Price Floors and Employer Preferences: Evidence from a Minimum Wage Experiment

John J. Horton
Leonard N. Stern School of Business
New York University*

July 17, 2018

Abstract

Minimum hourly wages were randomly imposed on firms posting job openings in an online labor market. A higher minimum wage raised the wages of hired workers substantially. However, there was some reduction in hiring and large reductions in hours-worked. Treated firms hired more productive workers, which can explain, in part, the reduction in hours-worked: with more productive workers, projects were completed in less time. At the conclusion of the experiment, the platform imposed a market-wide minimum wage. A difference-in-differences analysis shows that, in equilibrium, firms still substitute towards more productive workers, adversely affecting less productive workers.

*Author contact information is available at http://john-joseph-horton.com/. Thanks to seminar participants at the University of Minnesota; the University of Chicago; FAÜ Nürnberg; Microsoft Research New York; Texas A&M University, Cornell University, Arizona State University and Harvard Business School. Thanks to Andrey Fradkin, Arin Dube, Bo Cowgill, Bill Kerr, Chris Stanton, David Autor, Dominic Coey Isaac Sorkin, Jeffey Clemens, Joe Golden, Jonathan Meer, Mandy Pallais, Michael Luca, Ramesh Johari, Richard Zeckhauser, and Scott Stern for helpful comments and suggestions. Slides: http://john-joseph-horton.com/papers/minimum_wage.slides.html Video Presentation: http://john-joseph-horton.com/papers/minimum_wage.video.html

1 Introduction

There is little consensus among economists about the effects of minimum wages. To the extent a consensus exists, it seems to be that small increases in the minimum wage do not reduce employment very much. If disemployment effects are modest, a likely reason is that firms can adjust in ways that do not necessarily reduce employee headcounts: firms can cut back on hours, reduce non-monetary compensation, increase prices, and so on. Other non-wage labor costs might fall, such as through reduced turn-over or increased productivity from enhanced worker morale that increased productivity. Given enough time, firms might also change what "kinds" of workers they hire through labor-labor substitution, perhaps leaving headcounts unchanged.

The adjustments firms might make are clear enough, but credibly measuring some of these adjustments is challenging. Although US state-level variation in minimum wages provides a plausible identification strategy, it is not a perfect strategy. Changes to state minimum wage levels are not made at random, and so much of the back-and-forth in the literature is about how to address the resulting selection issues. An additional empirical problem is measurement—some plausible firm adjustments would simply not show up in conventional administrative or survey datasets, either because the needed data is not recorded, or it is recorded with too much error to be useful.

In this paper, I report the results of a minimum wage experiment conducted in an online labor market. During the experiment, treated firms were prohibited from hiring a worker at a wage below that firm's randomly assigned minimum.[1] Job applicants were automatically instructed to raise their wage bids—if needed—when submitting applications. The existence of this minimum wage was not announced to firms or to workers. At the end of the experiment, the platform announced its intention to impose a platform-wide minimum wage, and then imposed that minimum wage several months later.

Because of the empirical context of the experiment and the exogenous

[1] I use the terms "worker", "firm", "employer", "hired", "wage" and "employer" for consistency with the literature and not as an indication of my views on the legal status of the relationships created in the marketplace.

source of variation, many of the challenges of conventional minimum wage research are not challenges in this study. With individual employers as the unit of randomization, the experimental sample is enormous, consisting of nearly 160,000 job openings. For each job opening, I observe whether anyone was hired, at what wage, and for how many hours; I also have detailed measures on the pre-experiment attributes of all workers. These measurements are made essentially without error because of the computer-mediated nature of the empirical context.

Despite many advantages, the empirical context also creates new challenges. For one, a minimum wage that only applies to some firms is quite different from a minimum wage that binds market-wide, as the latter scenario would have clear equilibrium effects not relevant in the former scenario. For these equilibrium questions, the platform's announcement and imposition of the minimum wage serves as a useful natural experiment, which I analyze after presenting the experimental results.

The main results of the experiment are as follows. Imposing a minimum wage raised the wages of hired workers, but this imposition also reduced hiring, albeit not by very much. In contrast, hours-worked fell sharply, with reductions as large as 30% in some sub-populations of job openings expected to pay low wages. Large reductions in hours-worked occurred even in sub-populations that saw no reduction in hiring. Presumably some of the reduction in hours-worked was caused by employers economizing on labor, and perhaps from improved worker morale. However, hours-worked also likely fell because treated employers hired substantially more productive workers, with productivity measured by pre-experiment worker attributes.

Employers facing minimum wages hired workers with greater past earnings, higher profile rates, and higher past average wages—all proxies for worker productivity. This labor-labor substitution towards more productive workers occurred even at minimum wages for which there was no detectable decline in hiring, ruling out a pure selection explanation for reductions (i.e., jobs that went unfilled would have only been for "small" projects taking few hours to complete). The extent of labor-labor substitution is large enough to explain about half of the reduction in hours-worked.

The labor-labor substitution I find is only detectable because of the proxies for individual productivity available in this empirical context—proxies that would not be available in conventional settings. Mirroring the conventional minimum wage literature results, I find only small changes in the composition of hired workers with respect to demographic characteristics, and these compositional changes were found only at the highest minimum wage, suggesting variation in individual productivity is mostly within demographic groups rather than between groups.

In the experiment, the minimum wage only applied to treated firms. With a market-wide minimum wage policy, if all firms tried to hire more productive workers, sought-after workers would see their wages bid up, with the amount depending on the labor supply elasticity of more productive workers. To explore these equilibrium issues, I use the platform-wide announcement and imposition of a minimum wage. Simply announcing the upcoming minimum wage apparently did little; there is little evidence that employers tried to hire workers quickly for relatively low-wage jobs or post more such jobs. In contrast, the imposition of the minimum wage had strong effects on several market outcomes.

Following the imposition of the minimum wage, the wage of hired workers increased, employers shifted towards hiring more productive workers, and hours-worked fell substantially. One difference from the experiment is that I find no evidence of a reduction in hiring in equilibrium—if anything, the probability that a job opening is filled *increases*. However, I present some suggestive evidence that employers were less likely to post job openings likely to pay low wages post-imposition, suggesting the increase in hiring could be a selection effect.

A shift in employer preferences towards relatively more productive workers could adversely affect less productive workers. Indeed, I find that workers that had been working for less than the new platform minimum wage raised their wage bids after the platform-wide minimum wage was imposed. These same workers experienced a substantial decrease in their probability of being hired.

I find no evidence that the the minimum wage had any spill-over effects on workers previously working just above the minimum wage. Despite the fall-off in hiring probability for those workers, I find no evidence that workers affected by the minimum wage were more likely to exit the market or change their job application intensity.

The most important finding of the paper is the extent of labor-labor substitution as a form of adjustment. While generalization to conventional settings should be done with the caution, the finding offers a parsimonious explanation for why conventional minimum wage studies find such modest or non-existent disemployment effects, despite the implausibility of a highly inelastic labor demand curve or a monopsonistic conception of the labor market.[2] Substantial substitution could occur in conventional markets, leaving headcounts unchanged, but changing the kinds of workers employed. This kind of substitution might be overlooked simply because it is difficult to detect without very rich individual productivity data.

The plan of the paper is as follows. Section 2 describes the empirical context of the experiment. Section 3 introduces the experimental design and explores threats to internal validity; it also focuses on the methodology for identifying job openings likely to pay low wages and thus be affected by the active treatments. Section 4 presents a simple conceptual framework for understanding the experimental results. Section 5 presents the main experimental results of the paper. Section 6 presents results from the announcement and imposition of a market-wide minimum wage. Both the experimental and non-experimental results are discussed in Section 7 and some concluding thoughts are offered.

[2] In some search-focused models of the labor market, a minimum wage might lead to more filled vacancies, raising the value to firms of posting such vacancies, as workers reject fewer offers (Burdett and Mortensen, 1998) or search more intensively (Flinn, 2006).

2 Empirical context

In online labor markets, firms contract with workers to perform tasks that can be done remotely, such as computer programming, graphic design, data entry, and writing (Horton, 2010). Platforms differ in their scope and focus, but common services provided by the platforms include publishing job listings, hosting user profile pages, arbitrating disputes, certifying worker skills, and maintaining reputation systems.

The experiment in this paper was conducted in a large online labor market. In this market, a would-be employer writes job descriptions, labels the job opening with a category (e.g., "Administrative Support"), lists required skills, and then posts the job opening to the platform website. Workers generally learn about job openings via electronic searches. Workers submit applications, which generally include a wage bid (for hourly jobs) or a total project bid (for fixed-price jobs) and a cover letter. In addition to worker-initiated applications, employers can also search worker profiles and invite workers to apply. After a worker submits an application, the employer screens his or her applicants and can decide to make an offer or offers.[3]

The marketplace used for this study is not the only marketplace for online work. As such, a worry is that job openings are simultaneously posted on several platforms, and perhaps in conventional markets as well. However, surveys conducted by the platform suggest that online and offline hiring are only very weak substitutes, and that "multi-homing" of job openings is relatively rare. Supporting this view, a finding of the experiment is that hiring reductions were small or non-existent, implying that displacement to other platforms was not an important margin of adjustment, at least in the short-run. Furthermore, as I will discuss later, there is no evidence that treated job openings were subsequently posted on another online labor market that, at the time the experiment was run, had a lower minimum wage.

[3]Although they can bargain over the wage, there is relatively little wage bargaining, with most employers and workers treating wage bids as take-it-or-leave it offers (Barach and Horton, 2017a). Interestingly, Fradkin (2016) finds surprisingly little bargaining on Airbnb. There is perhaps some reluctance to begin a relationship with haggling over price.

There has been some research that uses online labor markets as an empirical context for research. Pallais (2014) shows via a field experiment that past on-platform worker experience is an excellent predictor of being hired for future job openings. Stanton and Thomas (2016) show that agencies (which act as quasi-firms) help workers find jobs and break into the marketplace. Agrawal et al. (2016) investigate what factors matter to firms in making selections from an applicant pool and present some evidence of statistical discrimination, which can be ameliorated by better information. Horton (2017b) explores the effects of making algorithmic recommendations to would-be employers. Barach and Horton (2017b) report the results of an experiment in which employers no longer had access to applicant wage history when making hiring decisions.

2.1 Variable construction and measurement

Measures of individual worker productivity are important for studying the effects of the experiment. Some of these measures are straightforward—such as past wages earned on the platform—but others require some explanation. One important productivity measure is a worker's hourly "profile rate," which is listed on his or her platform profile. This profile rate is a worker's default bid for hourly job openings, in that the application form is pre-populated with it, though workers are free to tailor their bids to each job opening. Workers can set their profile rate and change it whenever they like, but they have an incentive to keep it "close" to what they think their market rate is, as firms searching in the market for workers use the profile rate in their decision-making in deciding who to invite (Horton, 2017a). If a worker is hired, it is at an agreed-upon hourly wage.

Most relationships formed on the platform are quite short (the median contract is on the order of a week). However, these relationships have no set end date and some relationships could continue, causing measures like the count of hours-worked to grow. To work on hourly contracts, workers must install software that precisely records hours-worked, To stabilize the data for analysis purposes, I stop measurements at 6 months after the formation of the contract; only about 3.8% of filled contracts extend beyond this cut-off.[4]

3 Experimental design and internal validity

During the experimental period, firms posting an hourly job opening were immediately assigned to an experimental cell.[5] The experiment consisted of four experimental cells: a control group with the platform status quo of no minimum wage, which received 75% of the sample ($n = 121,704$), and three active treatment cells, which split the remaining 25% of the sample. A total of 159,656 job openings were assigned. Neither employers nor workers were told they were in an experiment. The active treatments had minimum wages of \$2/hour in MW2 ($n = 12,442$), \$3/hour in MW3 ($n = 12,705$), and \$4/hour in MW4 ($n = 12,805$). If the firm posted additional job openings, these openings also received the same experimental assignment as the original opening. However, I do not include these follow-on openings in the analysis.

The minimum wage was implemented by not allowing workers to submit wage bids below the assigned opening-specific minimum wage. Prior to the experiment, wage bids were restricted to positive numbers via an automated check of the job application form. For job openings in the active treatment cells, this \$0 floor was simply raised to the appropriate minimum. If an applying worker tried entering a wage below the minimum wage, he or she was instructed (via a dialog box) that the proposed wage was too low and needed to be raised. The worker was not told the precise amount the wage bid had to increase by, in order to reduce bunching at the exact cutoff.[6] The worker's

[4]No results are sensitive to this restriction. The analysis is available upon request.

[5]Firms posting fixed-price jobs were not eligible for the experiment. Firms could have posted a subsequent fixed-price job to avoid the minimum wage, which is part of the reason I only use the first job opening in the analysis. Despite this possibility, there is no evidence that firms switched to using fixed-price contracts as an adjustment strategy. This analysis is not reported here, but it is available upon request. A small number of very large platform employers were exempted pre-randomization.

[6]Given this bunching-prevention design choice, along with the fact that applications are shown to employers 10 at a time and not in wage bid order, with no visualization made of the distribution of wage bids, it is unlikely that employers would infer they were in some kind of experiment.

application was not sent to the employer until the minimum wage condition was met. The experimental intervention was effective at preventing contracts from being formed below the cell-specific minimum wage (see Appendix A.1),

3.1 Threats to internal validity

The internal validity of the experiment would be compromised if any of the following occurred: (1) a failed randomization, (2) applicant attrition at the wage bidding stage, (3) workers sorting across job openings based on the experimental cell of that opening, (4) firms sorting across time (i.e., posting the same opening again some time later to get a better "draw" of applicants), and (5) firms sorting across platforms, including the "platform" of the conventional labor market. For issues (4) and (5), the concern is that any observed reduction in hiring could actually be displacement to other platforms. However, as will be discussed, there was very little reduction in hiring, so both of these concerns are somewhat moot. For issues (2) and (3), the concern is that different cells would have selected applicant pools. For some of these issues, the relatively small active treatment cells were a useful design feature, as it reduced the potential for market-moving violations of the SUTVA condition—a common concern in experiments conducted in a true marketplace (Blake and Coey, 2014).

For issue (1)—failed randomization—there is no evidence this occurred. The software used to randomize openings has been used for numerous prior experiments on the platform without issue. Job openings are well-balanced on pre-randomization attributes, and the counts of job openings per cell is consistent with randomization. See Appendix A.2 for this analysis.

For issue (2)—applicants abandoning the application process before submitting an application—this was regarded, ex ante, as unlikely. Submitting a wage bid was the last part of the application process. For this reason, workers had already borne the costs of application, making those costs sunk, and so workers had little incentive not to comply with the instructions to raise their wage bid.[7] Consistent with this sunk cost argument, there is no evidence that the count of applicants differed across experimental cells. See Appendix A.3 for this analysis.

For issue (3)—worker sorting across openings—the potential problem is that workers were free to apply to any job opening in the marketplace, and if workers knew the assignment of a job opening, they might seek out their preferred opening. While a concern in principle, this kind of sorting would be exceedingly difficult in practice, and the lack of differences in applicant counts by experimental group is consistent with a lack of sorting.[8]

The reason this worker sorting is unlikely is that firm treatment assignments were not publicly known to workers, nor was the existence of the experiment. As such, few workers learned there was some opening-specific difference to seek out—much less what preferences they should have over these differences. The only way to learn about any particular job opening's assignment was to apply at a low enough rate. Compounding the difficulty for would-be sorting workers, recall that only 25% of job openings had any minimum wage at all, which would make finding a preferred opening challenging. In addition to high costs, there would not be much incentive to seek out these openings, as few workers face a binding constraint on the number of applications sent. Workers have little incentive to save applications for the "best" job opening,

For issue (4)—firms "sorting" across time by re-posting their job opening to get another draw of applicants—the problem is that such sorting would look like a reduction in hiring, as the first job goes unfilled. Firms might post another job opening if they thought they received an idiosyncratically bad "draw" of applicants.[9] Despite the possibility, there is actually a slight

[7] Application costs include finding a job opening, deciding to apply, writing a cover letter, and so on. Although workers do have a time-based quota of job applications they can send, it is set so high that it is almost never binding and so withdrawing an application because of a too-high minimum wage would be unlikely.

