from our analysis (0.09% of sample, 48 workers). One possible limitation of taking this approach is excluding workers who may legitimately have no prior work experience off Upwork, although that number would be very small and unlikely to affect our overall conclusions.

Given that users of a gig economy platform may need to build an online reputation to find online work [58, 81], online reputation may be especially important, so we use an array of signals to quantify online work experience. To account for online work experience, we additionally captured the following factors: number of skill tests completed through the Upwork platform (e.g., tests of English grammar, XML), number of portfolio items listed by the worker, average feedback rating, total number of feedback ratings received, number of hours worked on hourly assignments, number of assignments completed total, and tenure (in days) on Upwork. Unfortunately, Upwork's public API does not allow us to collect information on workers' Job Success scores; nevertheless, we expect that average feedback rating will capture at least some of the effects of these scores.

3.4.4 Level of Education. We inferred highest level of education based on the education history reported on workers' profiles. For each item in workers' education history, we extracted the institution name, degree, and any comments added by the worker, and classified workers into four levels of education determined by the US BLS [72] using a set of rules (see Table 1 for a summary). We could not identify the education history for 3.04% of the sample (1687 workers) and 3.90% of the sample (2167 workers) did not specify any education history. These workers were removed from consideration for our final analyses. We were also unable to determine education level for overseas institutions whose names and degrees conferred may be in a different language. For the remaining workers, our simple keyword-based classifier outlined in Table 1 obtained high average precision (0.93) and recall (0.94) compared to a ground truth dataset of 200 randomly sampled workers. This dataset was created by one independent rater who looked at workers' profile descriptions and determined highest level of education using four levels determined by the US BLS [72].

3.4.5 Age Range. It is possible age has an influence on pay gaps; as noted in Background, as workers age, women tend to earn less on average than men of the same age [38]. We sought to determine workers' age ranges as a potential covariate for our analyses, but the data and methods available were insufficiently reliable. We used Microsoft's Face API to determine age based on workers' public profile photos, an emerging approach to estimating age ranges when other data is not available (e.g., [55]). The API returns an estimated age based on the face detected in a photo. To ensure faces were large enough to be detected, we resized photos using the OpenCV Python package. Unfortunately, the average performance of the age classifier was very low (precision = 0.31, recall = 0.27) on an initial ground truth dataset. This dataset was created by one independent rater who looked at workers' profile descriptions and determined perceived age range using ranges determined by the US BLS [72]. As such, we did not include age as a covariate in our causal analysis. Our expectation is that years of work experience will capture at least some of the effect of age.

Level of Education	Institution Name	Degree	Comments Field
Bachelor's degree or higher	"university," "institute," or "college"	Bachelor's degree, master's degree, doctorate, or professional degree	<i>No keywords used</i>
Some college or associate's degree	"university," "institute," or "college"	None of the degrees above listed	<i>No keywords used</i>
High school degree	"high school" or "high"	High school degree or diploma, or stated the word "graduated"	<i>No keywords used</i>
Less than a high school degree	"high school" or "high", or none mentioned	None mentioned	"self-taught" or "homeschool"

Table 1. A description of keywords used to infer level of education for items in workers' education history.

3.5 Causal Analysis

To investigate possible causal relationships between gender and requested hourly bill rate, we used causal analysis techniques [45]. Causal analysis has frequently been used to approximate the effect of a treatment variable on an outcome variable in non-experimental settings by comparing data points that have similar covariate values (e.g., [34, 75]). As noted above, we focused on four groups of variables that have been identified as particularly important covariates of earning rates in traditional labor markets (see Background section): job category, number of years of offline work experience, highest level of education, and measures of online work experience through the Upwork platform.

3.5.1 Propensity Score Stratified Regression. We used the open source CausalInference [2] Python package to run a matched propensity-score-stratified regression [45]. Propensity scores reflect how likely a worker in the dataset is to "receive treatment" (i.e., be female) based on the values of the worker's covariates. To account for possible imbalance in covariates in the dataset, we trimmed the sample of workers with very high propensity scores, which would prevent us from finding equivalent workers who did not "receive treatment" (i.e., were not female; [33]). In other words, the final regression did not include workers with propensity scores very close to 0 or 1, because these workers cannot be compared fairly to workers from the opposite treatment status. Following Rosenbaum and Rubin's [65] original approach to propensity score stratification, we stratified our dataset into five subsets based on similar propensity scores to simulate a randomized blocked trial. Through blocking, we compared male and female workers with similar propensity scores (and therefore similar covariate values) to estimate the average "effect" of being female on requested bill rate. In total, our dataset for the causal inference analysis included 47,396 workers (49.5% female), after removing workers for whom we could not identify work experience, job category, or education and trimming workers with extreme propensity scores.

3.5.2 Validity Checks. To check the validity of our causal analysis, we also repeated the analysis with different modeling choices: we performed an analysis without trimming extreme propensity workers and also performed an analysis by matching workers one-to-one using nearest neighbor covariate matching instead of propensity score stratification. In each case, we verified that observed effects were roughly consistent in direction and scale to that which we report below. This suggests our results are robust to small model changes (including extreme propensity workers) and large

model changes (matching directly on covariates with replacement). For each analysis, we also inspected the covariate balance after matching: in each case, the standardized bias (the difference in means divided by the pooled standard deviation for a single covariate), a quantitative measure for assessing covariate balance, was substantially reduced, and no covariates in any strata had a standardized bias of greater than 0.25, a recommended threshold for "large" differences [69]. We report results from the propensity score blocked regression with extreme propensity scores trimmed, as this method strikes a balance between minimizing covariate bias without removing many workers from the analysis.

3.6 Online and Offline Comparisons

To compare gender representation and hourly bill rate discrepancy in online and offline work settings, we first matched job categories on Upwork to job categories surveyed by the BLS [72]. Job categories on Upwork are not automatically aligned with categories surveyed by the BLS, which are classified according to the 2010 Standard Occupational Classification (SOC) system. Following prior work in epidemiology that also involved labeling large occupational datasets, we mapped job categories on Upwork to equivalent SOC job categories using the Standardized Occupation Coding for Computer-assisted Epidemiologic Research (SOCcer) tool [32, 67]. This tool returns the 10 closest SOC job categories for each Upwork job title and job description given. Each of the returned SOC categories is given a probability score that reflects how close the Upwork job title and description is to the matched category. In this case, we used as job titles the names of the 12 primary job categories on Upwork. The input job descriptions were the names of the sub-categories within each Upwork job category.