[8] However, the lack of a difference in counts is not decisive, as different workers could have different preferences that left total counts unchanged.

[9] Or they could re-post to avoid their treatment cell if they (a) believed they were in an experiment and (b) mistakenly thought the level of randomization was the job post rather than firm or (c) thought the experiment would conclude shortly.

decrease in the probability of posting a subsequent job opening for treated employers. See Appendix A.4 for this analysis.

For issue (5)—firms sorting across platforms—the concern is that it would look like a reduction in hiring. To assess this concern, I checked whether firms in the highest minimum wage cell were more likely to post their job openings on another online labor market that, at the time, had a lower minimum wage. I find no evidence of increased cross-posting of job openings assigned to the highest minimum wage. See Appendix A.5 for this analysis.

I have no evidence on whether any work was displaced to offline hiring, but as discussed in Section 2, survey evidence suggests that few firms see offline hiring as a substitute for online hiring. Given the nature of work and the typical wages on online labor platforms, it is unlikely that local hiring was a feasible alternative for most firms. To re-iterate, there was little or no reduction in on-platform hiring, consistent no displacement to "offline" hiring.

Given the lack of internal validity issues, there is a simple way to interpret the experiment: firms got the same applicants they would have gotten, regardless of experimental cell, but with the distribution of wage bids differing based on their treatment assignment—namely with workers that would have submitted non-complying wage bids bidding up.

3.2 Low wage sub-populations of job openings

I use two approaches to find sub-populations of job openings that were likely to pay low wages: (1) I use the lowest paying sub-categories of work, which are largely found in the "Administrative Support" category, or ADMIN,[10] (2) I fit a predictive model with data from historical job openings. I then use the fitted model out-of-sample to label all experimental job openings with low predicted wages ($\leq$ \$5/hour) as the low-predicted wage sample, or LPW.

[10] The sub-categories used are: "Web Research", "Sales & Lead Generation", "Data Entry", "Other - Business Services", "Other - Sales & Marketing", "SEO - Search Engine Optimization", "Personal Assistant", "Other - Administrative Support", "Advertising", "SMM - Social Media Marketing", "Customer Service & Support", "Order Processing", "Market Research & Surveys", "Other - Customer Service", "Technical Support", "Email Marketing", "Email Response Handling", and "Payment Processing."

The left panel of Figure 1 shows the kernel density estimate of the log wages of hired workers in the control. This distribution is decomposed into ADMIN and non-ADMIN job openings in the right panel of the figure.[11] The three levels of the minimum wage are overlaid as dashed vertical lines. We can see that the wage distribution in ADMIN is considerably left-shifted relative to the non-ADMIN job openings. In ADMIN, the 1st quartile of the wage distribution is below \$3/hour, the median is near \$4/hour, and the 3rd quartile is only slightly above \$5/hour. Note that the highest minimum wage of \$4/hour is above the median wage in ADMIN.

Figure 1: Wage distributions of hired workers in the control group

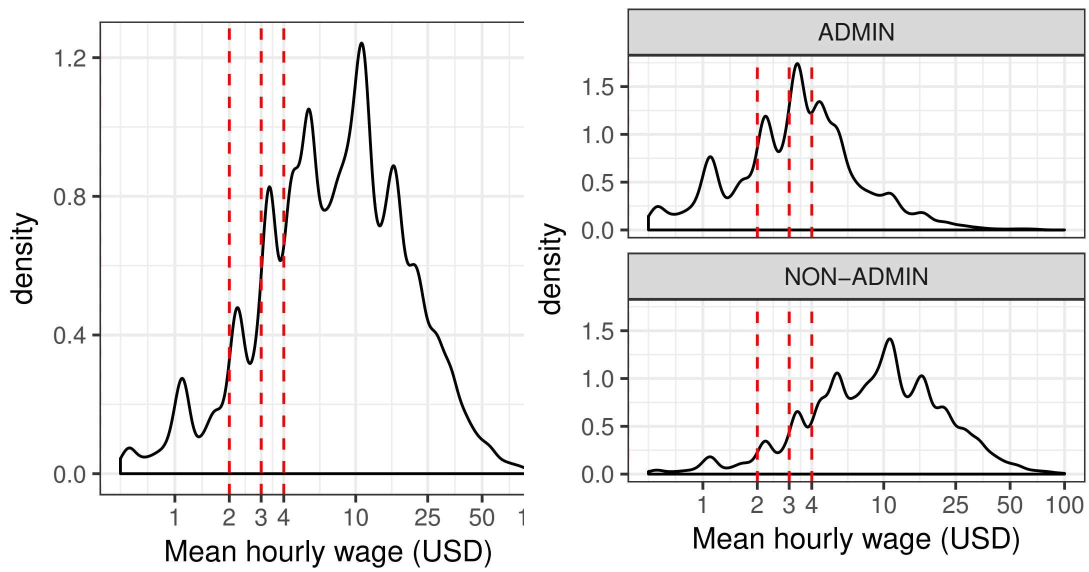

(a) Control cell of ALL

(b) Control cell of ADMIN jobs (top) and non-ADMIN openings (bottom)

Notes: This figure shows the distribution of the hourly wages of hired workers in the control group, on a log scale. The kernel density estimate shown in the left panel is for all workers. In the right panel, the top density estimate is workers hired to ADMIN job openings in the control. The bottom density estimate is for all other job openings in the control group.

For the LPW sample predictive model, the training data was 100,000 pre-experiment job openings for which a hire was made. The outcome was log

[11]The bandwidth for the kernel density estimate is selected using Silverman's rule of thumb (Silverman, 1986).

hourly wage for the hired worker. The candidate predictors included the category of work, skills required, the anticipated duration, and the job opening title.[12] To estimate the model, I used the glmnet package developed by Friedman et al. (2009), using LASSO for regularization and variable selection (Tibshirani, 1996), with the optimal tuning parameters selected via cross validation. Using the fitted model, I made predictions for every job opening in the experiment, and then selected those predicted to pay less than $5/hour.

4 Conceptual framework

With the experiment described, I now consider the connections between the experimental design and other empirical and theoretical work. One prediction common to all competitive labor market models is that fewer hours of labor are demanded when the minimum wage is binding. In conventional minimum wage research, this prediction is typically tested at the "market" level, with quantities measured not with the actual number of hours worked, but instead with the headcounts of employed workers. Whether this measure changes following a change to the minimum wage is the focus of much of the controversy in the modern minimum wage literature.

The first wave of quasi-experimental evidence showed little or no short-run disemployment effects for small increases in the minimum wage (Card and Krueger, 1994; Card, 1992; Katz and Krueger, 1992), but this revision never reached a consensus, as other work using more or less the same methods did show dis-employment effects (Neumark and Wascher (1992); Neumark et al. (2004)).[13] At present, the debate is both active and unsettled, with new debates about what is the proper way to account for state-specific differences in growth (Allegretto et al., 2011; Neumark et al., 2014), and whether using other control methods, such as contiguous counties, is more attractive (Dube

[12]For textual predictors, I used the RTextTools package, developed by Jurka et al. (2012) to create a document term matrix.

[13]Sorkin (2015) argues that most of the empirical literature in the US context has focused on short-run effects and that if adjustment costs are high, short-run estimates would be unable to detect long-run effects.

et al., 2010).

Newer empirical approaches are characterized by alternative ways of defining the populations of interest but still rely on state variation in US minimum wages. For example, Meer and West (2015) look at the flow of new openings rather than the stock of all existing relationships. They find a substantial reduction in job growth caused by higher minimum wage levels. Both Powell (2016) and Dube and Zipperer (2016) adopt a synthetic control approach, using contiguous counties as comparison units, though the papers reach different conclusions about dis-employment effects. Another paper using the contiguous counties approach is Aaronson et al. (Forthcoming), which looks at how changes in the minimum wage changed the composition of restaurants, finding evidence consistent with the putty-clay model of firm dynamics.

There are also attempts to break out of state/county panel framework. Clemens and Wither (2014) look at the career trajectories of workers right below and right above newly imposed minimum wages, finding that workers on the "wrong" side of the new minimum suffered substantial reductions in earnings and employment probabilities.

There is some research on non-employment adjustments to minimum wages. Schmitt (2013) provides an overview of the various hypothesized adjustment margins. One paper in this vein is Draca et al. (2011), which finds a reduction in firm profits following the UK minimum wage implementation. Another is Hirsch et al. (2011), which studies the effects of an increased minimum wage on fast-food restaurants. They find that the minimum wage leads to higher prices for customers and lower profit margins, but that employers are partially compensated with reduced turn-over. Consistent with this finding, Luca and Luca (2017) finds that minimum wages tend to drive lower quality (and presumably lower profit) restaurants out of business. There is also work on labor-labor substitution as a margin of adjustment, which I will discuss later, after presenting evidence from the experiment on labor-labor substitution.

4.1 Hourly hiring for project-based work

In the empirical context of the experiment, bids are for hourly work, but the work itself is still project-based, which requires some re-framing of the employer's problem. Consider a firm with a project of "size" Y, meaning the project can be completed with Y efficiency units of labor. A worker with technical productivity y will complete the project in Y/y hours. When completed, the firm will sell the output for pY. The firm receives wage bids from a pool of I applicants with heterogeneous technical productivity. Let a worker i have technical productivity y_i. Each worker submits a take-it-or-leave-it hourly wage bid of w_i. If worker i is hired, the project is completed in Y/y_i hours and the wage bill is $w_i Y/y_i$.

The firm has an outside option of $\underline{u}$ if it chooses not to hire anyone. The firm's decision problem is to hire the profit-maximizing applicant, if any, or

$$\arg\max_{i \in I} \quad \left\{ \underline{u}, pY - w_i \frac{Y}{y_i} \right\}. \tag{1}$$

As Y and p are the same for each applicant, the firm selects the applicant that minimizes the ratio of wages to technical productivity, so long as the payoff obtained from hiring that applicant exceeds $\underline{u}$. Let $G(\cdot)$ be the cdf of $\pi_i = pY - w_i Y/y_i$.

Proposition 1 predicts a hiring effect from minimum wages, while Proposition 2 predicts an hours-worked and wage effect (proofs are in Appendix B). The propositions say nothing about the relative magnitudes of these two effects, but they provide a framework for interpreting the experimental results.

Proposition 1. *Under a firm-specific minimum wage, the firm is less likely to hire anyone.*

Proposition 2. *If a firm facing a minimum wage, $\underline{w}$, still makes a hire, the expected number of hours-worked falls and the observed wage of the hired worker increases.*

Proposition 2 is essentially about the substitution that can happen within an applicant pool. Given the nature of labor, variation in y is expected, but if the market is competitive, why is w/y not the same for all applicants?[14] One explanation is that workers regard job openings as being more or less attractive than their other options at that moment in time, and these differences are reflected in their wage bids.[15] Similarly, firms might infer different levels of productivity in the applicants—differences that the applicants themselves are unaware of and do not incorporate into their bids. Supporting this view, there is evidence of substantial heterogeneity in productivity among workers receiving the same hourly wage (Lazear et al., 2015). Whatever the source of idiosyncratic variation in w/y, the result is a distribution of payoffs the firm would get from hiring different workers, which creates the possibility of substitution when a price floor is imposed.

5 Experimental results

The main experimental outcomes of interest are: whether the firm hired anyone; the number of hours-worked for hired workers; the wages of hired workers; and the pre-randomization attributes of hired workers, particularly those attributes that are proxies for productivity. The effects on earnings are reported in Appendix C.1, as they are already implied by the changes in wages and hours-worked.

For each outcome, I analyze the experiment in two ways: (1) with treatment cell indicators as regressors and the control group as the omitted category, and (2) with the numerical minimum wage as a regressor. For each outcome, I present results for all job openings, labeled ALL, for administrative openings, labeled ADMIN, and for jobs predicted to pay low wages, or LPW.

Some outcomes, such as the wage of the hired worker, are only observed if a hire is made. Other outcomes are well defined even if a hire is not made,

[14]This is similar to the question posed by Romer (1992) about why firms try to hire the "best" applicants if workers are simply paid their marginal product. The fact that firms bother to screen and evaluate candidates before making a hire is evidence that they are not indifferent over the pool they receive, and that there is latent information beyond what is reflected in the wage bid.

[15]See Horton (2017a) for evidence on how workers adjust their wages depending on how busy they are when they apply.

such as hours-worked (with zeros for jobs where no hire was made), though these outcomes are more naturally expressed in logs. For these cases, I restrict the sample to only those observations with non-zero quantities. I deal with the resulting interpretation issue by (1) flagging it when it is relevant, (2) highlighting results from cells that had no reduction in hiring. I also report all regression outcomes in "levels" in appendices.

Firms are free to hire multiple workers for the same job opening. However, multiple hires are fairly rare in the experimental data: of employers making a hire, 85% only hire one worker, while 9% hire two workers. There is no evidence that the minimum wage altered the number of hires per opening, conditional upon the employer making at least one hire. When there are multiple hired workers per job opening and the outcome of interest is a rate, such as the wage of hired workers, I use the average for all hired workers. For outcomes that are quantities, such as the number of hours-worked, I use the sum.

5.1 Effects on hiring

Let h_j be the number of hours-worked by a worker hired for job opening j. I define a job opening as being "filled" as any hours worked, or $\mathbb{1}\{h_j > 0\} = 1$. Consider a regression of this indicator on the treatment indicators, or

$$\mathbb{1}\{h_j > 0\} = \beta_0 + \beta_2 \text{MW2}_j + \beta_3 \text{MW3}_j + \beta_4 \text{MW4}_j + \epsilon, \tag{2}$$

where $\text{MW}x_j$ is an indicator for whether job opening j had a minimum wage of x. The left panel of Figure 2 reports the $\hat{\beta}$ coefficients for each of the three active treatment cells (i.e., MW2, MW3 and MW4), for ALL, ADMIN, and LPW. Around each point estimate, a 95% confidence interval is shown, calculated with robust standard errors. All regression results are additionally presented as tables in Appendix F.

Starting with ALL, we can see that hiring was lower across minimum wage cells compared to the control. The reduction in hiring is statistically significant in MW4 and nearly so in MW3. In the MW4 cell, which had the largest reduc-

Figure 2: Effects of the minimum wage on whether anyone was hired for the job opening

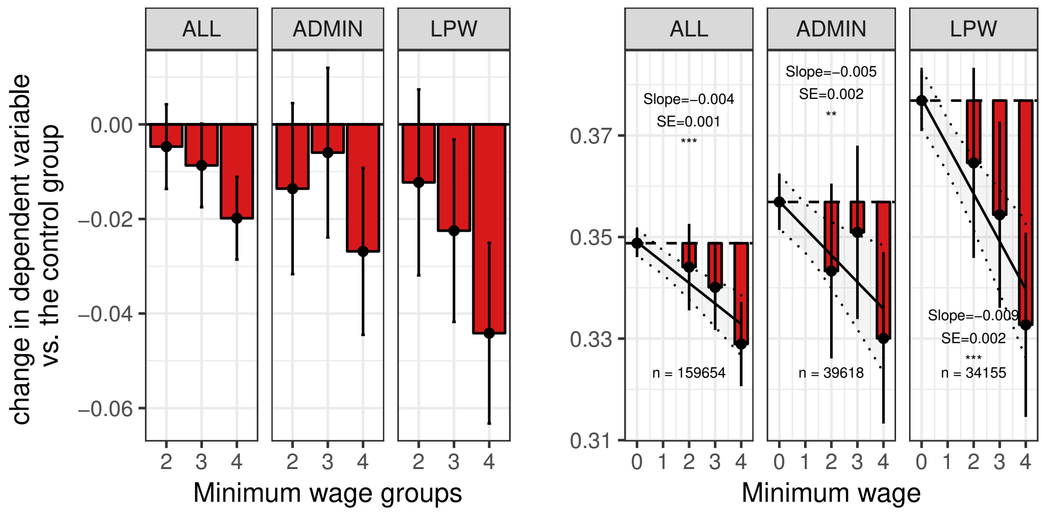

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is whether the employer hired anyone and paid them some amount of money. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$..

tion in hiring, the decrease is about 2.0 percentage points.[16] This reduction is from a baseline hiring rate of 35%, so the treatment caused approximately a 7% reduction in hiring.

Among the ADMIN and LPW sub-populations, there are larger reductions in hiring due to MW4. However, the largest reduction, in LPW, is still only about 4.0 percentage points. Because LPW job openings have a higher baseline hire rate (about 38%), the percentage reduction in hiring is only about 10%, despite the minimum wage in MW4 being substantially above the median wage for filled control cell job openings in LPW. In MW3, the reduction is close to 5%, and in MW2 it is close to 2.5%.

[16]Note that here—and throughout the paper—for differences in levels where the outcome is naturally discussed as a fraction, I label level differences as "percentage points," whereas for true percentage changes from the control, I use the "%" symbol. When the outcome is in logs, I describe changes in log points as percentage changes using the $\log(1 + x) \approx x$ approximation.

The right panel of Figure 2 reports regression results where the outcome is still $\mathbb{1}\{h_j > 0\}$, but the regressor is the imposed minimum wage, or

$$\mathbb{1}\{h_j > 0\} = \alpha_0 + \alpha \underline{w}_j + \epsilon, \tag{3}$$

where $\underline{w}_j$ is the minimum wage assigned to job opening j. The fitted regression line is plotted for each sample, with a 95% prediction interval for the conditional expectation. The associated cell-specific effects (from the left panel) are overlaid on the plot, but with the origin placed at the mean value for the outcome in the control cell.[17]

In the right panel of Figure 2, we can see that the effect of the minimum wage on hiring is negative and highly significant in ALL, as well as in the subpopulations, ADMIN and LPW. For LPW, the slope is such that each $1 increase in the minimum wage lowers the hiring probability by about 1 percentage point. The slopes in ALL and ADMIN are about half as large in magnitude as the slope for LPW. Although it is tempting to calculate a hiring elasticity with respect to the minimum wage from this data, the control-to-MW2 jump has an undefined denominator, and each subsequent difference is imprecisely estimated.

5.2 Effects on hours-worked

The left panel of Figure 3 shows the effects of the minimum wage on log hours-worked. The sample is restricted to openings where a worker was hired and he or she billed at least one quarter of an hour, the minimum amount of billable time on the platform.[18] In the full population ALL, hours-worked fell in every cell with minimum wages. The magnitude of the effect ranges from a little less

[17]Although the bars illustrating effects are now visually smaller, they can be compared relative to the baseline in the control—something not possible in the left panel, which is better suited for comparisons across the different minimum wage cells. The size of the sample for each regression is indicated in each panel, left and right. The R^2 values are omitted as they are generally very close to zero.

[18]I also use the count of hours-worked, with zero hour contracts included, as the outcome in Appendix C.2. The pattern of results is the same as the log hours-worked analysis presented here.

than a 5% reduction in MW2 to a nearly 10% reduction in MW3. The effects are conventionally statistically significant in MW3 and MW4, and nearly so in MW2.

In the sub-populations, the story changes dramatically: the reductions in hours-worked are significant in every cell, in both ADMIN and LPW. The magnitudes are substantial, with reductions of more than 25% in both MW3 and MW4 in the LPW sample. In ADMIN, the decrease in hours-worked is about 20% in both MW3 and MW4, and more than 15% in MW2. It is important to note that hours-worked fell even in cells that had little or no reduction in hiring. For example, MW3 in ADMIN had almost no reduction in hiring, but a 25% decrease in hours-worked, ruling out a pure selection explanation.

Figure 3: Effects of the minimum wage on log hours-worked, conditional upon a hire

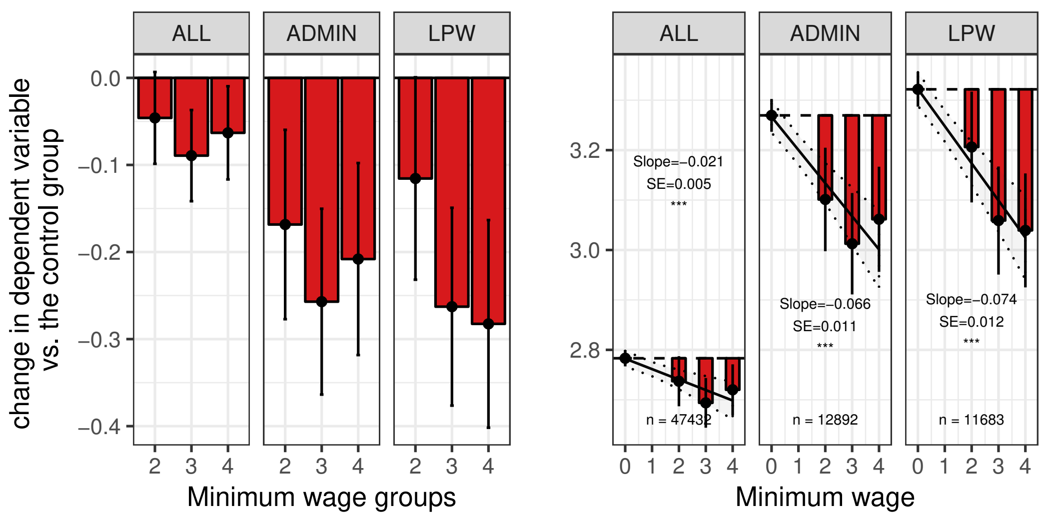

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the log hours worked, conditional upon a hire. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$..

The stronger effects on hours-worked in the sub-populations can be seen in the right panel of Figure 3. A \$1 increase in the minimum wage leads to about 7% fewer hours-worked in both ADMIN and LPW, while the slope in ALL is much flatter, with reductions of only about 2% per \$1 increase in the minimum wage.

5.3 Effects on wages of hired workers

If a worker is hired, I can observe his or her hourly wage. The left panel of Figure 4 reports results from regressions of the log wage of the hired worker on the cell indicators. In ALL, the minimum wage increased hourly wages: there is nearly a 10% increase in MW2 and MW3 and a 15% increase in MW4. For the sub-populations, the effects are stronger. In ADMIN hired wages rose nearly 40% in MW4, 25% in MW3, and 15% in MW2. The effect sizes are similar in LPW. In all samples, effect sizes are increasing in the level of the imposed minimum wage.

Figure 4: Effects of the minimum wage on log mean wage, conditional upon a hire

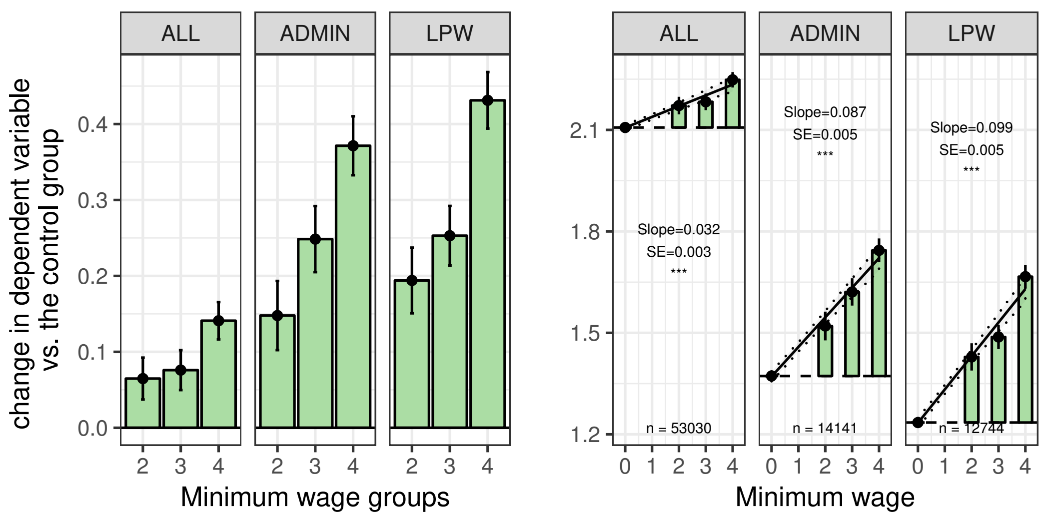

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the log mean wage paid, conditional upon a hire. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$..

In the right panel of Figure 4, the slope is highly significant in ALL, with each \$1 increase in the minimum wage associated with 4% higher wages. The slope is much steeper in the sub-populations: in ADMIN and LPW, each additional dollar in the minimum wage is associated with about 9% higher wages for the hired worker.

It is clear that imposing a minimum wage increased the wages of hired workers. There are several potential reasons for the increase: (1) the job openings that do not fill would have paid low wages—what is "left" are the relatively higher-paying jobs; (2) firms hire the same workers they would have hired anyway, but at a higher wage; (3) the firms select higher productivity workers, who command higher wages. Although these explanations are not mutually exclusive, I can rule out (1) as the sole explanation: recall that in the MW3 cell in ADMIN, there was almost no reduction in hiring, and yet the average wage increased by nearly 25%.

Although I do not know who the firm would have counter-factually hired when assigned to some other cell, I can test whether hired workers have higher or lower than expected wages by cell, given their productivity-relevant attributes. This test requires having some notion of a worker's expected wage. One attractive predictor for a worker's wage is his or her pre-experiment profile rate.[19] The difference between the profile rate and the hired wage can be thought of as a markup. If hired workers have higher average markups when their hiring firm faced a minimum wage, it suggests that some of the observed increase in wages came from workers bidding more but still being hired. I find that imposing a minimum wage strongly increased the average markup of the hired worker, with markups increasing by as much as 25 percentage points in the MW4 group in ADMIN and LPW—see Appendix C.3 for the full analysis of these markup effects.

Firms paying higher wages for the "same" workers could explain, in part, the reduction in hours-worked if firms simply economized on labor such as by reducing the scope of their projects. Another possibility is that workers exhibit greater productivity in response to the "gift" of high wages, ala Akerlof and Yellen (1990). However, as Gilchrist et al. (2016) show with a field experiment in a very similar empirical context, paying higher wages, per se, has no discernible effect on measured productivity.

[19]The profile rate is discussed in Section 2.1.

Another possibility is that firms simply hired more productive workers—a hypothesis we can explore in part by examining the productivity-relevant attributes of workers hired in treated cells. I explore these potential selection effects in the next section.

5.4 Effects on the composition of hired workers

To test for substitution towards more productive workers, I use the average past wage rate of the hired worker as the outcome, calculated using jobs completed before the start of the experiment. This past wage is arguably the most direct measure of a worker's marginal productivity.

Figure 5 reports regressions where the outcome is the log past average wage of the hired worker. The average wage is calculated by dividing total hourly earnings by total hours-worked. In ALL, in the left panel, hired workers in MW2 and MW4 had higher past wages, with the effect significant or nearly significant in both cells. The effect is slightly negative in MW3. Consistent with this mixed evidence, the slope in the ALL sample in the right panel is positive but not conventionally significant.

In the sub-populations, hired workers had substantially higher past average wages in the active treatment cells. In the MW4 cell, in both ADMIN and LPW, hired workers had approximately 15% higher past wages compared to those hired in the control. The MW2 effects were positive and close to 5%, and nearly conventionally significant. In the right panel, in both LPW and ADMIN, each $1 increase in the minimum wage is associated with about a 3% increase in the average past wage of the hired worker.

I also look at selection with respect to the profile rate and cumulative past earnings (in Appendices C.4, C.5, respectively). These other productivity proxies show the same pattern—firms hired substantially more productive workers.[20]

[20] In Appendix C.6 I examine whether the treatments affected the probability the employer hired a worker with no past on-platform experience. There is no evidence that treated workers were less likely to hire a worker without experience, though given the high baseline rate (over 90% have experience in the control group), there is not much "room" for large effects.

Figure 5: Effects of the minimum wage on the log past wage of the hired worker

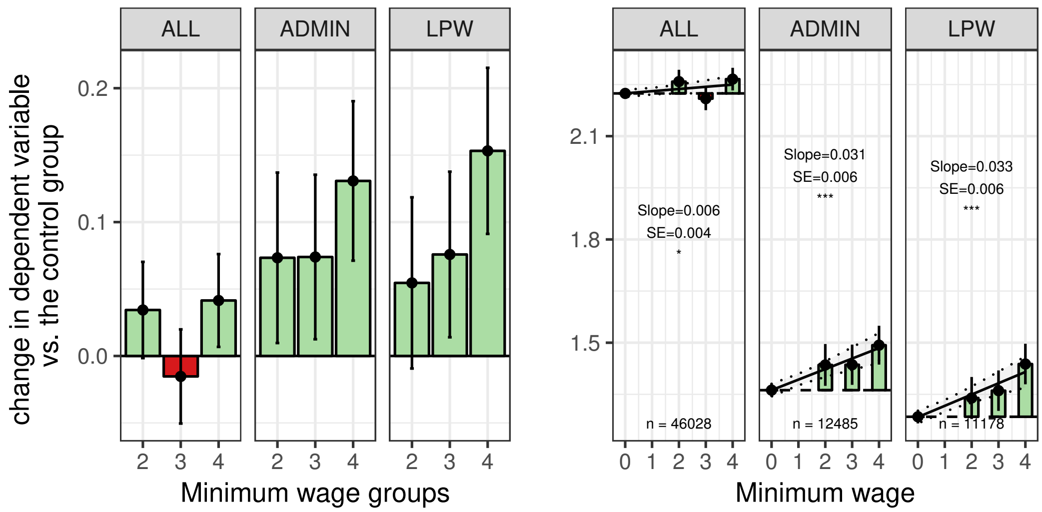

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is log past wage of the hired worker. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$..

If the productivity proxies are proportional to the marginal product of the worker, then in the model sketched out in Section 4, the extent of substitution is large enough to explain about the half of the reduction in hours-worked. Several other studies have found evidence of labor-labor substitution in conventional markets in response to minimum wages or wage floors.[21] In these studies, changes in the composition of hired workers are detected with respect to demographic characteristics. For example, using personnel data from a single large firm, Giuliano (2013) finds that teenagers from zip codes where socioeconomic status is higher displaced older workers following a minimum wage increase. Fairris and Bujanda (2008) find that a Los Angeles living wage that applied to city contractors caused those vendors to substitute in favor

[21] Although modern empirical work has given relatively little attention to labor-labor substitution, some of the earliest empirical work on the minimum wage considered the possibility: the remarkable study of the introduction of a minimum wage in Oregon by Obenauer and von der Nienburg (1915) looked at changes in employment by workers of different experience levels, which had different associated minimum wages.

of workers with characteristics associated with a wage premium in that local labor market.

Extant work on labor-labor substitution focuses on demographics, but detecting labor-labor substitution through changes in demographics is potentially challenging if most of the variation in individual productivity is within—rather than between—demographic groups. If this is the case, the kind of productivity-focused substitution found in the experiment might not result in much evidence of substitution if measured by changes in demographics. In the next section, I mirror the demographic approach to detecting substitution, using the hired worker's country, which is associated with large differences in hourly wages.

5.5 Country of the hired worker

Table 1 reports regressions where the outcomes are indicators for whether the hired worker was from a particular country. The independent variables are indicators for the treatment cell, with the control group as the omitted category. The countries are, from left to right, the US, India, Philippines, and Bangladesh, corresponding to Columns (1) through (4). Countries are ordered by the average hourly wage of workers from that country. The four countries used in this analysis made up about 80% of the hired workers, with the plurality coming the Philippines (about 30%), with India and Bangladesh next, each with about 20%, followed by the US at only 7%. The sample is the same in each regression and consists of job openings in LPW, the sub-population in which we would expect the strongest substitution effects.