Not all of the top matched job categories had complete reported earnings data in the BLS survey. Therefore, when earnings data was not available for the top match, we took the subsequent match with data available and were able to find matches for all job categories except for the Translation category. For example, the Customer Service job category on Upwork was matched to "customer service representative occupations" in the SOC system and Accounting and Consulting on Upwork matched to "accountants and auditors" in the SOC system.

4 RESULTS

In this section, we discuss the gender differences in hourly rate-setting overall (RQ1) and between occupations (RQ2) for independent workers in the US in the online labor marketplace, Upwork.

4.1 RQ1: Measuring Gender Differences in Hourly Rate-Setting Across All Occupations

We began our analyses by examining the distribution of hourly rates in our dataset. Bill rates showed positive skewness (7.52, $z = 213.12$, $p < 0.001$) and a high kurtosis (131.18, $z = 140.41$, $p < 0.001$). About 3.8% of our sample (1797 workers) set bill rates at least 2 standard deviations above the mean; among these outliers, the majority were male (70.6%) and the median bill rate was \$150.00 per hour. Given these results, we either performed non-parametric tests or normalized our data in the remainder of our analyses.

4.1.1 Women Asked for Lower Median Hourly Bill Rates. Without taking other covariates into account, female workers asked for lower median hourly bill rates overall (\$26.00, mean (M) = \$35.61, standard deviation (SD) = \$31.77) than male workers in our dataset (\$35.00, $M = \$46.35$, $SD = \$44.16$; Table 2 shows more details). In other words, female workers in the US on Upwork asked for 74.3% of what male workers asked for in median hourly bill rate and 76.8% in mean hourly bill rate. By (loose) comparison, in the broader economy, the US BLS estimated that female wage and salary

Gender determined by genderComputer	Median Hourly Bill Rate (\$)	Mean Hourly Bill Rate (\$)	Standard Deviation Hourly Bill Rate (\$)
Male	35.00	46.35	44.16
Female	26.00	35.61	31.77
Unisex	30.00	37.51	34.58
Not identified	28.00	37.97	41.49

Table 2. Women in our final sample asked for \$9.00 less in median and \$10.74 less in mean hourly bill rate than men.

workers who were paid by the hour earned 87% in median hourly earnings relative to those of male workers in 2016 [72].

The results of our causal analyses also suggest a marked difference in expected Upwork hourly bill rate based primarily on gender. After matching workers on job category, years of offline work experience, highest level of education, and online work experience, our causal analysis showed an average treatment effect on the control estimate (i.e., male workers) of \$6.28 ($p < 0.001$). This means that in a counterfactual world in which male workers were instead female, they would experience a \$6.28 or 13.6% reduction on average in mean hourly bill rate (M for male workers included in this analysis = \$46.08).

We verified that the above results were not driven entirely by outlier hourly rates by rerunning our analysis with outliers removed (i.e., workers with bill rates at least 2 standard deviations above the mean). Given that there were male users with outlier bill rates (70.6% of outliers were male) that drove up the mean bill rate, we might expect to see a diminished effect size when removing outliers. The overall results were as expected: we saw a moderate decline in average treatment effect on the control (\$3.46, $p < 0.001$). This is equal to an 8.8% reduction in mean hourly bill rate of male workers included in this analysis ($M = \$39.21$). The presence of a non-trivial and significant negative effect suggests that outliers were responsible for only a portion of the pay discrepancy we observed.

4.1.2 No Significant Evidence of Women Using Unisex or Unidentifiable Names to Set Higher Hourly Rates. We also explored the possibility that some women were using unisex or unidentifiable names as a strategy for setting higher bill rates. Workers with a unisex or unidentifiable name had a \$2.26 higher mean bill rate than female users on average. However, when we used a propensity-score blocked regression, which is the same approach described above in the Methods, but with female users as the control and unknown gender users as the treatment group, the estimate did not show a statistically significant difference in hourly bill rate compared to having a female name ($p = 0.413$). Therefore, we found no significant support for the notion that women are using unisex or unidentifiable names as a way to ask for substantially higher hourly bill rates on Upwork. We discuss future qualitative work to further explore this possibility in the Discussion.

4.1.3 Hourly Bill Rates versus Overall Earnings. While the above analyses show that women ask for less per hour than their male counterparts, this does not necessarily mean that they earn less overall on Upwork. In particular, women could earn less per hour but work more hours, translating at least some of the hourly rate difference into time costs rather than financial costs and/or perhaps implementing an undercutting strategy. Indeed, comparing workers who had worked more than zero hours on the platform (9747 workers, 52.5% female), the median total number of hours worked

by women was 48.83 hours ($M = 436.55$, $SD = 1,325.70$), while the median total number of hours worked by men was 32.50 hours ($M = 335.42$, $SD = 1049.90$).

To understand the impact of these increased hours worked on overall revenue, we estimated the total hourly revenue workers have earned by multiplying each worker's hourly bill rate with their total number of hours worked on hourly assignments on Upwork. Looking again at workers who have logged more than zero hours of work on the platform (9747 workers, 52.5% female), the median US woman appears to have earned slightly more than the median US man on Upwork. The median hourly revenue women earned was \$1,386.67 ($M = \$13,382.39$, $SD = \$40,896.30$), while the median hourly revenue men earned was \$1,278.08 ($M = \$15,276.44$, $SD = \$52,900.29$). However, a Mann-Whitney U test did not find a significant difference between these distributions ($z = 1.18 \times 10^7$, $p = 0.23$). Additionally, transforming the data to a normal distribution using Box-Cox transformation and running a t-test did not show a significant difference between male ($M = 6.73$, $SD = 1.96$) and female ($M = 6.74$, $SD = 1.94$) transformed total hourly revenue ($z = 0.27$, $p = 0.79$).

These results suggest that, in general, the cost of the Upwork hourly rate gap for women may not be money, but rather time. We observed that women earn less per hour, but work more hours, leading to approximately the same overall earnings between the genders.

4.2 RQ2: Gender Differences Within and Across Job Categories

Because job category plays such a large role in the gender pay gap offline (e.g., [38]), our second research question seeks to explore these same dynamics in more detail on Upwork. As noted above, for our second research question, we use as a baseline patterns of gender representation and pay offline from a survey completed by the US BLS in 2016. At a high level, we find that more workers on Upwork were female and attained higher education levels than those in the BLS survey. Compared to the BLS dataset ($N = 111,091,000$, 44.3% female), our final sample ($N = 48,019$, 48.8% female) was slightly closer to gender parity. Workers in our final sample were also more likely to have at least a bachelor's degree compared to individuals 25 years and older in the BLS dataset ($N = 101,015,000$); nearly 80% of workers in our sample reported having at least a bachelor's degree, while just more than 40% of workers surveyed by the BLS had a bachelor's degree. Because the BLS survey did not collect data on work experience, we were not able to compare our datasets on this measure.

4.2.1 Women Were Overrepresented in Jobs With the Lowest Median Hourly Bill Rates. We found that women on Upwork tended to participate most in job categories with the lowest median hourly bill rates. We observed a strong negative correlation between female representation and median hourly bill rate between job categories ($r = -0.79$, $n = 12$, $p < 0.01$). Additionally, women outnumbered men in the four lowest-paying job categories: Writing, Translation, Administrative Support, and Customer Service (see Figures 1 and 2). Additionally, men outnumbered women in the highest-paying job categories (e.g., IT and Networking; Engineering and Architecture; and Web, Mobile and Software Development).

4.2.2 Representation of Women in Job Categories on Upwork Aligns with BLS Dataset. Are the differences in gender representation between job categories on Upwork reflected in the offline labor market? Our data suggest that the representation of women in job categories on Upwork aligns with the representation of women in equivalent occupations offline in the BLS dataset. We found a strong positive correlation between percentage of women in job categories on Upwork and percentage of women in the equivalent category in the BLS survey ($r = 0.82$, $n = 11$, $p < 0.01$, see Figure 3). In the BLS survey, women were most highly represented in Administrative Support (secretaries and administrative assistants; 94.0%), Customer Service (customer service representatives; 64.1%), Accounting and Consulting (accountants and auditors; 61.5%), and Translation (miscellaneous media and communication workers; 61.0%). They were least represented in Engineering and Architecture

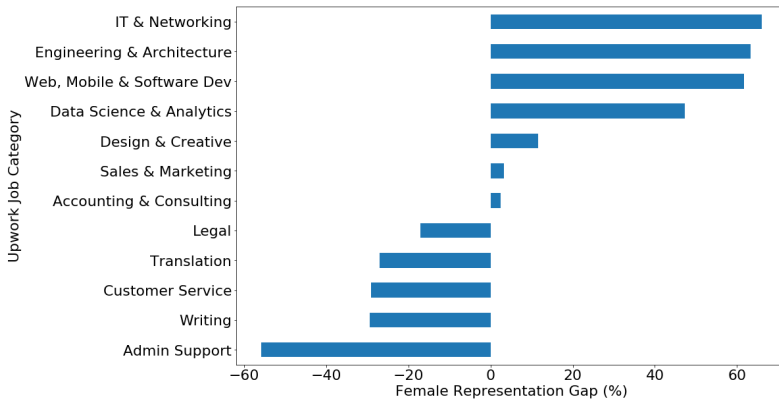


Fig. 1. Representation of female workers varied across job categories on Upwork, with women being over-represented in Administrative Support and underrepresented in IT and Networking and Engineering and Architecture.

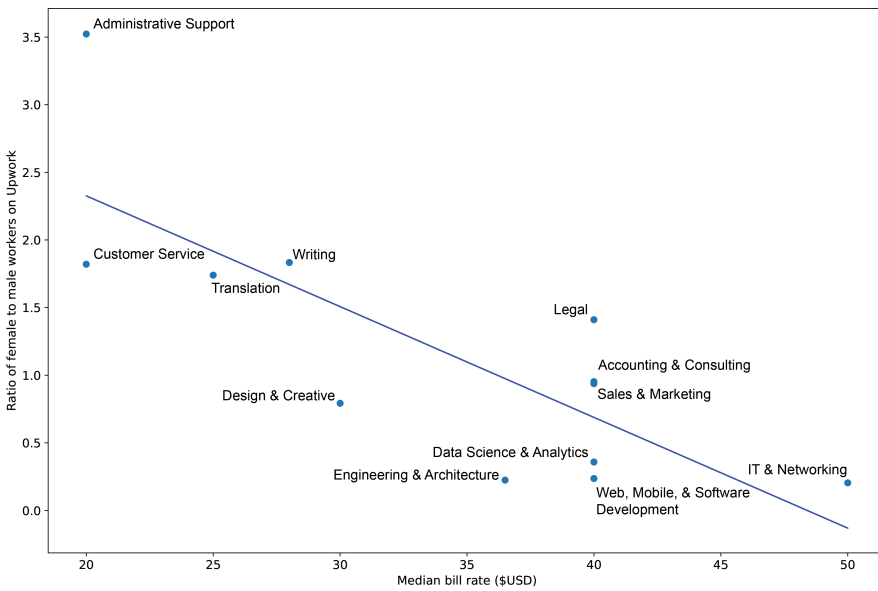


Fig. 2. The job categories with the lowest median hourly bill rate - Writing, Translation, Administrative Support, and Customer Service - also had among the highest ratios of female to male workers on Upwork ($r = -0.79, n = 12, p < 0.01$)




Fig. 3. Across most domains, the percentage of women in job categories on Upwork reflected the percentage of women in equivalent job categories in the BLS survey ($r = 0.82$, $n = 11$, $p < 0.01$). We were unable to find equivalent BLS job category matches for Translation and did not include this in our analyses.

(engineers, all other; 12.8%), IT and Networking (computer and information systems managers; 25.3%), and Web, Mobile, and Software Development (computer programmers; 25.8%) categories.

4.2.3 Women Asked for Lower Hourly Bill Rates Compared to Men. Not only were women less represented in higher-paying job categories on Upwork, they also asked for lower hourly bill rates compared to men within job categories. Following the same approach as our causal analysis described above, we used propensity score blocking and covariate nearest neighbor matching with and without extreme propensity trimming to estimate the treatment effect on the control for workers within each job category. When comparing workers within job categories in which women were underrepresented, these analyses suggested that women still asked for significantly lower rates. Specifically, after estimating the treatment effect within each job category, we found that in all but two categories (Translation; Engineering and Architecture), at least one causal model suggested there was a statistically significant negative treatment effect on the control. For example, within the job category of Web, Mobile, and Software Development, our causal analysis estimated a treatment effect on male workers of \$6.14. We note that unlike our primary causal analysis with the full dataset, which was extremely robust to model changes, the analyses within each job category were not robust to model changes (i.e., sometimes blocking estimated a significant effect but matching did not) and for some categories with especially imbalanced gender representation, blocking was

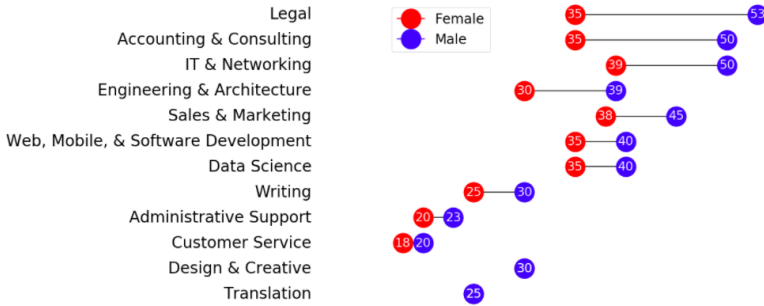


Fig. 4. In all but two job categories, women asked for lower median hourly bill rates than men. The gender bill rate gap was highest in Legal, Accounting & Consulting, and IT & Networking categories. There was no gender gap in median bill rate in Translation and Design and Creative categories.