Column (1) of Table 1 shows that in MW4, the fraction of hires from the US increased from about 7% to 10%. Column (4) shows that workers from Bangladesh saw their share of hires reduced by about 2.5 percentage points. Both shifts are conventionally significant. Comparing across countries, the MW2 and MW3 coefficients are not conventionally significant. Furthermore, the magnitudes are all close to zero, though both MW2 and MW3 indicators are positive for the US and negative for Bangladesh. The MW4 point estimate

Table 1: Effects of the minimum wage on the country of the hired worker

	Hired worker from:			
	US	India	Philippines	Bangladesh
	(1)	(2)	(3)	(4)
MW4	0.032***	−0.003	−0.011	−0.025*
	(0.008)	(0.012)	(0.014)	(0.012)
MW3	0.009	0.007	−0.006	−0.016
	(0.008)	(0.012)	(0.014)	(0.012)
MW2	0.003	0.001	0.009	−0.008
	(0.008)	(0.012)	(0.014)	(0.012)
Constant	0.073***	0.187***	0.303***	0.209***
	(0.002)	(0.004)	(0.004)	(0.004)
Observations	14,131	14,131	14,131	14,131
R^2	0.001	0.00003	0.0001	0.0004

Notes: This table reports regressions where the dependent variable is an indicator for whether the hired worker was from the indicated country. The countries are, from left to right, the US, India, Philippines, and Bangladesh. This is also the descending ordering of average wages on the platform by worker country. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$.

for India is close to zero, and while the Philippines estimate is negative, it is only about 1 percentage point and is not conventionally significant.

The US versus Bangladesh comparison in MW4 is suggestive of substitution, and perhaps a "conventional" analysis would have detected it. However, the magnitudes are not large, and recall that MW4 in LPW had a non-trivial reduction in hiring, and the shift could be viewed as due to selection. That substitution is barely detectable with respect to demographic measures but easily detectable with respect to individual productivity measures illustrates the importance of individual productivity measures in detecting substitution.

6 Effects of market-wide imposition

After the experiment concluded, the platform implemented a universal \$3/hour minimum wage. Unlike the experiment, this minimum wage policy was publicly announced. The announcement was made about two and half months before the minimum wage was imposed. As the minimum wage was universally applied, it is not possible to report experimental estimates of its effects. However, I can compare various market outcomes before and after the announcement and imposition. To control for any seasonal differences, I can use market data from one calendar year prior to construct difference-in-differences estimates.

The basic empirical strategy is to estimate a regression of the form

$$y_j = \beta_0 + \beta_{\text{ACTUAL}} \text{POST}_j + \epsilon, \tag{4}$$

where POST_j is an indicator that job opening j was posted after the announcement of the \$3/hour minimum wage. I then estimate the same regression with the same method of constructing the sample, but with data from one year prior. The implied difference-in-differences treatment effect is $\hat{\beta}_{\text{ACTUAL}} - \hat{\beta}_{\text{PLACEBO}}$. As there are numerous choices that can be made about the "window" size to use in both the pre- and post-periods, I simply report a range of estimates using different values.

6.1 Hiring and job opening composition

Figure 6 plots a collection of difference-in-differences estimates.[22] The left "column" shows announcement effects and the right column shows imposition effects. Each estimate uses a different pre- and post-window length. I use three different pre-period windows: four weeks, five weeks, and seven weeks. These different windows are indicated by different point symbols. For each different pre-period window, I show four different post-period window estimates: one week, three weeks, five weeks, and seven weeks. The length of the post-period window is indicated on the x-axis (in days). Note that estimates for a given post-period length are "dodged" for clarity to prevent over-plotting; the actual post-period used is the same for each cluster of estimates.

Starting with hiring, in the top left panel of Figure 6, there is no evidence that fill rates changed substantially post-announcement. Although all point estimates are negative, only the one week post-period estimates are conventionally significant for all pre-period bandwidths. These are also the least precise estimates. The four week pre-period bandwidth is always negative and significant, albeit marginally, for each post-period bandwidth.

For the imposition, only the estimates using the one week post-period window are negative. However, these estimates are close to zero and not conventionally significant. With a larger post-period window, the effects become positive and highly significant. In short, not only is there no decline in hiring, there is evidence of a substantial *increase* in hiring.

One possible explanation for the increase in hiring is that the composition of job openings changed. In the panel below the hiring results, I report the same set of coefficients but for regressions where the outcome is an indicator for whether the job opening was posted in the ADMIN category. Note that job opening compositional changes were not possible in the experiment, as job opening type was fixed pre-randomization. For the announcement, there is no strong evidence of a compositional shift, as the estimates are generally close to zero for all post-period windows. In contrast, following the imposition, the

[22] See Appendix C.7 for the difference-in-differences results on hiring decomposed into actual and placebo year event studies.

Figure 6: Estimates of the effects of the platform-wide $3/hour minimum wage on hiring and job composition

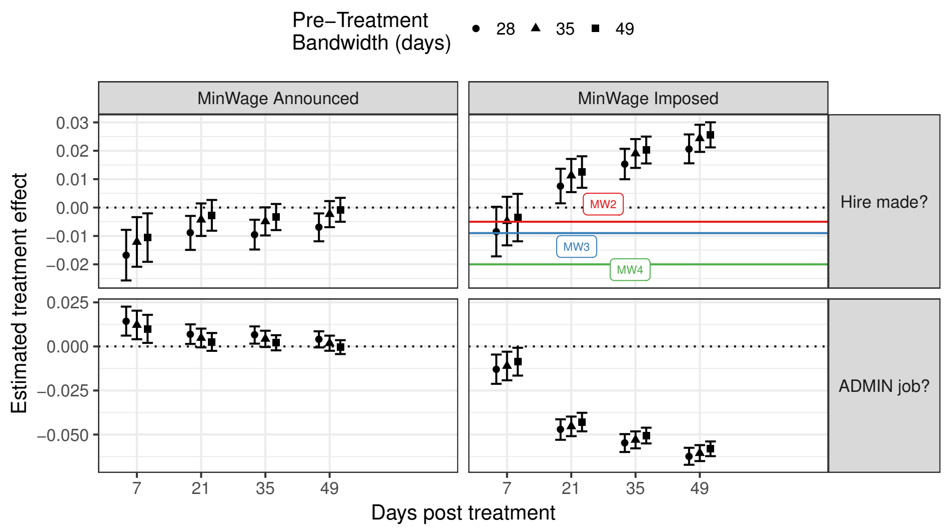

Notes: This figure plots the difference-in-differences estimate of the effect of (1) announcing the minimum wage and (2) implementing the minimum wage. The left column shows the "announcement" estimates and the right column the "imposition" estimates. The x-axis shows the estimated treatment effect using different post-period windows around the event (in days). The y-axis is the estimated treatment effect taking the actual year estimate minus the estimate calculated from the placebo year (one year prior), i.e., $\hat{\beta}_{\text{ACTUAL}} - \hat{\beta}_{\text{PLACEBO}}$. When applicable, experimental estimates are overlaid on the plot.

fraction of jobs posted in ADMIN fell substantially. All estimates, regardless of the size of the pre- and post-windows, are negative and highly significant, with reductions of about 5 to 7 percentage points for the largest post-period bandwidth. However, the fraction of job openings in a category waxes and wanes, and so it is unclear how credible difference-in-differences results are for this outcome. See Appendix C.8 for more exploration of this issue and the daily time series of job openings posted in ADMIN.

6.2 Employer selection and post-hire outcomes

As in the experiment, many of the outcomes of interest are only observed if a worker is hired. The effects of the announcement and imposition on these outcomes are shown in Figure 7. It shows difference-in-differences estimates using the same methodology as for the hiring and job opening composition results from Figure 6.

The outcome in the top panel of Figure 7 is the hourly rate of the hired worker. There is no evidence that the announcement had any effect, with most of the confidence intervals comfortably including zero. In contrast, hired wages increased substantially after the imposition, with point estimates ranging from 10% to 15%, depending on the post-period window length. The estimates are somewhat sensitive to the pre-period window used, with larger windows implying smaller effects, but the estimates do not seem sensitive to length of the post-period windows used. These estimates are larger than the MW3 experimental estimates but are close to the MW4 experimental estimates (recall Figure 4).

In the next panel down, the outcome is the profile rate of the hired worker. There is no evidence of an announcement effect but strong evidence of an imposition effect. Profile rates were about 10% higher, regardless of the post-period window length. This increase is substantially higher than the experimental increase in any of the cells. Although this would seemingly imply even greater labor-labor substitution in equilibrium, it is important to note that workers are free to change their listed profile rates at any time. Post-imposition, work-

Figure 7: Estimates of the effects of the platform-wide \$3/hour minimum wage on filled opening outcomes

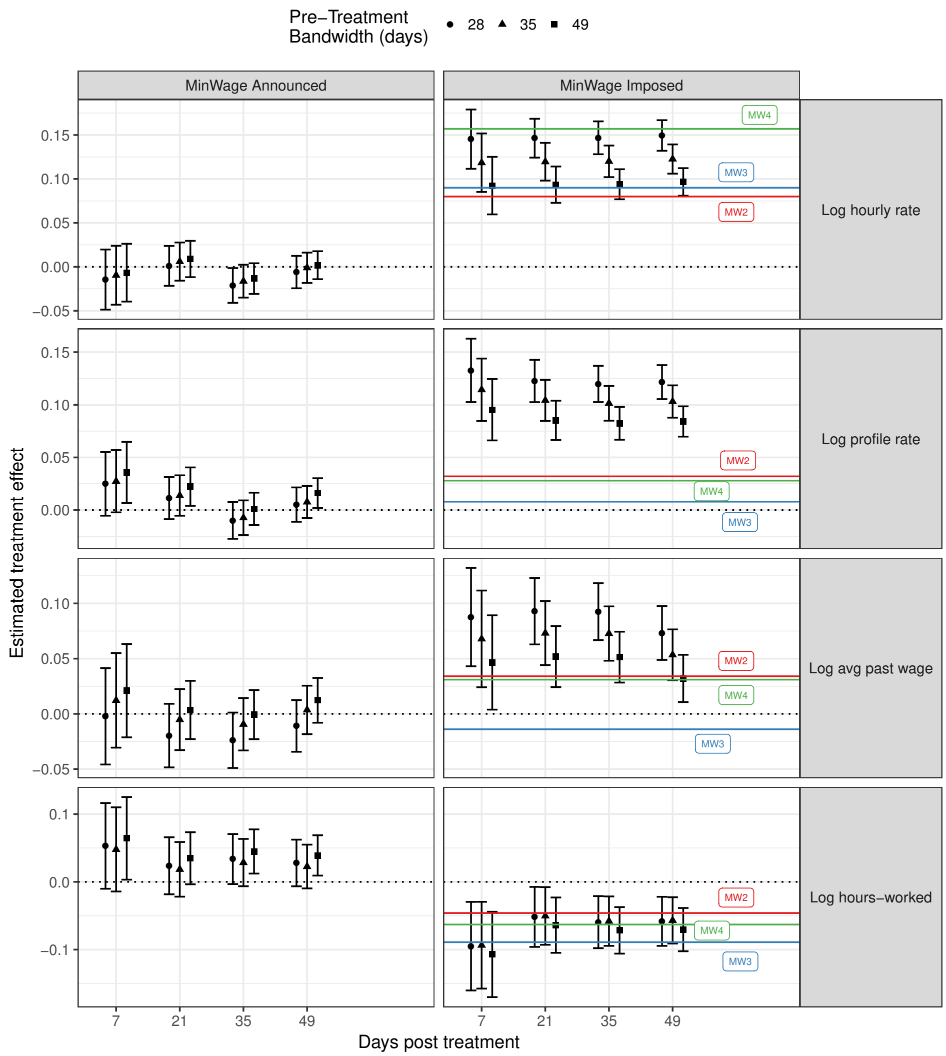

Notes: This figure plots difference-in-difference estimates of the effect of announcing and imposing a \$3/hour platform-wide minimum wage. The left column shows the "announcement" estimates and the right column shows the "imposition" estimates. The x-axis shows the length of the post-period window (in days). The y-axis is the estimated treatment effect taking the actual year estimate minus the estimate calculated from the placebo year (one year prior), or $\hat{\beta}_{\text{ACTUAL}} - \hat{\beta}_{\text{PLACEBO}}$. When applicable, experimental estimates are overlaid on the plot.

ers presumably changed their profile rates to reflect the new minimum wage policy.

The past average wage of the hired worker does not suffer from the same limitations as the profile rate, as workers cannot change it. This past average wage of the hired worker is the outcome in the third panel of Figure 7. There is no evidence of an announcement effect, but strong evidence of a positive imposition effect. The point estimates vary, but they are about 5%. This is higher than the MW3 experimental estimates, which were actually negative for MW3 in ALL. The results suggest there was also substitution towards higher wage workers after the market-wide imposition.

In the bottom panel Figure 7, the outcome is the number of hours-worked. Interestingly, there is perhaps some weak evidence of more hours-worked after the announcement, which would be consistent with employers trying to get work done in anticipation of the upcoming policy change. However, the effect is not large, and not all specifications give point estimates that are conventionally significant. In contrast, following the imposition of the minimum wage, hours-worked fall substantially. The point estimates imply a 6% reduction in hours-worked in the post-period. The experimental estimate for MW3 was about a 9% reduction, though this was the largest reduction among the active treatment cells—in MW4 and MW2 the reduction was closer to 5%, suggesting that MW3 was a high estimate of the true causal effect due to sampling variation.

The market-wide imposition difference-in-differences estimates generally match the experimental outcomes in both sign and magnitude, with the notable exception that hiring does not seem to decrease at all post-imposition. It is clear that the average past wage of the hired worker increased substantially, with point estimates somewhat larger than those found in the experiment.

6.3 Effects of market-wide imposition on workers

Both the experimental and difference-in-differences evidence show that firms adjusted to the platform-wide minimum wage by hiring more productive workers. A natural question is how this substitution affected workers in different parts of the pre-imposition wage distribution. To explore this question, I construct a dataset of all applications sent to job openings in the 14 days before and 14 days after the imposition date, in both the actual year and the placebo year (one year prior), and then compare the wage bid workers proposed and whether the application leads to a hire, conditioned on pre-period wage bidding.

As workers can send multiple applications, applications are nested within the worker. This multiple applications per worker structure is useful, as it allows for a within-worker estimate of the effects of the minimum wage on application behavior. To account for the nested structure of the data, I include worker-specific fixed effects and cluster standard errors at the level of the individual worker. To capture the pre-imposition place of a worker in the wage bid distribution, I segment workers into K "bands" based on their average wage bid in the pre-period. I then estimate a regression of the form

$$y_{ij} = \sum_{k \in K} \beta_k \left(\textsc{Post}_{ij} \times \textsc{PreWageBand}_i^k \right) + c_i + \epsilon, \tag{5}$$

where i indexes workers, j indexes job openings applied to, and $\textsc{Post}_{ij}$ is an indicator that the application was sent to a job opening posted after the imposition date of the platform-wide minimum wage. The $\textsc{PreWageBand}_i^k$ is an indicator for whether worker i had an average wage bid in the pre-period that was in band k. The c_i is an individual worker fixed effect. The coefficients of interest are the collection of β_k coefficients. Figure 8 plots the $\hat{\beta}_k$ coefficients for a collection of worker-level outcomes. For each outcome, the estimates are plotted using a solid line for the actual imposition year and a dashed line for the placebo year.