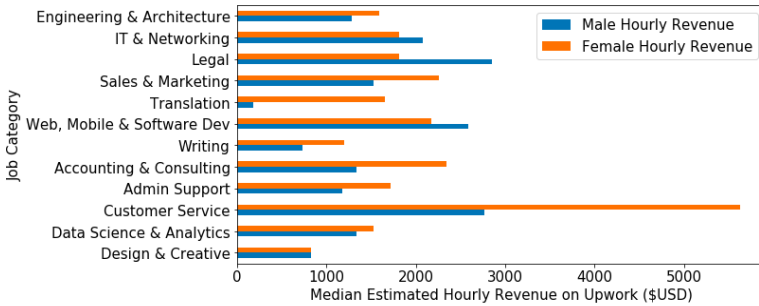


Fig. 5. In eight Upwork job categories, women earned more in estimated hourly revenue than men. We continue discussing these complex dynamics below.

impossible or resulted in increased covariate bias. However, the combination of our descriptive analyses (i.e., comparison of median wages shown in Figure 4), our primary causal analysis, and these job-specific causal estimates strongly suggests that the gender bill rate gap on Upwork is not being driven by choices related to job category alone.

4.2.4 Women in Some Job Categories Earned More Revenue than Men. Importantly, we further estimated the total hourly revenue workers earned in different job categories. Women worked sufficiently more hours to surpass men in median total hourly revenue in eight job categories: Data Science and Analytics; Customer Service; Administrative Support; Accounting and Consulting; Writing; Translation; Sales and Marketing; and Engineering and Architecture. Men still earned more in hourly revenue than women in four categories: Design and Creative; Web, Mobile, and Software Development; Legal; and IT and Networking (see Figure 5). We examine this relationship between hourly bill rates and revenue and its implications in the Discussion section.

5 DISCUSSION

5.1 Equal Pay for Equal Work in Online Labor Marketplaces?

Through our analyses, we identified a relationship between gender and hourly bill rates in the online labor marketplace, Upwork; the median woman on Upwork asks for significantly less per hour than the median man on Upwork. This difference persists even when controlling for key

covariates such as offline and online work experience, highest education level, and job category. Our results suggest that online labor marketplaces, though providing flexible, on-demand work opportunities, may continue to be shaped in part by the same gender dynamics that we see in the traditional offline labor market. By examining self-determined hourly bill rates, our results raise the possibility that women on Upwork may be undervaluing their work compared to men, either intentionally (to increase their prospects of employment) or unintentionally.

5.2 Lower Hourly Rates and Higher Overall Revenue?

We found that women's lower hourly bill rates are coupled with a higher number of total hours worked. Indeed, in eight of the 12 job categories, the differential in total hours worked on hourly assignments across genders is sufficiently large that the median woman earns more total (estimated) revenue than the median man. While these results are subject to limitations in our ability to accurately approximate total revenue (see above), this finding raises the possibility that some women may be deliberately setting lower prices for their time to increase their prospects of employment, and are doing so sufficiently effectively that it increases their overall revenue earned.

Regardless of intent, these actions come at a cost. Women on Upwork are not selling products, they are selling their time. More hours worked, therefore, equates to less time available for other pursuits. And unlike the number of products that can be produced in many cases, the number of hours available per worker is highly constrained. This is particularly concerning as research has identified that one key attraction for women to gig work over traditional labor arrangements is time to attend to family responsibilities [44]. Our results clearly indicate that future qualitative research is needed to better understand workers' decision-making processes when setting their hourly bill rates and deciding how many hours to work.

The relationship between hourly rates and hours worked that we observed also presents another intriguing possibility: men may be overestimating their value. Informally, we observed a higher likelihood of embellishment in men's self-descriptions of their work experience on their profile pages. It is possible that this embellishment extends to hourly bill rates, and it may be hurting men in online marketplaces. This is also a direction worthy of more study, and we expand on further directions for future work below.

5.3 Future Work

5.3.1 Rate-setting Success in Online Labor Marketplaces. Future qualitative and quantitative work is necessary to understand hourly rate-setting success and strategies in online labor marketplaces, and whether those strategies are economically optimal for each individual. For example, although our study suggests that women were unlikely to have used unisex names as a way to set higher hourly bill rates, a future study should specifically interview male and female workers about the factors they consider when determining their hourly rates. Such studies will ideally be collaborations between social computing researchers (who understand the gig economy context), behavioral economists (who understand the relationship between human choices and economic principles), and industry collaborators (who can easily access data).

5.3.2 Understanding the Role of Online Work Experience. Our results also suggest that more research is needed to understand the impact of online work experience on worker success. Despite having higher levels of education compared to workers surveyed by the US BLS, women on Upwork still asked for significantly less pay than men. This is intriguing, given that past research suggests higher levels of education correlate with smaller gender pay gaps [61]. However, these results make more sense when considered in the context of online work. Educational qualifications may play a relatively small role in online labor marketplaces when compared to offline work performance [40].

Perhaps success in these marketplaces is subject to different signals of work experience beyond those used in the traditional labor market. Future work is needed to determine the extent to which online work experience leads to work success compared to more "traditional" signals such as education level and offline work experience.

5.3.3 Hourly Bill Rate Dynamics by Worker Status and Country. Future work should be devoted to understanding these results in the context of workers' full-time or part-time status and their countries of origin, as these may also influence levels of pay. In offline labor markets, workers can expect to receive significantly lower compensation per hour for the same job when they work part-time compared to full-time [56]. Hence, we predict that workers who intend to use Upwork as their primary source of income will set higher bill rates than those who do not. Furthermore, because we limited the scope of this analysis to workers in the US, our results may not generalize to workers in other countries. In most countries, women earn less on average than men, and gender differences in pay vary by country [13] due in part to differences in women's participation in the labor force [60]. In countries such as Ireland and France where women earn almost as much as men, fewer women participate in the labor force to begin with compared to countries with larger pay gaps, such as the US and the United Kingdom [60]. More work is needed to understand gender discrepancies in bill rate on Upwork across the globe, as well as how the transparent flow of work *between* countries (i.e., workers working with international clients [52]) further influences these discrepancies.