The top panel outcome is the worker's individual wage bid in logs. In both the actual and placebo years, workers bidding a low wage in the pre-period bid higher in the post-period. However, in the actual year, workers with below-minimum wage bids in the pre-period bid substantially higher in the post-period. For example, workers in $[2, 3)$ in the placebo year increased their

Figure 8: Changes in wage bids, hire probability, and search intensity after the implementation of a platform-wide minimum wage

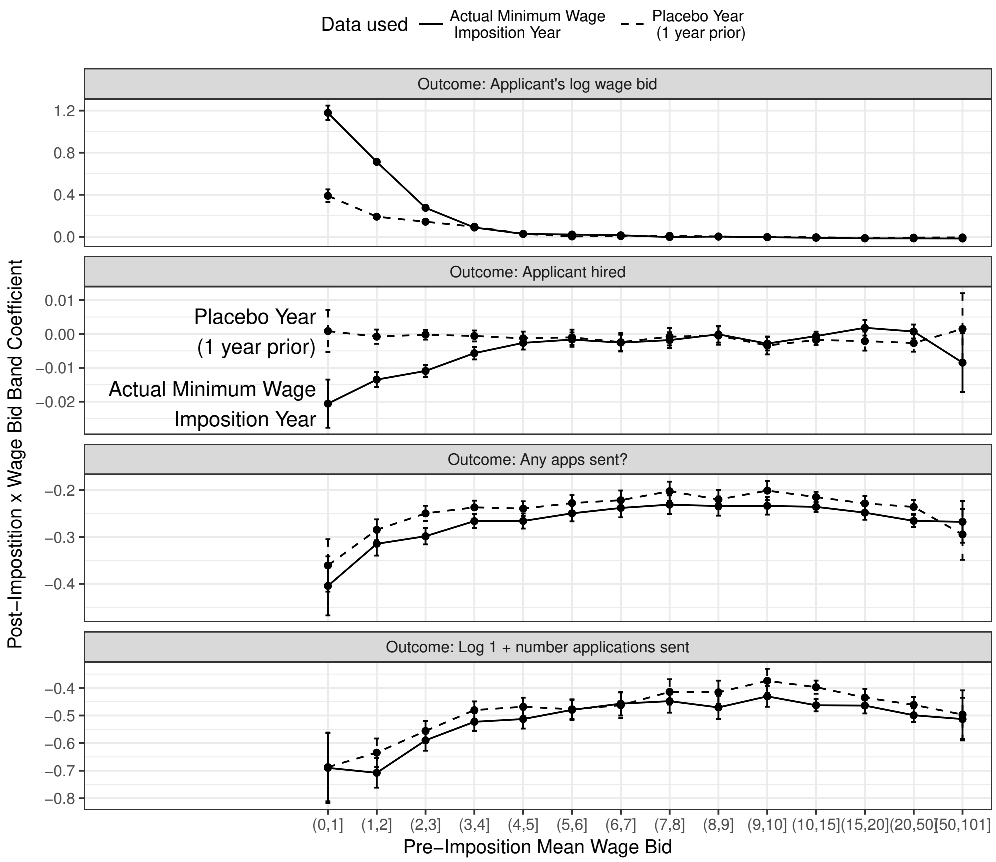

Notes: This figure shows the β^k coefficients from Equation 5 (in the top two panels). The sample consists of all job applications to hourly job openings 14 days before and 14 days after the minimum wage imposition. The top panel shows the change in wage bids in the post-period relative to the pre-period, by pre-period average wage bid. The next panel down shows the change in the application success rate relative to the pre-period, by pre-period average wage bid. In the bottom two panels, the coefficients are shown for an analogous regression where the outcome is the log count of applications (plus one), or any indicator for whether any applications were sent, respectively. For all regressions, standard errors clustered at the level of the individual worker. 95% confidence intervals are shown around each point estimate.

bids by about 20%, whereas in the treatment year, those workers increased their wages bids by nearly 80%. Workers who were well above the \$3/hour minimum had essentially no change in their wage bids relative to the placebo (or the pre-period for that matter—all points estimates are close to zero).

In the next panel down, the outcome is an indicator for whether the applying worker was hired for the associated job opening. We can see that in the placebo year there is essentially no change from the post- to pre-periods, with all point estimates close to zero. In contrast, for the imposition year results, we can see that those same workers who had to bid up to meet the new minimum wage suffered a decrease in their success probability.

To get a sense of the magnitude, consider the $(2, 3]$ band workers, who bid about 10% higher relative to what they "should" have bid, given the increase in the placebo. This led to about a 1 percentage point decrease in the per-application win probability. While this may not seem large, the average per-application hire rate for workers in this band is just 0.015, implying that the per-application success probability is less than half of what it was before the change.

Workers might potentially offset this reduction in hiring probability with more intensive search and application intensity, but they also might exit the market. If other workers increase the number of applications sent but do not exit the market, the equilibrium reduction in success probabilities might be even greater (or this already-large reduction in hire probability already reflects this equilibrium adjustment).

In the next panel down, the outcome is an indicator for whether a worker active in the pre-period sent at least one application in the post-period. We can see that in the treatment year workers are somewhat less likely to send an application in the post period relative to the placebo year. However, there is no evidence this differs by pre-period wage band. In the bottom panel, the outcome is the log count of applications, plus 1. Across wage bands, we see that the point estimates are all negative, which is expected given mean reversion. Furthermore, in the actual year, there is some evidence of a reduction in application count, but does not seem to depend on the pre-period wage band.

7 Discussion and Conclusion

To summarize the experimental findings, the experiment showed that for a firm facing a minimum wage: (1) the wages of hired workers increases, (2) at a sufficiently high minimum wage, the probability of hiring goes down, (3) hours-worked decreases at much lower levels of the minimum wage, and (4) the size of the reductions in hours-worked can be parsimoniously explained in part by the substantial substitution of higher productivity workers for lower productivity workers.

The observational findings are that there is little decrease in hiring after the imposition of the minimum wage, but some evidence of a reduction in the posting of job openings likely to pay minimum wages. The wage of hired workers increased substantially after the imposition of the minimum wage, in line with the experimental estimate. As in the experiment, firms substitute towards more productive workers and hours-worked fall. After the imposition, workers that historically bid below the minimum wage raised their wage bids substantially and experienced a reduction in their probability of being hired, per application. There is no evidence that these affected workers exited the market or changed their application intensity. However, it is important to note that the post-period is only two weeks.

A key finding of the paper is that labor-labor substitution is an important margin of adjustment for firms in this market facing a minimum wage. This kind of substitution is conceptually distinct from the typical framing of labor-labor substitution, in which workers have "types" but are imperfectly substitutable in the productive process, as in Katz and Murphy (1992). The substitution I find occurred within a pool of applicants that had all self-classified as being suitable for that particular job opening. The tasks are well-defined and so applicants are unlikely to offer radically different ways of performing the same task. The substitution here seems to be happening with respect to technical productivity, and yet in a competitive market, these workers should already be getting their marginal product, creating something of a puzzle. I explore these issues in Appendix D, sketching out a simple model of a competitive labor market that still allows for labor-labor substitution. In a nutshell, the degree of substitution and the effects of a minimum wage depend on the elasticity of labor supply of more productive workers. If they are highly elastic, then employers can readily shift production to these workers with little increase in labor costs, as there are no price spill-over effects for these workers.

The labor-labor substitution findings from both the experiment and the observational analysis are readily apparent when using individual worker productivity proxies. This substitution is much harder to see when using demographic proxies, which are typically the only proxies available in conventional settings.

Although labor-labor substitution was important to explaining the increase in wages, it is important to keep in mind that substitution was not wholly responsible for the increase in wages: observationally similar workers were paid a higher wage under the minimum wage treatment, though earnings did not detectably increase. However, if the reduction in hiring was also small in equilibrium, those relatively higher-productivity workers would presumably see an earnings increase, as they can do more projects at a higher wage.

Two directions of research grounded in conventional settings would be particularly welcome: (1) more research that quantified the extent to which firms observe the productivity of both applicants and their existing workforce (not already captured by wages), and to what extent can firms adjust this composition through hiring and firing[23]; (2) at a market level, what is the labor supply elasticity of "high-type" workers, particularly those currently out of the labor force (such as teens from high socioeconomic backgrounds). Of course, researching (1) in a conventional setting would be challenging, but with increasing computer-mediation of all aspects of the labor market matching process, this could change, perhaps making it possible to collect data on the pool of applicants available for jobs and the hired workers, before and after a minimum wage change.

[23] Altonji and Pierret (2001) is clearly related to this question.

References

Aaronson, Daniel, Eric French, Isaac Sorkin, and Ted To, "Industry dynamics and the minimum wage: A putty clay approach," *International Economic Review*, Forthcoming.

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.

Akerlof, George A and Janet L Yellen, "The fair wage-effort hypothesis and unemployment," *The Quarterly Journal of Economics*, 1990, *105* (2), 255–283.

Allegretto, Sylvia A, Arindrajit Dube, and Michael Reich, "Do minimum wages really reduce teen employment? Accounting for heterogeneity and selectivity in state panel data," *Industrial Relations*, 2011, *50* (2), 205–240.

Altonji, Joseph G and Charles R Pierret, "Employer learning and statistical discrimination," *The Quarterly Journal of Economics*, 2001, *116* (1), 313–350.

Barach, Moshe and John J Horton, "How Do Employers Use Compensation History?: Evidence from a Field Experiment," *Working Paper*, 2017.

Barach, Moshe E. and John J. Horton, "How do employers use compensation history? Evidence from a field experiment," *Working Paper*, 2017.

Blake, Thomas and Dominic Coey, "Why marketplace experimentation is harder than it seems: The role of test-control interference," in "Proceedings of the fifteenth ACM conference on Economics and computation" ACM 2014, pp. 567–582.

Burdett, Kenneth and Dale T Mortensen, "Wage differentials, employer size, and unemployment," *International Economic Review*, 1998, pp. 257–273.

Card, David, "Using regional variation in wages to measure the effects of the federal minimum wage," *Industrial and Labor Relations Review*, 1992, *46* (1), 22–37.

— **and Alan B. Krueger**, "Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania," *The American Economic Review*, 1994, *84* (4), 772–793.

Clemens, Jeffrey and Michael Wither, "The minimum wage and the great recession: Evidence of effects on the employment and income trajectories of low-skilled workers," *Working Paper*, 2014.

Draca, Mirko, Stephen Machin, and John Van Reenen, "Minimum wages and firm profitability," *American Economic Journal: Applied Economics*, 2011, *3* (1), 129–151.

Dube, Arindrajit and Ben Zipperer, "Pooling multiple case studies using synthetic controls: An application to minimum wage policies," *Working Paper*, 2016.

— **, T William Lester, and Michael Reich**, "Minimum wage effects across state borders: Estimates using contiguous counties," *The Review of Economics and Statistics*, 2010, *92* (4), 945–964.

Fairris, David and Leon Fernandez Bujanda, "The dissipation of minimum wage gains for workers through labor-labor substitution: Evidence from the Los Angeles living wage ordinance," *Southern Economic Journal*, 2008, *75* (2), 473—496.

Flinn, Christopher J, "Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates," *Econometrica*, 2006, *74* (4), 1013–1062.

Fradkin, Andrey, "Search, matching, and the role of digital marketplace design in enabling trade: Evidence from Airbnb," *Working Paper*, 2016.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani, "glmnet: Lasso and elastic-net regularized generalized linear models," *R package version*, 2009, *1.*

Gilchrist, Duncan S., Michael Luca, and Deepak Malhotra, "When 3 + 1 >4: Gift structure and reciprocity in the field," *Management Science*, 2016, *62* (9), 2639–2650.

Giuliano, Laura, "Minimum wage effects on employment, substitution, and the teenage labor supply: Evidence from personnel data," *Journal of Labor Economics*, 2013, *31* (1), 155–194.

Hirsch, Barry T, Bruce E Kaufman, and Tetyana Zelenska, "Minimum wage channels of adjustment," *Working Paper*, 2011.

Horton, John J., "Online labor markets," *Internet and Network Economics: 6th International Workshop, WINE 2010, Stanford, CA, USA, December 13-17, 2010. Proceedings*, 2010.

— , "Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution," *Working Paper*, 2017.

— , "The effects of algorithmic labor market recommendations: Evidence from a field experiment," *Journal of Labor Economics*, 2017, *35* (2), 345–385.

Jurka, Timothy P, Loren Collingwood, Amber Boydstun, Emiliano Grossman, and Wouter van Atteveldt, "RTextTools: Automatic text classification via supervised learning," 2012.

Katz, Lawrence F and Alan B Krueger, "The effect of the minimum wage on the fast food industry," *Industrial and Labor Relations Review*, 1992, *46* (1), 6–21.

— **and Kevin M Murphy**, "Changes in relative wages, 1963–1987: Supply and demand factors," *The Quarterly Journal of Economics*, 1992, *107* (1), 35–78.

Lazear, Edward P, Kathryn L Shaw, and Christopher T Stanton, "The value of bosses," *Journal of Labor Economics*, 2015, *33* (4), 823–861.

Luca, Dara Lee and Michael Luca, "Survival of the fittest: The impact of the minimum wage on firm exit," *Working Paper*, 2017.

Meer, Jonathan and Jeremy West, "Effects of the minimum wage on employment dynamics," *Journal of Human Resources*, 2015.

Neumark, David and William Wascher, "Employment effects of minimum and subminimum wages: Panel data on state minimum wage laws," *Industrial and Labor Relations Review*, 1992, *46*, 55.

— **, JM Ian Salas, and William Wascher**, "Revisiting the minimum wage–employment debate: Throwing out the baby with the bathwater?," *ILR Review*, 2014, *67* (3 Supplemental), 608–648.

— **, Mark Schweitzer, and William Wascher**, "Minimum wage effects throughout the wage distribution," *The Journal of Human Resources*, 2004, *39* (2), pp. 425–450.

Obenauer, Marie L. and Bertha von der Nienburg, "Effect of minimum-wage determinations in Oregon," *Bulletin of the United States Bureau of Labor Statistics*, 1915, (176).

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

Powell, David, "Synthetic control estimation beyond case studies: Does the minimum wage reduce employment?," *Working Paper*, 2016.

Romer, David, "Why do firms prefer more able workers?," *Working Paper*, 1992.

Schmitt, John, "Why does the minimum wage have no discernible effect on employment?," *Center for Economic and Policy Research*, 2013.

Silverman, Bernard W., *Density estimation for statistics and data analysis*, Vol. 26, CRC press, 1986.

Sorkin, Isaac, "Are there long-run effects of the minimum wage?," *Review of Economic Dynamics*, 2015, *18* (2), 306–333.

Stanton, Christopher T. and Catherine Thomas, "Landing the first job: The value of intermediaries in online hiring," *The Review of Economic Studies*, 2016, *83* (2), 810.

Teulings, Coen N., "Aggregation bias in elasticities of substitution and the minimum wage paradox," *International Economic Review*, 2000, *41* (2), 359–398.

Tibshirani, Robert, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, 1996, pp. 267–288.

A Online Appendix: Internal validity

A.1 Compliance

To test whether the experimental intervention was effective—in the sense of preventing contracts from being formed below the cell-specific minimum—in Figure 9 I plot the distribution of hired worker hourly wages by experimental cell, on a log scale. The hourly wage is calculated by taking the total wage bill for each contract and dividing by the total number of hours-worked. The bars in the histogram are each \$1 wide, with intervals $[a, a+1)$, where a is an integer.

The top panel of the figure shows the distribution for the control group. In the control, there are substantial numbers of hired workers that received less than the lowest minimum wage found in MW2. Below the control, the hired worker wage distribution is shown for the three active treatment cells. In each of the active treatment cells, the mass of observed hourly wages is nearly all to the right of the imposed minimum wage for that cell.[24]

A.2 Randomization

Table 2 shows the means for a host of pre-treatment job opening outcomes for both the control and MW4. We can see that there differences are all close to zero and none of the differences are conventionally statistically significant. In terms of job opening attributes, "Technical" is an indicator for whether the job opening required some kind of computer programming. "Admin" and "Software Dev." are indicators for more-refined self-assess categories.

For the other job opening attributes, "New buyer?" is an indicator for whether the buyer had ever used the platform before by posting a job opening; "Prefers high quality" is an indicator for whether the buyer stated ex ante that they were looking for the most experienced, highest wage workers; the job description length is the length of the buyer's job description measured in

[24]The small number of non-complying observations is due to workers offering refunds to their firms, lowering the wage bill but keeping the hours-worked the same, lowering the effective wage.