5.3.4 New Socio-Technical Interventions to Increase Equality. Our findings also highlight the possibility for interesting new socio-technical interventions in online labor marketplaces. Although independent contractors in online labor marketplaces may not be subject to the same legal protections as employees in the US, experiments have shown that they have similar expectations for fair compensation [25] and may find value from tools that increase equality in pay. For example, platforms could present data analytics to help employers understand their hiring practices for men and women in online labor marketplaces and adjust these practices to increase equality. Researchers could also develop per-worker pricing strategies to help individuals maximize the value of their work time while achieving other goals and responsibilities (e.g., preserving time for family responsibilities and professional development) [51].

5.3.5 Gender Dynamics in Computer-Supported Cooperative Work (CSCW). Emerging research suggests that women face barriers to participation in various socio-technical systems, including peer production communities (e.g., Wikipedia) and social Q&A sites (e.g., StackOverflow) [31, 37, 42]. Yet, as feminist human-computer interaction (HCI) scholars point out, few socio-technical systems are designed and studied with gender differences in mind [18, 62]. In this study, we find that gender is a robust factor in predicting bill rate that needs to be considered in the design of future online labor marketplaces. Gender influences workers' bill rates partly by replicating patterns of gender representation in offline job categories, but possibly also through more nuanced online mechanisms. As mentioned earlier, previous qualitative research with workers in online labor marketplaces from different countries has shown that workers often undervalue their services as a way of securing employment[40]. Women on Upwork may be doing the same by underbidding as a strategy to gain more employment online. As CSCW and HCI researchers, we must continue to study the role of gender, not only as it relates to participation in socio-technical systems such as online labor marketplaces, but also as it relates to strategies for success in these systems [59, 77].

5.4 Questions of Responsibility

Our research raises a difficult question: whose responsibility is it to create more equal online labor marketplaces? On the one hand, our results suggest that factors associated with the bill rate gap on Upwork include structural differences in society (e.g., representation across job categories), the choices made by men and women in pricing their time, and the amount of time men and women work. All of these factors stand somewhat apart from the Upwork platform itself. However, our work also highlights opportunities for data transparency that have not yet been leveraged by Upwork. For instance, as noted above, Upwork could implement specific pricing guidelines for women to increase visibility of the bill rate gap, and analytics for employers to understand hiring practices. Trivially, Upwork could also show women what men with the most similar characteristics in the most recent job are charging using approaches similar to our causal analysis.

Researchers interested in designing online labor marketplaces to support more equal bill rates between men and women should consider intervening in workers' bill rate decision making process. For example, platforms could provide more explicit access to other workers' bill rates. In prior research, descriptive norms, or information about what others do [63], have been effectively used to influence people's eating [68] and pro-environment behaviors [35]. Nonetheless, industry research suggests efforts must go beyond merely providing access to other workers' bill rates, as women still tend to place themselves at the lower end of the pay spectrum [43]. Workers may also be hesitant to negotiate higher rates when they are compared against other workers [40]. Further qualitative work is needed to understand differences in rate-setting and the frequency with which workers update their rates based on experiences gained, particularly in categories where gender bill rate gaps are the highest (e.g., Legal; IT and Networking; Accounting and Consulting).

6 LIMITATIONS

Although our quantitative analysis provides insight into hourly bill rates for US workers across job categories on Upwork, we do not have insight into how many fixed-price jobs workers have completed and at what rates. We also cannot see the outcome of any bidding processes. These data could substantially affect our above results. However, access to these data are restricted on Upwork.

One of the major limitations of our approach is our reliance on perceived gender. True gender is difficult information to capture, as Upwork does not collect information on workers' genders during registration. Furthermore, our approach to gender inference is limited to binary gender identification, and may exclude non-binary genders that are more easily expressed in online systems that allow for pseudonymous and nuanced identities (e.g., Archive of Our Own (AO3) [36]). Future work should be devoted to verifying the current findings with surveys of smaller representative samples of workers across online labor platforms who provide self-reported gender.

Lastly, while the gender pay gap differs based on age [38], we did not include age range as a potential covariate in our analyses because of the low reliability of our age range classifier. Other research on MTurk suggests that workers in online labor marketplaces in the US tend to be younger ($M = 33\text{-}35$ years, [66]) than the population surveyed by the US BLS [72]. Therefore, the age of our final sample may have also contributed to the observed bill rate gap between men and women.

7 CONCLUSION

As technology increasingly supports work across the globe, online labor marketplaces have been of increasing interest to CSCW and social computing researchers (e.g., [70, 73]). Our results lead our understanding of the relationship between gender, bill rate-setting, and earned revenue in these marketplaces. In this study, we found that women ask for only 74% of what men ask for in median hourly bill rate on Upwork, one of the largest online labor marketplaces in the world. However, we

also found that by working more hours at these lower rates, women earn as much or more than men overall. The current work contributes to a developing body of research by gender studies scholars in CSCW and social computing who study barriers to participation in other social technologies, such as peer production sites, online communities, social media, and blogs [31, 42, 59]. As scholars, we must better understand how women decide on hourly rates and hours worked if online labor marketplaces are to more equally engage women in the paid labor force.

ACKNOWLEDGMENTS

This work was generously supported by the Segal Design Institute Design Research Cluster, the School of Communication at Northwestern University, and NSF IIS-1707296. We would like to thank our anonymous reviewers, members of the Delta Lab at Northwestern University, Paola Sapienza, Ho Kwan Cheung, Denae Ford, Jamie Gorson, Emily Wang, Diego Gomez Zara, and Daniel Shams for their useful comments, suggestions, and support.