Figure 9: The realized wage distributions for hired workers in ALL, by experimental group

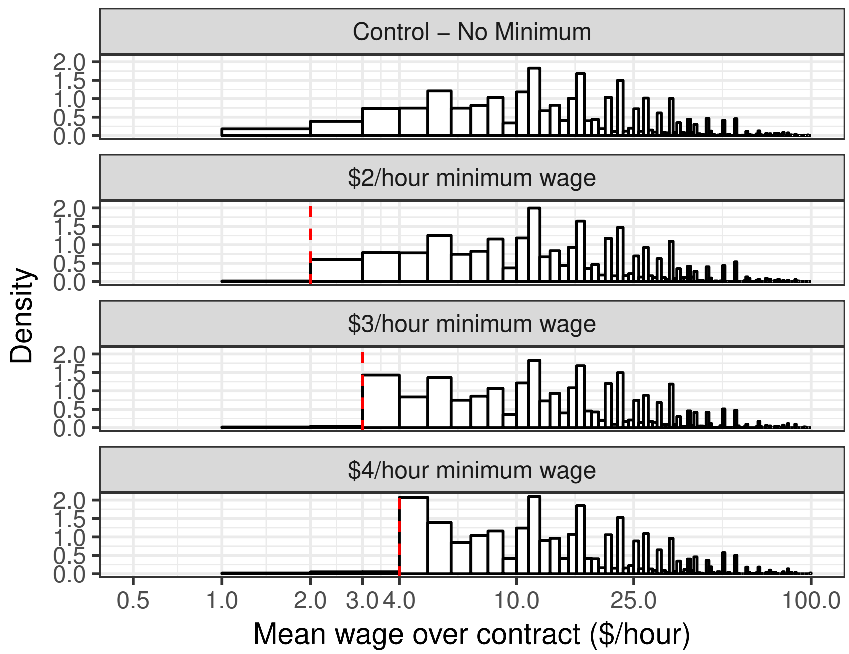

Notes: This figure shows a density histogram of log observed hourly wages in each of the experimental cells. The x-axis is on a log scale. The bars in the histogram are each $1 wide, with intervals of $[a, a+1)$, where a is an integer.

Table 2: Comparison of pre-treatment covariates for the control and MW4 groups as a check of randomization

	Treatment mean: $\bar{X}_{TRT}$	Control mean: $\bar{X}_{CTL}$	Difference in means: $\bar{X}_{TRT} - \bar{X}_{CTL}$	p-value
Observation Counts				
	9,725	91,781		
Type of work				
Technical (1 if yes, 0 otherwise)	0.426 (0.005)	0.422 (0.002)	0.004 (0.005)	0.471
Type of work—(more detailed)				
Admin	0.113 (0.003)	0.114 (0.001)	-0.001 (0.003)	0.832
Software Dev.	0.124 (0.003)	0.122 (0.001)	0.002 (0.004)	0.530
Vacancy attributes				
New employer?	0.783 (0.004)	0.782 (0.001)	0.002 (0.004)	0.716
Prefers high quality?	0.211 (0.004)	0.209 (0.001)	0.002 (0.004)	0.595
Has employees already?	0.075 (0.003)	0.079 (0.001)	-0.004 (0.003)	0.139
Log job description length (chars)	5.734 (0.011)	5.731 (0.004)	0.004 (0.012)	0.770
Log prior spend + 1	1.459 (0.030)	1.477 (0.010)	-0.018 (0.032)	0.583

Notes: This table reports pre-treatment covariate means for the MW4 and control groups.

characters of text; and "prior spend" is the cumulative amount of money paid by the buyer on wages prior to the experiment.

A.3 Workers sorting across openings

The first test of sorting is whether applicant counts differed by experimental cell. Figure 10 reports the regression results where the outcome is the log number of applications per job opening. The sample is restricted to job openings that received at least one application. The right panel shows the fitted regression line with the minimum wage as a regressor.

In the population, the counts are very slightly negative, but far from significant. In the sub-populations where would expect larger results, the estimates are simply less precise, with the MW2 cell in LPW now having a positive sign. As we expected, it seems that workers neither avoided nor queued for job openings with imposed minimum wages.

Figure 10: Effects of the minimum wage on log number of organic applications received

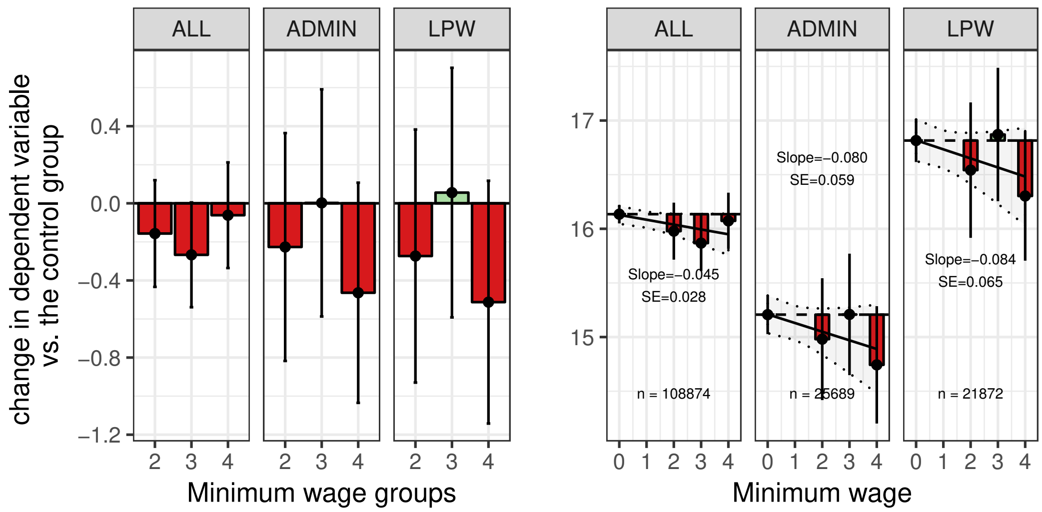

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the log number of organic applications received. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$..

A.4 Firms seeking to avoid the minimum on the platform

In the conventional minimum wage research there is a hard-to-address concern which is that firms have beliefs that affect their choices about how to react to a minimum wage. For example, a firm anticipating a minimum wage might make different hiring and capital choices compared to a firm not anticipating the change. This "expectations" problem is less likely to be a concern on the platform for several reasons. Unlike a conventional minimum wage, there was no observable public political process around the minimum wage, nor any public announcement. The platform used for the present study never had a minimum wage and made no public indications that it was considering such a change. However, this lack of knowledge raises another concern: perhaps firms thought they received an idiosyncratically bad "draw" of applicants and by re-posting the job opening, they could do better.

Figure 11: Effects of the minimum wage on a follow-on opening by the employer

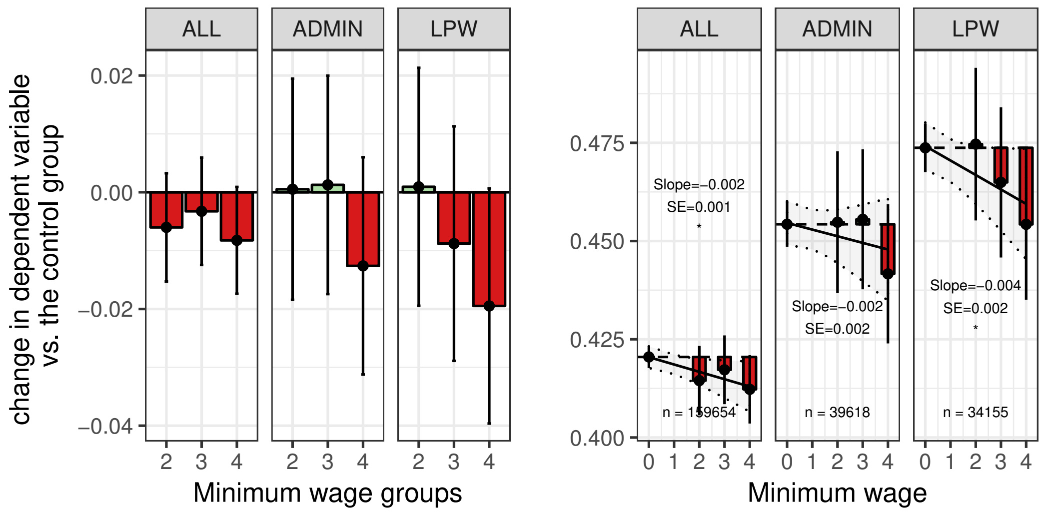

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is whether the employer posted a second job opening after his or her first opening. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$..

This re-posting hypotheses would tend to over-state the extensive margin reductions in hiring. Given that these effects are already minuscule, it there is not much "room" for this kind of adjustment by firms. There is also no evidence for this phenomenon: Figure 11 reports the regression results when the outcome is an indicator for whether the buyer posted one or more follow-on job posts. If firms were re-posting because they thought they received an idiosyncratically bad draw, this effect should be positive. In the population, each active treatment cell has a negative coefficient, though none are significant. In the sub-populations ADMIN and LPW, the estimates are less precise and are not always the same sign. If anything, the negative effects are stronger in MW4 where the incentive to "hunt" would be strongest. This pattern is consistent with firms acting like price-takers, with firms getting many high-price applicants deciding not to abandon the platform (though again, the caveat here is that none of these effects are conventionally significant and all are close to zero).

A.5 Firms sorting across platforms

Although would-be employers have several options for low-wage, hourly administrative work, survey evidence suggests that relatively few firms "multi-home" by posting jobs on multiple platforms (see Section 2 for a discussion of how prevalent this in practice). However, if firms did respond to the minimum wage by posting their job opening on another market platform, they would have essentially two other options. During the period of the experiment, all of the major alternative platforms had minimum wages as well, though they differed in their level. Each opening in the experimental sample has a job title e.g., "Java Developer Needed for Short Project." Assuming firms posting on multiple sites would re-use their job titles, for each MW4 job title, I constructed an indicator for whether that exact job title appeared and alternative online labor market whose collection of job titles is available. The resultant fitted model is

$$\Pr\left(\text{Title Match on Alt. Plaform}\right) = \underbrace{0.0034}_{0.0069} \cdot \mathbb{1}\{\underline{w} = 4\} + \underbrace{0.155}_{0.0018} \qquad (6)$$

which shows that the minimum wage on the platform did not simply displace firms to the most natural alternative and closest substitute, at least in the short-run; the coefficient is close to zero.

B Online Appendix: Proofs

Proof of Proposition 1

Proof. The probability that a firm wants to hire a randomly selected applicant is $1 - G(\underline{u})$. If the firm gets n applicants, the probability of not hiring anyone is $\Pr(\text{No Hire}) = G(\underline{u})^n$.

If the workers who would have otherwise bid below the minimum respond to a minimum wage by increasing their wage bids to the minimum or higher, then there is a new distribution of profits. This new distribution is first order stochastically dominated by $G(\cdot)$, and so

$$\Pr(\text{No Hire}|\text{MinWage imposed}) > \Pr(\text{No Hire}).$$

□

Proof of Proposition 2

Proof. Suppose that in the absence of the minimum wage, the firm would hire a worker with wage bid w^* and technical productivity y^*. Assume that all workers that previously bid below $\underline{w}$ raised their bids to $\underline{w}$, and all other workers kept their bids the same. If $w^* > \underline{w}$, the raised minimum does not change the firm's choice and hours-worked stays the same. However, if $w^* < \underline{w}$, the firm has three options: (1) not hire anyone, (2) hire the same worker as before, or (3) hire a worker with higher technical productivity. The firm would not hire a worker with a lower productivity, say y', because this worker would also now be bidding the minimum wage, and if $y' < y^*$ then $\underline{w}/y^* < \underline{w}/y'$, and so the firm would prefer its original worker. If the firm hires a worker with higher technical productivity, then hours-worked falls, as the more productive worker takes less time to complete the project. □

C Online Appendix: Additional outcomes

C.1 Effects of the minimum wage on the hired worker earnings

For each job opening, I can calculate the total wage bill, which is the amount paid to hired workers. Figure 12 reports regressions with the total wage bill as the outcome, with zeros for unfilled job openings included. There is some weak evidence of a lower wage bill in MW3 (significant in ADMIN and LPW), but this result is somewhat undermined by the absence of effects in MW2 or MW4. Further, there is no evidence of a stronger pattern in ADMIN or LPW compared to ALL, as was the case with many other results. With no clear pattern in the earnings results across populations and experimental groups, it seems that the wage and hours effects are essentially offsetting, at least given the available precision.

Figure 12: Effects of the minimum wage on total wage bill

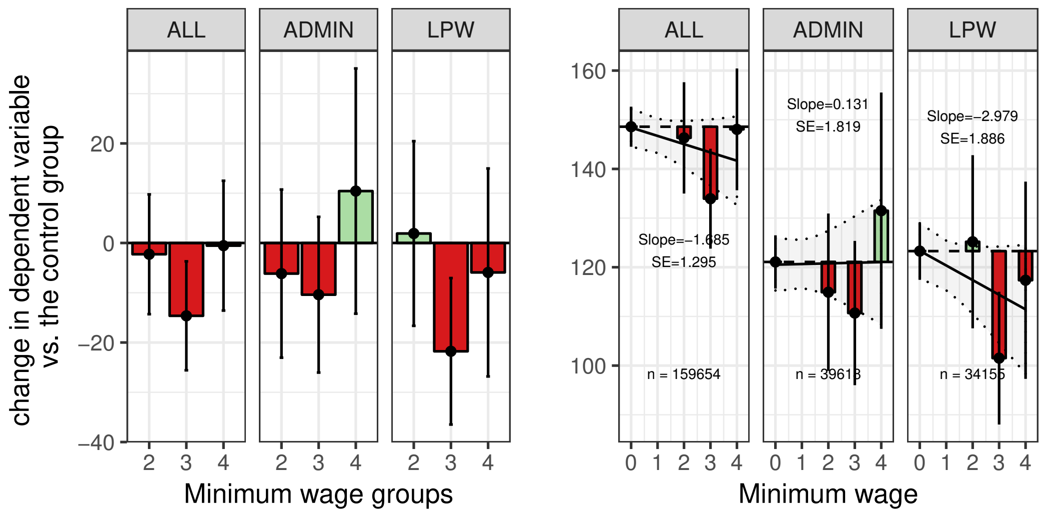

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the total wage bill. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$..

Given the imprecision of total earning estimates, I also examine changes in log earnings, conditional upon a job opening being filled. Figure 13 shows the results when the outcome variable is log earnings, conditional upon the hired worker earning some amount of money. With these estimates, there is fairly strong evidence for a larger wage bill in MW4, though this is also the cell that had a non-trivial reduction in hiring.

C.2 Hours-worked in total

In the left panel of Figure 14, we can see that the minimum wage reduced the number of hours-worked in all groups, across all sub-populations—only the MW2 group in LPW is marginally significant. The effects were largest

Figure 13: Effects of the minimum wage on log total wage bill

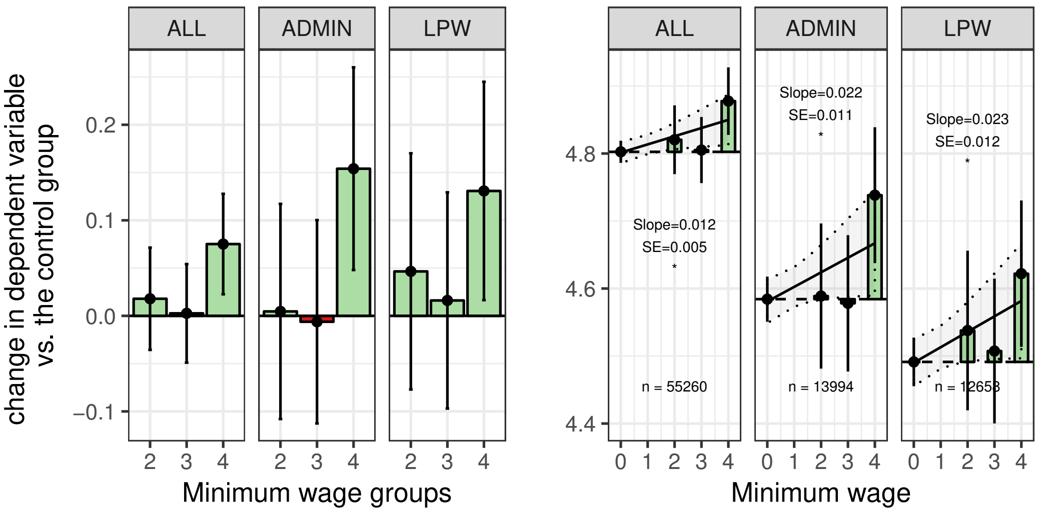

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the log total wage bill. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$..

in MW3—a cell that showed little to no extensive margin reduction in labor demand (recall Figure 2). Effects were larger in the target sub-populations of ADMIN and LPW, though given that these groups are composed of different work types, changes in absolute numbers cannot be interpreted as stronger effects.