REFERENCES

- [1] [n. d.]. Amazon Mechanical Turk. ([n. d.]). <https://www.mturk.com/>
- [2] [n. d.]. Causal Inference in Python: CausalInference 0.1.2 documentation. ([n. d.]). <http://causalinferencepython.org/>
- [3] [n. d.]. Fiverr - Freelance Services Marketplace for The Lean Entrepreneur. ([n. d.]). <https://www.fiverr.com/>
- [4] [n. d.]. InnoCentive | Open Innovation & Crowdsourcing Platform. ([n. d.]). <https://www.innocentive.com/>
- [5] [n. d.]. Making career moves? Sign up to be a Uber Driver or get a ride to the airport | Uber. ([n. d.]). <https://www.uber.com/>
- [6] [n. d.]. My Job Success Score. ([n. d.]). <http://support.upwork.com/hc/en-us/articles/211068358-My-Job-Success-Score>
- [7] [n. d.]. Skills Tests. ([n. d.]). <http://support.upwork.com/hc/en-us/articles/211063198-Skills-Tests>
- [8] [n. d.]. Stack Overflow - Where Developers Learn, Share, & Build Careers. ([n. d.]). <https://stackoverflow.com/>
- [9] [n. d.]. TaskRabbit connects you to safe and reliable help in your neighborhood. ([n. d.]). <https://www.taskrabbit.com/>
- [10] [n. d.]. Upwork - Hire Freelancers & Get Freelance Jobs Online. ([n. d.]). <https://www.upwork.com/>
- [11] 2018. Freelancer - Hire & Find Jobs. (2018). <https://www.freelancer.com/>
- [12] 2018. Stock photos, royalty-free images & video clips. (2018). <https://www.istockphoto.com/>
- [13] ActionAid. 2015. *Close the gap! The cost of inequality in women's work*. Technical Report. ActionAid. https://www.actionaid.org.uk/sites/default/files/publications/womens_rights_on-line_version_2.1.pdf
- [14] Ajay Agrawal, John Horton, Nicola Lacetera, and Elizabeth Lyons. 2015. Digitization and the Contract Labor Market: A Research Agenda. In *Economic Analysis of the Digital Economy*. University of Chicago Press, 219–250. <https://doi.org/10.7208/chicago/9780226206981.003.0008>
- [15] Emily T. Amanatullah and Michael W. Morris. 2010. Negotiating gender roles: Gender differences in assertive negotiating are mediated by women's fear of backlash and attenuated when negotiating on behalf of others. *Journal of Personality and Social Psychology* 98, 2 (2010), 256–267. <https://doi.org/10.1037/a0017094>
- [16] Linda Babcock and Sara Laschever. 2009. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press. Google-Books-ID: k6vd7PJUFmYC.
- [17] Roxana Barbulescu and Matthew Bidwell. 2012. Do Women Choose Different Jobs from Men? Mechanisms of Application Segregation in the Market for Managerial Workers. *Organization Science* 24, 3 (June 2012), 737–756. <https://doi.org/10.1287/orsc.1120.0757>
- [18] Shaowen Bardzell. 2010. Feminist HCI: Taking Stock and Outlining an Agenda for Design. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 1301–1310.
- [19] Francine D. Blau and Kahn, Lawrence M. 2016. *The Gender Wage Gap: Extent, Trends, and Explanations*. Technical Report 21923. National Bureau of Economic Research. 1–77 pages. <https://www.nber.org/papers/w21913.pdf>
- [20] Beth Ann Bovino and Jason Gold. 2017. *The Key to Unlocking U.S. GDP Growth? Women*. Technical Report. S&P Global Inc. https://media.spglobal.com/documents/Women_at_Work2.pdf
- [21] Hannah Riley Bowles, Linda Babcock, and Lei Lai. 2007. Social incentives for gender differences in the propensity to initiate negotiations: Sometimes it does hurt to ask. *Organizational Behavior and Human Decision Processes* 103, 1 (May 2007), 84–103. <https://doi.org/10.1016/j.obhdp.2006.09.001>
- [22] Daren C. Brabham. 2008. Crowdsourcing as a Model for Problem Solving: An Introduction and Cases. *Convergence: The International Journal of Research into New Media Technologies* 14, 1 (Feb. 2008), 75–90. <https://doi.org/10.1177/1354856507084420>

- [23] Daren C. Brabham. 2008. Moving the crowd at iStockphoto: The composition of the crowd and motivations for participation in a crowdsourcing application. *First Monday* 13, 6 (May 2008). <http://journals.uic.edu/ojs/index.php/fm/article/view/2159>
- [24] Sarah Green Carmichael. 2017. Women Dominate College Majors That Lead to Lower-Paying Work. (April 2017). <https://hbr.org/2017/04/women-dominate-college-majors-that-lead-to-lower-paying-work>
- [25] Daniel L Chen and John J Horton. 2016. Research Note: Are Online Labor Markets Spot Markets for Tasks? A Field Experiment on the Behavioral Response to Wage Cuts. *Information Systems Research* 27, 2 (June 2016), 403–423. <https://doi.org/10.1287/isre.2016.0633>
- [26] Le Chen, Ruijun Ma, Aniko Hannak, and Christo Wilson. 2018. Investigating the Impact of Gender on Rank in Resume Search Engines. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*.
- [27] YoonKyung Chung, Barbara Downs, Danielle H. Sandler, and Robert Sienkiewicz. 2017. *The Parental Gender Earnings Gap in the United States*. Technical Report 17-68. Center for Economic Studies, U.S. Census Bureau. <https://ideas.repec.org/p/cen/wpaper/17-68.html>
- [28] Pauline R. Clance and Maureen A. O’Toole. 1987. The Imposter Phenomenon: An internal barrier to empowerment and achievement. *Women & Therapy* 6, 3 (1987), 51–64. https://doi.org/10.1300/J015V06N03_05
- [29] Cody Cook, Rebecca Diamond, Jonathan Hall, John A. List, and Paul Oyer. 2018. *The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers*. Technical Report Working Paper No. 3637. <https://www.gsb.stanford.edu/faculty-research/working-papers/gender-earnings-gap-gig-economy-evidence-over-million-rideshare>
- [30] Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin* 70, 4 (1968), 213–220. <https://doi.org/10.1037/h0026256>
- [31] Benjamin Collier and Julia Bear. 2012. Conflict, Criticism, or Confidence: An Empirical Examination of the Gender Gap in Wikipedia Contributions. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW ’12)*. ACM, New York, NY, USA, 383–392. <https://doi.org/10.1145/2145204.2145265>
- [32] Roxanne Connelly, Vernon Gayle, and Paul S. Lambert. 2016. A Review of occupation-based social classifications for social survey research. *Methodological Innovations* 9 (April 2016), 2059799116638003. <https://doi.org/10.1177/2059799116638003>
- [33] Richard K. Crump, V. Joseph Hotz, Guido W. Imbens, and Oscar A. Mitnik. 2009. Dealing with limited overlap in estimation of average treatment effects. *Biometrika* 96, 1 (Jan. 2009), 187–199. <https://doi.org/10.1093/biomet/asn055>
- [34] Munmun De Choudhury, Emre Kiciman, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. Discovering Shifts to Suicidal Ideation from Mental Health Content in Social Media. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM Press, New York, New York, USA, 2098–2110. <https://doi.org/10.1145/2858036.2858207>
- [35] Leila Elgaaied-Gambier, Elisa Monnot, and Fanny Reniou. 2018. Using descriptive norm appeals effectively to promote green behavior. *Journal of Business Research* 82 (Jan. 2018), 179–191. <https://doi.org/10.1016/j.jbusres.2017.09.032>
- [36] Casey Fiesler, Shannon Morrison, and Amy S. Bruckman. 2016. An Archive of Their Own: A Case Study of Feminist HCI and Values in Design. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM Press, 2574–2585. <https://doi.org/10.1145/2858036.2858409>
- [37] Denae Ford, Alisse Harkins Harkins, and Chris Parnin. 2017. Someone like me: How does peer parity influence participation of women on stack overflow?. In *2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 239–243. <https://doi.org/10.1109/VLHCC.2017.8103473>
- [38] Claudia Goldin. 2014. A Grand Gender Convergence: Its Last Chapter. *American Economic Review* 104, 4 (April 2014), 1091–1119. <https://doi.org/10.1257/aer.104.4.1091>
- [39] Claudia Goldin and Lawrence F. Katz. 2016. A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation. *Journal of Labor Economics* 34, 3 (July 2016), 705–746. <https://doi.org/10.1086/685505>
- [40] Mark Graham, Isis Hjorth, and Vili Lehdonvirta. 2017. Digital labour and development: impacts of global digital labour platforms and the gig economy on worker livelihoods. *Transfer: European Review of Labour and Research* 23, 2 (May 2017), 135–162. <https://doi.org/10.1177/1024258916687250>
- [41] Aniko Hannak, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. Bias in Online Freelance Marketplaces: Evidence from TaskRabbit and Fiverr. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing*. ACM, New York, New York, USA. <https://doi.org/10.1145/2998181.2998327>
- [42] Libby Hemphill and Jahna Otterbacher. 2012. Learning the Lingo?: Gender, Prestige and Linguistic Adaptation in Review Communities. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW ’12)*. ACM, New York, NY, USA, 305–314. <https://doi.org/10.1145/2145204.2145254>
- [43] Hired.com. 2017. Women, Work, and the State of Wage Inequality. (2017). <https://hired.com/wage-inequality-report-2017>
- [44] Hyperwallet. 2017. *The Future of Gig Work is Female*. Technical Report. Hyperwallet. 18 pages. https://www.hyperwallet.com/app/uploads/HW_The_Future_of_Gig_Work_is_Female.pdf?mkt_tok=

- eyJpJjoiTVRjMU9UQmlOakk1TW1WaSIsInQiOjYVQRNEtTzUzNWliUzZOSTQ3R2wxTnlwY00xZG9MZmErTnVXUkJVdGhMRm9l3D%3D
- [45] Guido W. Imbens and Donald B. Rubin. 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press. Google-Books-ID: Bf1tBwAAQBAJ.
- [46] McKinsey Global Institute. 2017. *Technology, Jobs, and the Future of Work*. Technical Report. McKinsey Global Institute. <https://www.mckinsey.com/~media/McKinsey/Global%20Themes/Employment%20and%20Growth/Technology%20jobs%20and%20the%20future%20of%20work/MGI-Future-of-Work-Briefing-note-May-2017.ashx>
- [47] Institute for Women’s Policy Research. 2018. *The Gender Wage Gap by Occupation 2017 and by Race and Ethnicity*. Technical Report C467. Institute for Women’s Policy Research. https://iwpr.org/wp-content/uploads/2018/04/C467_2018-Occupational-Wage-Gap.pdf
- [48] Edelman Intelligence. 2017. *Freelancing in America*. Technical Report. Edelman Intelligence. <https://www.slideshare.net/upwork/freelancing-in-america-2017/1>
- [49] Lilly C. Irani and M. Six Silberman. 2013. Turkopticon: interrupting worker invisibility in amazon mechanical turk. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 2013)*. ACM Press, 611. <https://doi.org/10.1145/2470654.2470742>
- [50] Otto Kassi and Vili Lehdonvirta. 2016. *Online Labour Index: Measuring the Online Gig Economy for Policy and Research*. Technical Report 74943. <https://mpr.aub.uni-muenchen.de/74943/>
- [51] Aniket Kittur, Jeffrey V Nickerson, Michael Bernstein, Elizabeth M Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In *Proceedings of the 16th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, New York, New York, USA, 1301–1318. <https://doi.org/10.1145/2441776.2441923>
- [52] Marios Kokkodis, Panagiotis Papadimitriou, and Panagiotis G Ipeirotis. 2015. Hiring Behavior Models for Online Labor Markets. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*. ACM Press, New York, New York, USA, 223–232. <https://doi.org/10.1145/2684822.2685299>
- [53] Siou Chew Kuek, Cecilia Maria Paradi-Guilford, Toks Fayomi, Saori Imaizumi, and Panos Ipeirotis. 2015. *The global opportunity in online outsourcing*. Technical Report ACS14228. The World Bank. 1–74 pages. <http://documents.worldbank.org/curated/en/138371468000900555/The-global-opportunity-in-online-outsourcing>
- [54] Kristine M. Kuhn. 2016. The Rise of the Gig Economy and Implications for Understanding Work and Workers. *Industrial and Organizational Psychology* 9, 1 (2016), 157–162. <https://doi.org/10.1016/j.bushor.2012.09.002>
- [55] Haewoon Kwak and Jisun An. 2016. Revealing the Hidden Patterns of News Photos: Analysis of Millions of News Photos Using GDELT and Deep Learning-based Vision APIs. In *First Workshop on NEws and public Opinion (NECO ’16)*. <http://arxiv.org/abs/1603.04531> arXiv: 1603.04531.
- [56] Michael K. Lettau. 1997. Compensation in part-time jobs versus full-time jobs What if the job is the same? *Economics Letters* 56, 1 (Sept. 1997), 101–106. [https://doi.org/10.1016/S0165-1765\(97\)00098-0](https://doi.org/10.1016/S0165-1765(97)00098-0)
- [57] Bin Lin and Alexander Serebrenik. 2016. Recognizing Gender of Stack Overflow Users. In *Proceedings of the 13th International Conference on Mining Software Repositories (MSR ’16)*. ACM, New York, NY, USA, 425–429. <https://doi.org/10.1145/2901739.2901777>
- [58] David Martin, Benjamin V. Hanrahan, Jacki O’Neill, and Neha Gupta. 2014. Being a Turker. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW ’14)*. ACM, New York, NY, USA, 224–235. <https://doi.org/10.1145/2531602.2531663>
- [59] Amanda Menking and Ingrid Erickson. 2015. The Heart Work of Wikipedia: Gendered, Emotional Labor in the World’s Largest Online Encyclopedia. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI ’15)*. ACM, New York, NY, USA, 207–210. <https://doi.org/10.1145/2702123.2702514>
- [60] Claudia Olivetti and Barbara Petrongolo. 2005. *Unequal pay or unequal employment? a cross-country analysis of gender gaps*. Technical Report. Centre for Economic Performance, London School of Economics and Political Science, London. <http://cep.lse.ac.uk/pubs/download/dp0711.pdf>
- [61] June O’Neill and David M. O’Neill. 2012. *The Declining Importance of Race and Gender in the Labor Market: The Role of Employment Discrimination Policies*. Rowman & Littlefield. Google-Books-ID: ilt0Hi_hrcIC.
- [62] Jennifer A. Rode. 2011. A theoretical agenda for feminist HCI. *Interacting with Computers* 23, 5 (Sept. 2011), 393–400. <https://doi.org/10.1016/j.intcom.2011.04.005>
- [63] Todd Rogers, Noah J. Goldstein, and Craig R. Fox. 2018. Social Mobilization. *Annual Review of Psychology* 69, 1 (2018), 357–381. <https://doi.org/10.1146/annurev-psych-122414-033718>
- [64] Jakob Rogstadius, Vassilis Kostakos, Aniket Kittur, Boris Smus, Jim Laredo, and Maja Vukovic. 2011. An Assessment of Intrinsic and Extrinsic Motivation on Task Performance in Crowdsourcing Markets. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*. 321–328.
- [65] Paul R. Rosenbaum and Donald B. Rubin. 1984. Reducing Bias in Observational Studies Using Subclassification on the Propensity Score. *J. Amer. Statist. Assoc.* 79, 387 (1984), 516–524. <https://doi.org/10.2307/2288398>