C.3 Effects of the minimum wage on the wage bid "markup" of the hired worker

The left panel of Figure 15 reports the results from regressions of the markup of the hired worker on the cell indicators. The sample is restricted to hired workers with a listed profile rate. Across all cells and samples, we can see that wage bid markups were higher in the active treatment cells. The effects were larger in ADMIN and LPW sub-populations, with markups in MW4 as much as 30 percentage points higher in LPW and about 25 percentage points higher in ADMIN.

Figure 14: Effects of the minimum wage on hours-worked

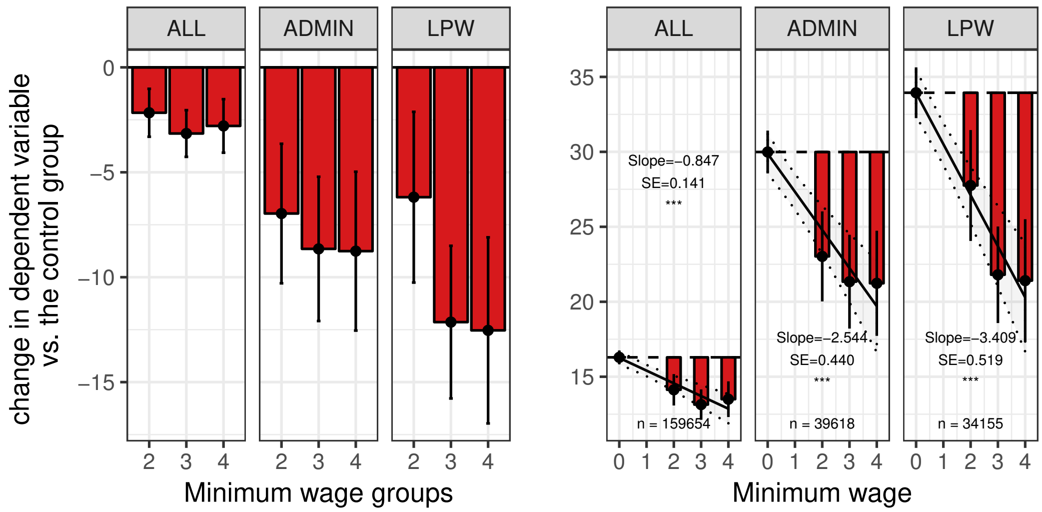

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the hours worked (including 0) 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$..

In the right panel Figure 15, the dashed horizontal line indicates the mean markup in the control. The increase in hired worker markup is large, even in cells with almost no reduction in hiring, such as MW3 in ADMIN. Note that the mean markup in the control is negative, as most workers bid below their profile rate. Also note that in ADMIN and LPW, the baseline markups are substantially lower. The markup results imply that firms facing minimum wages paid higher than expected wages.

C.4 Profile rate of the hired worker

One advantage of using the profile rate as a proxy for worker productivity is that it is available for all workers, even if they have never worked on the platform. Furthermore, it can potentially give a more accurate measure of the worker's market wage compared to average past wage, which can include wages from many long-completed jobs. Figure 16 reports regressions where the outcome is the log profile rate of the hired worker. In the left panel, in

Figure 15: Effects of the minimum wage on markup in the wage bid of the hired worker

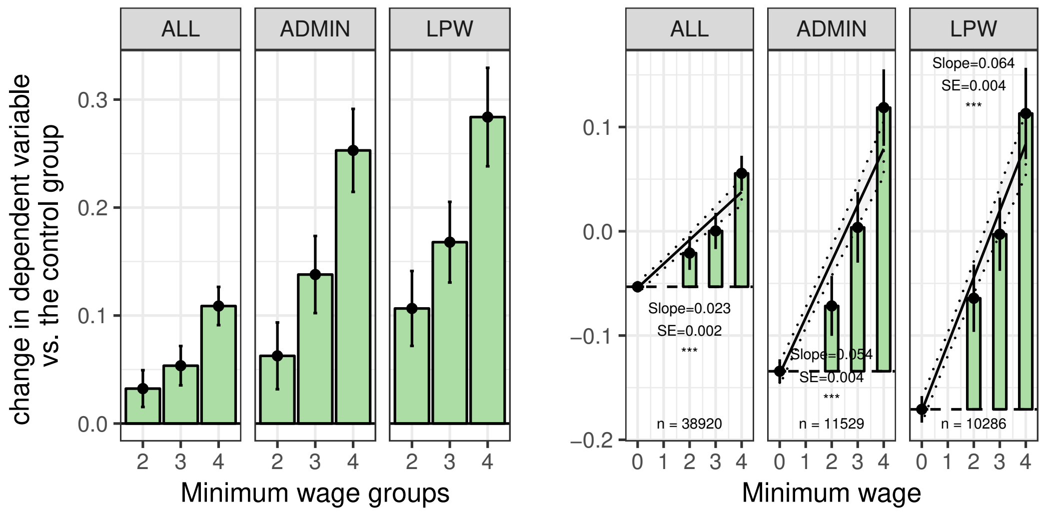

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the markup in the wage bid of the hired worker. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$..

ALL, hired worker profile rates are higher for each minimum wage, though only significantly so in MW2 and marginally so in MW4.

As with the other productivity proxies, in the sub-populations, the treatment effects are larger: in ADMIN, profile rates are about 7% higher in MW2 and MW3, and nearly 15% higher in MW4. In LPW, the MW4 effect size is about the same as in ADMIN. In the right panel of the figure, the regression line implies that in ADMIN and LPW, each additional $1 in the minimum wage is associated with about 3% higher profile rates for the hired worker.

C.5 Past earnings of the hired worker

Figure 17 reports regression results using log cumulative earnings of the hired worker as the outcome, conditional upon having at least $1 in past earnings. From the left panel of Figure 17, we can see that in ALL, log past earnings are higher in MW2 and MW4 but slightly negative in MW3. Only the MW4

Figure 16: Effects of the minimum wage on log profile rate of the hired worker

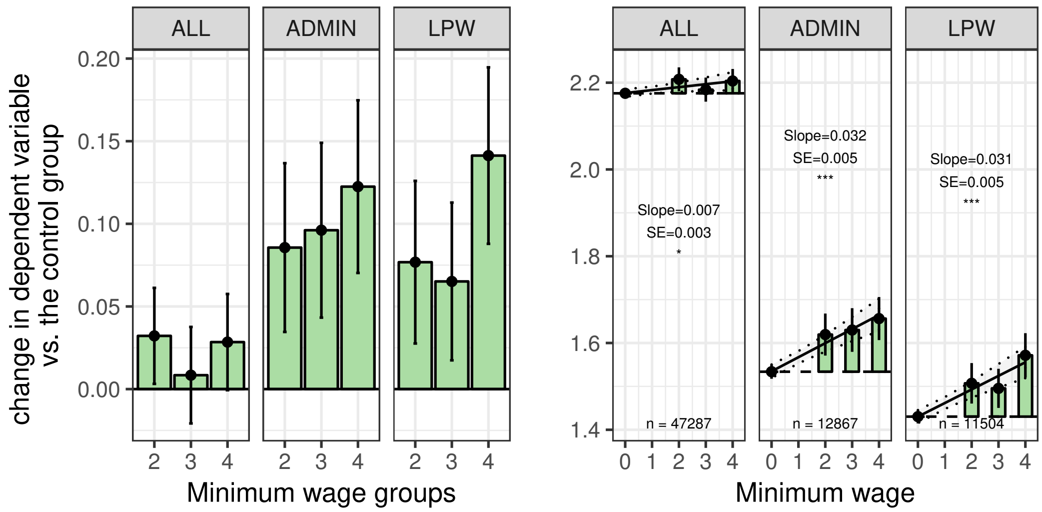

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the log profile rate of the hired worker. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$..

result is conventionally significant, though the magnitude is large—the increase is nearly 10%.

In the sub-populations, the effects are much stronger. In both ADMIN and LPW, hired workers in MW4 had nearly 25% higher cumulative past earnings compared to the control, while in MW2 and MW3, past earnings were 15% higher. From the right panel, for ADMIN and LPW, the slope is such that each additional $1 in the minimum wage was associated with 7% higher past earnings for the hired worker. The slope in ALL is also positive and significant, though the magnitude is smaller, with each $1 in the minimum wage associated with about 1.5% higher past earnings.

C.6 Any past experience of the hired worker

The outcome variable in the analysis shown in Figure 18 is an indicator for whether the hired worker had any on-platform work experience at the time they were hired. There is no evidence that the treatment affected the probability

Figure 17: Effects of the minimum wage on log past earnings of the hired worker

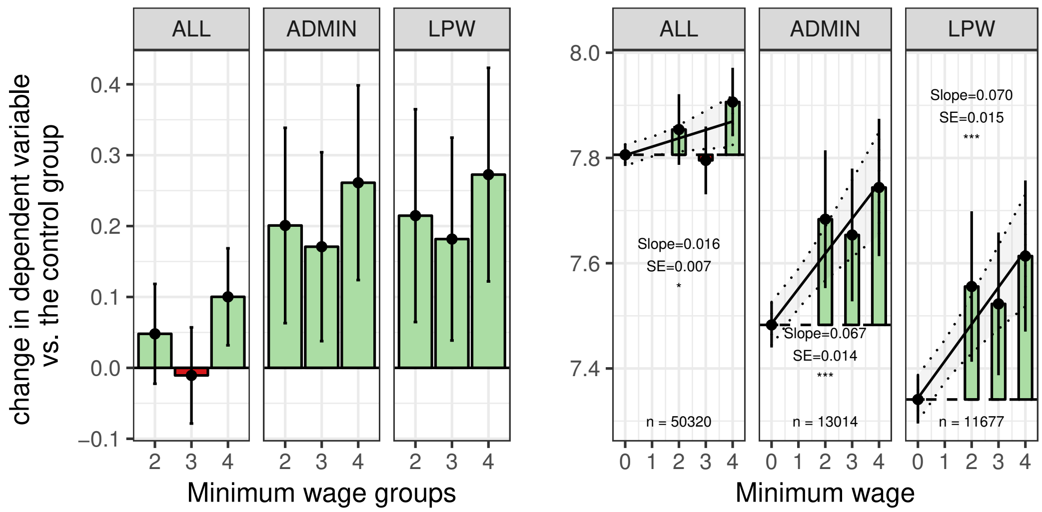

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is the log past earnings of the hired worker. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$..

that the hired worker had prior experience.

C.7 Effects of the announcement and imposition of the minimum wage on hiring

I constructed samples of hourly job openings posted some number of days before and after the imposition and announcement dates. I then constructed equivalently defined samples, but from one calendar year prior to serve as a comparison or "placebo" group.

To start, I present event study estimates. Table 3, Column (1) reports an estimate of

$$\mathbb{1}\{h_j > 0\} = \beta_0 + \beta_{\text{ACTUAL}}\text{POST}_j + \epsilon, \tag{7}$$

where POST_j is an indicator that job opening j was posted after the announce-

Figure 18: Effects of the minimum wage on an indicator for any past on-platform experience

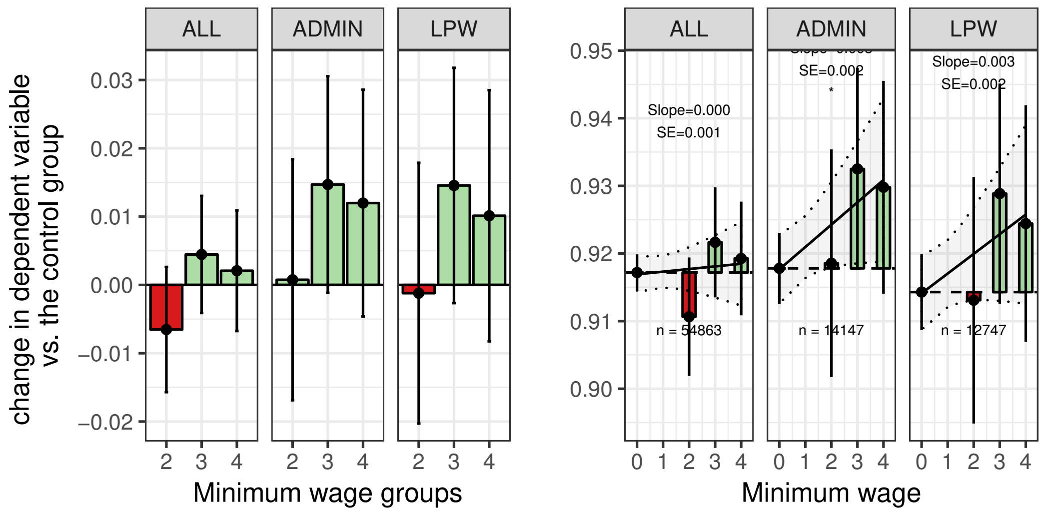

Notes: The left panel of this figure shows the treatment effects for each of the active treatment cells. The right panel shows the estimation regression line using the minimum wage as a regressor. The dependent variable is whether the hired worker had any past on-platform experience. 95% CI and 95% prediction intervals are shown in the left and right panels, respectively. Each panel shows results in three facets, labeled ALL, ADMIN, and LPW, corresponding to the sample used in that regression. For more details on these sample definitions, see Section 3.2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$..

ment of the \$3/hour minimum wage. The sample contains job openings two weeks before and two weeks after. Column (2) reports the same regression, but with a "placebo" sample constructed one calendar year prior. Columns (3) and (4) have the same outcomes as Columns (1) and (2) but use the corresponding samples for the minimum wage imposition rather than the announcement. The implied difference-in-differences treatment effect would be $\hat{\beta}_{\text{ACTUAL}} - \hat{\beta}_{\text{PLACEBO}}$.

Starting with the announcement, we can see in Column (1) that there is almost no difference in hiring following the imposition. The Column (2) regression from one calendar year prior also shows no effects, suggesting that there was no seasonal trend that would make an event study estimate misleading. The implied difference-in-difference estimate would be very close to zero, implying that the announcement had no effect on the probability that a job opening was filled.

Turning to the imposition, Column (3) shows a small negative effect on

Table 3: Event study on the effects of the platform-wide minimum wage announcement and imposition on whether a job opening was filled

	Dependent variable:			
	Anyone hired?			
	Announce (Actual)	Announce (Placebo)	Impose (Actual)	Impose (Placebo)
	(1)	(2)	(3)	(4)
Post	0.0001	0.007	−0.002	0.005
	(0.004)	(0.004)	(0.004)	(0.004)
Constant	0.395***	0.392***	0.382***	0.365***
	(0.003)	(0.003)	(0.003)	(0.003)
R^2	0.00000	0.0001	0.00001	0.00003

Notes: This table reports several OLS regressions where the dependent variable is an indicator whether a job opening was filled. The samples consist of all hourly job openings posted on the platform 14 days before and 14 days after a specified date. In Columns (1) and (2), the samples are built around the Julian date the platform-wide minimum wage was announced. Column (1) uses data from the actual announcement year, while Column (2) uses data from one year prior, or the "placebo" year. In Columns (3) and (4), the samples are built around the Julian date the platform-wide minimum wage was imposed. Column (3) uses data from the actual imposition year, while Column (4) uses data from the "placebo" year. The key independent variable is an indicator for whether the job opening appeared in the "Post" period i.e., after the specified date. Conventional standard errors are reported. Note that the sample size is intentionally omitted from the table. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$.

Figure 19: Fraction of hourly job openings that are administrative by day, in the year the minimum wage was imposed (solid line) and in the previous calendar year (dashed line).