- [66] Joel Ross, Lilly Irani, M Six Silberman, Andrew Zaldivar, and Bill Tomlinson. 2010. Who are the crowdworkers?: shifting demographics in mechanical turk. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 2863–2872. <https://doi.org/10.1145/1753846.1753873>
- [67] Daniel E. Russ, Kwan-Yuet Ho, Joanne S. Colt, Karla R. Armenti, Dalsu Baris, Wong-Ho Chow, Faith Davis, Alison Johnson, Mark P. Purdue, Margaret R. Karagas, Kendra Schwartz, Molly Schwenn, Debra T. Silverman, Calvin A. Johnson, and Melissa C. Friesen. 2016. Computer-based coding of free-text job descriptions to efficiently identify occupations in epidemiological studies. *Occupational and environmental medicine* 73, 6 (June 2016), 417–424. <https://doi.org/10.1136/oemed-2015-103152>
- [68] Stok F. Marijn, Ridder Denise T. D., Vet Emely, and Wit John B. F. 2013. Don't tell me what I should do, but what others do: The influence of descriptive and injunctive peer norms on fruit consumption in adolescents. *British Journal of Health Psychology* 19, 1 (Feb. 2013), 52–64. <https://doi.org/10.1111/bjhp.12030>
- [69] Elizabeth A. Stuart. 2010. Matching methods for causal inference: A review and a look forward. *Statistical science : a review journal of the Institute of Mathematical Statistics* 25, 1 (Feb. 2010), 1–21. <https://doi.org/10.1214/09-STS313>
- [70] Ryo Suzuki, Niloufar Salehi, Michelle S Lam, Juan C Marroquin, and Michael S Bernstein. 2016. Atelier: Repurposing Expert Crowdsourcing Tasks as Micro-internships. In *Proceedings of the 2016 Conference on Human Factors in Computing*. 2645–2656. <https://doi.org/10.1145/2858036.2858121>
- [71] Jacob Thebault-Spieker, Loren Terveen, and Brent Hecht. 2017. Toward a Geographic Understanding of the Sharing Economy: Systemic Biases in UberX and TaskRabbit. *ACM Trans. Comput.-Hum. Interact.* 24, 3 (April 2017), 21:1–21:40. <https://doi.org/10.1145/3058499>
- [72] U.S. Bureau of Labor Statistics. 2017. *Highlights of women's earnings in 2016*. Technical Report 1069. U.S. Bureau of Labor Statistics. <https://www.bls.gov/opub/reports/womens-earnings/2016/pdf/home.pdf>
- [73] Melissa A Valentine, Daniela Retelny, Alexandra To, Negar Rahmati, Tulsee Doshi, and Michael S Bernstein. 2017. Flash Organizations. In *the 2017 CHI Conference*. ACM Press, New York, New York, USA, 3523–3537. <https://doi.org/10.1145/3025453.3025811>
- [74] Bogdan Vasilescu, Andrea Capiluppi, and Alexander Serebrenik. 2014. Gender, Representation and Online Participation: A Quantitative Study. *Interacting with Computers* 26, 5 (Sept. 2014), 488–511. <https://doi.org/10.1093/iwc/iwt047>
- [75] Nicholas Vincent, Johnson, Isaac, and Hecht, Brent. 2018. Examining Wikipedia With a Broader Lens: Quantifying the Value of Wikipedia's Relationships with Other Large-Scale Online Communities. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. 1–13. <https://doi.org/10.1145/3173574.3174140>
- [76] Johannes Wachs, Aniko Hannak, Andras Voros, and Balint Daroczy. 2017. Why Do Men Get More Attention? Exploring Factors Behind Success in an Online Design Community. In *arXiv:1705.02972 [cs]*. <http://arxiv.org/abs/1705.02972> arXiv: 1705.02972.
- [77] Anne Weibert, Andrea Marshall, Konstantin Aal, Kai Schubert, and Jennifer Rode. 2014. Sewing Interest in E-textiles: Analyzing Making from a Gendered Perspective. In *Proceedings of the 2014 Conference on Designing Interactive Systems (DIS '14)*. ACM, New York, NY, USA, 15–24. <https://doi.org/10.1145/2598510.2600886>
- [78] Wendy Liu and Derek Ruths. 2013. What's in a Name? Using First Names as Features for Gender Inference in Twitter.. In *AAAI Spring Symposium: Analyzing Microtext*.
- [79] Matthew Wiswall and Basit Zafar. 2018. Preference for the Workplace, Investment in Human Capital, and Gender*. *The Quarterly Journal of Economics* 133, 1 (Feb. 2018), 457–507. <https://doi.org/10.1093/qje/qjx035>
- [80] World Economic Forum. 2017. *The global gender gap report: 2017*. Technical Report. World Economic Forum, Geneva. OCLC: 1012177619.
- [81] Ming Yin, Siddharth Suri, and Mary L. Gray. 2018. Running Out of Time: The Impact and Value of Flexibility in On-Demand Crowdwork. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 2018)*. ACM Press, 1–11. <https://doi.org/10.1145/3173574.3174004>

Received April 2018; revised July 2018; accepted September 2018