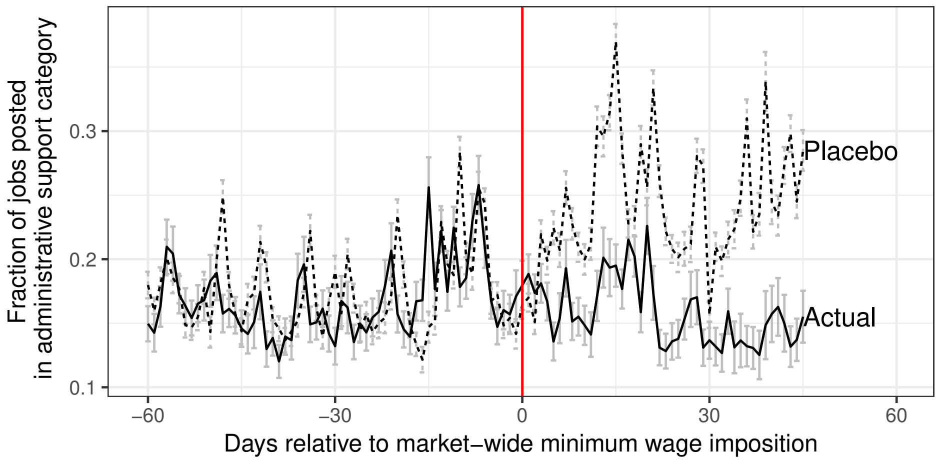

Notes: This figure plots the fraction of job openings posted in the administrative category. The solid line shows the fraction in the year the minimum wage was imposed. The dashed line shows the fraction in the "placebo year," which was one year prior. The "0" day, indicated with a vertical line, is the day the minimum wage was imposed.

hiring, while in the placebo year, the effect is small and positive. Although the implied difference-in-difference estimate is negative, implying a reduction in hiring, the effect is very small and far from significant. Consistent with the experimental results, there is little evidence of an equilibrium reduction in the fraction of posted jobs that are filled, at least with the two week pre- and post-period windows used.

C.8 Job compositional shift

Although the regression evidence suggests a compositional shift post-imposition, the credibility of this finding depends on the suitability of the placebo year as a counter-factual. To better assess this assumption, Figure 19 plots the daily fraction of jobs posted in the administrative category. The "0" day is the day the $3/hour minimum wage was imposed market-wide. The actual year fraction is shown as a solid line and the placebo year fraction is shown as a dashed line. A 95% CI is plotted for each daily estimate.

The two time series in Figure 19 track closely in the pre-period, but then diverge in the post-period. However, the post-period difference seems to be more caused by an increase in the placebo year not matched in the actual year. The post-period gap could, of course, be due changes in the market, in either year, not related to the minimum wage policy. The composition results would be more credible if both lines were more or less flat in the pre-period and the placebo continued to be flat in the post-period. Caveats aside, to the extent the placebo year is providing a reasonable counter-factual, there is some evidence of a post-imposition compositional shift away from relatively low-paying job openings.

C.9 Did affected workers shift to an "uncovered" sector of fixed price work?

One possible way that workers could adjust to platform-imposed minimum wage is to shift to an "uncovered" sector. On the platform, fixed price work is not covered by the minimum wage. However, Figure 20 shows, there is little evidence of such a shift. The outcome variable is an indicator for whether a worker's application was to a fixed price opening. Comparing the imposition year to th placebo year, there is no evidence that workers likely to be affected by the minimum wage began applying to more fixed price job openings.

D Online Appendix: Labor-labor substitution in competitive market

A natural question is whether the kind of labor-labor substitution observed in this experimental context can be reconciled with a competitive labor market model. The answer seems to depend on assumptions about the productive process. One plausible assumption is that a job has a fixed marginal technical productivity, regardless of who is hired e.g., all workers can produce one

Figure 20: Changes in application-focus post implementation of platform-wide minimum wage

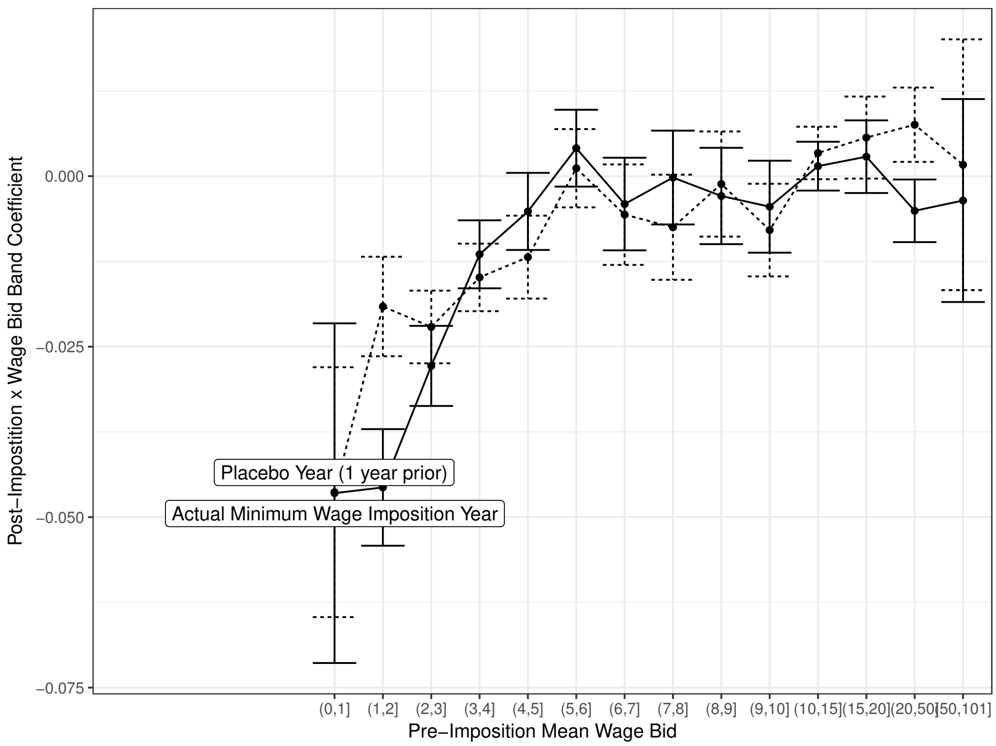

Notes: This figure shows the β^k coefficients from Equation 5, but with an outcome of being an indicator for whether the worker applied to a fixed price job opening. The sample consists of all job applications to hourly job openings 14 days before and 14 days after the experimental period. The equation is fit using OLS, with standard errors clustered at the level of the individual worker.

widget per hour. With this assumption, labor-labor substitution cannot be an adjustment strategy—either the minimum wage is above or below the marginal product of the job.

An alternative assumption is that workers can have heterogeneous technical productivity e.g., the job is making widgets, but one worker can produce one widget per hour, while another can produce two per hour, and so on. Wages would reflect these differences. If applicants to the same job opening have heterogeneous technical productivity, labor-labor substitution as a response to a minimum wage is not only possible—it explains too much, in the sense that the minimum wage could be completely undone by substitution, as firms buying labor in competitive market face a horizontal supply curve for every possible level of technical productivity. If each worker is paid their marginal product, then firms could always substitute towards higher productivity workers.

A possible solution to the puzzle is to assume that even if firms face a horizontal supply schedule, the market supply curve for different workers is upward sloping. To sketch a simple model, consider an industry with N firms, each with a production function $Y(L)$, where L is efficiency units of labor. Prices in the product market are normalized to one. There are two worker types, high and low, that differ only in their technical productivity, with the low-types offering $r < 1$ of the output in one unit of time as the high-types. The marginal rate of technical substitution between the two worker types is thus r.[25] Both types are inelastic on the intensive margin of labor supply, but elastic on the extensive margin.

The firm's profit maximization problem is to hire some combination of low-type workers, L_l and high-type workers, L_h, to maximize profits, or

$$\max_{L_h, L_l} \quad Y(L_h + rL_l) - w_l L_l - w_h L_l, \tag{8}$$

where w_h and w_l are the respective market wages of high- and low-types.

If all workers are paid their marginal revenue product, the wages of high- and low-type workers must satisfy $rw_H = w_L$. The firm demands L_E^* efficiency units of labor, but is indifferent over the precise worker composition, so long as $L_h^* + rL_l^* = L_E^*$. Despite this individual firm indifference, for simplicity I assume that all firms choose the same mix of high- and low-type workers. The market composition is then determined by the supply elasticities of the two worker types: in equilibrium, $S^h(w_h) = NL_h^*$ and $S^l(w_l) = NL_l^*$, where $S^h(\cdot)$ and $S^l(\cdot)$ are the high- and low-type labor supply curves, respectively.

Now consider how the equilibrium changes when a minimum wage, $\underline{w}$, is imposed. Assume this wage binds for the low-types but not for the high-types

[25] I have assumed workers are perfect substitutes, though models that allow imperfect substitution of even complementarity between different groups of workers might give richer insights—see Teulings (2000) for this kind of approach in the context of the minimum wage.

i.e., $w_l < \underline{w} < w_h$. Let the firm's demand for labor be such that each firm wants to hire $L'_E \leq L_E$ efficiency units of labor and assume that some low-type workers are still hired in equilibrium. If any low-types are still being hired, then $rw'_h = \underline{w}$, where w'_h is the new market wage of high-types.

The size of the $L_E - L'_E$ reduction in labor is "absorbed" by the two worker groups differently. To focus purely on labor-labor substitution, assume that labor demand is completely inelastic in efficiency units and so $L_E = L'_E$. If the high-type workers are fully inelastic in their labor supply, then there will be no compositional change in the workforce: an infinitesimal substitution in the direction of high-type workers will be enough to "restore" $rw'_H = \underline{w}$ and there will be no reduction in headcounts. At the other extreme, if the high-types are infinitely elastic, we will see the high-type workers completely replace the low-types, along with a reduction in headcounts. This reduction comes despite the firm's inelastic demand in efficiency units. There is a "pseudo" reduction in labor demand that comes from hiring more productive workers. Note that in both cases, the labor supply of the low-types is irrelevant, since $\underline{w} > w_l$ and so, at best, their unemployment levels will stay the same, as they are already willing to provide that much labor at the lower, original wage of w_l.

The model sketched above is, of course, highly stylized, but it does provide a potential explanation for how the kind of labor-labor substitution observed in this paper could exist in equilibrium in a competitive labor market.

E Online Appendix: Wages by category

Figure 21 shows box plots for the log wages for each on-platform category of work in the control group.[26]

[26] The sample is restricted to wages above 25 cents per hour which the worker worked at least one hour. There are a small number of contracts (0.2% of filled job openings) formed for very small hourly wages (usually 1 cent) though these are usually firms and workers that are using the platform's time tracking features but are not actually using the site for payment purposes.

Figure 21: Wages by category of work in the control group

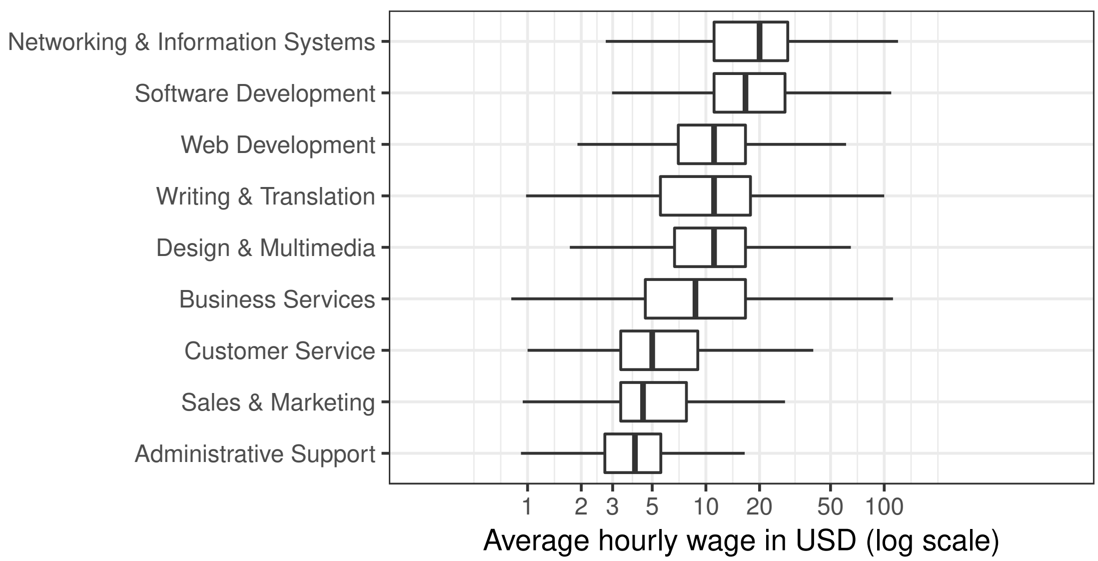

Notes: This figure shows the distribution of hourly wages hours for filled jobs in the control group, by category of work. The box indicates the 25th and 75th percentiles. The heavy center-line is the median. The whiskers are the highest and lowest values within 3/2 of the IQR, from the median.

F Online Appendix: Tables

Table 4: Effects of the minimum wage on whether anyone was hired for the job opening

	Dependent variable:		
	whether anyone was hired for the job opening		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	−0.020***	−0.027**	−0.044***
	(0.004)	(0.009)	(0.010)
MW3	−0.009	−0.006	−0.022*
	(0.004)	(0.009)	(0.010)
MW2	−0.005	−0.014	−0.012
	(0.004)	(0.009)	(0.010)
Constant	0.349***	0.357***	0.377***
	(0.001)	(0.003)	(0.003)
Observations	159,656	39,620	34,157
R^2	0.0001	0.0003	0.001
Panel B			
Minimum Wage, $\underline{w}$	−0.004***	−0.005**	−0.009***
	(0.001)	(0.002)	(0.002)
Constant	0.349***	0.357***	0.377***
	(0.001)	(0.003)	(0.003)
Observations	159,656	39,620	34,157
R^2	0.0001	0.0002	0.001

Notes: The dependent variable is whether the employer hired anyone and paid them some amount of money.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than $5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 2. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 5: Effects of the minimum wage on log hours-worked, conditional upon a hire

	Dependent variable:		
	log hours-worked, conditional upon a hire		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	-0.063^{*}	-0.208^{***}	-0.283^{***}
	(0.027)	(0.056)	(0.061)
MW3	-0.089^{***}	-0.257^{***}	-0.263^{***}
	(0.027)	(0.055)	(0.058)
MW2	-0.046	-0.168^{**}	-0.116^{*}
	(0.027)	(0.055)	(0.058)
Constant	2.783^{***}	3.270^{***}	3.322^{***}
	(0.008)	(0.016)	(0.017)
Observations	47,434	12,894	11,685
R^2	0.0004	0.003	0.003
Panel B			
Minimum Wage, $\underline{w}$	-0.021^{***}	-0.066^{***}	-0.074^{***}
	(0.005)	(0.011)	(0.012)
Constant	2.782^{***}	3.267^{***}	3.322^{***}
	(0.008)	(0.016)	(0.017)
Observations	47,434	12,894	11,685
R^2	0.0003	0.003	0.003

Notes: The dependent variable is the log hours worked, conditional upon a hire.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than $5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 3. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 6: Effects of the minimum wage on log mean wage, conditional upon a hire

	Dependent variable:		
	log mean wage, conditional upon a hire		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	0.141***	0.371***	0.431***
	(0.015)	(0.026)	(0.026)
MW3	0.076***	0.249***	0.253***
	(0.015)	(0.026)	(0.025)
MW2	0.065***	0.148***	0.194***
	(0.015)	(0.026)	(0.025)
Constant	2.107***	1.372***	1.235***
	(0.004)	(0.008)	(0.007)
Observations	53,032	14,143	12,746
R^2	0.002	0.020	0.030
Panel B			
Minimum Wage, $\underline{w}$	0.032***	0.087***	0.099***
	(0.003)	(0.005)	(0.005)
Constant	2.107***	1.371***	1.234***
	(0.004)	(0.008)	(0.007)
Observations	53,032	14,143	12,746
R^2	0.002	0.020	0.030

Notes: The dependent variable is the log mean wage paid, conditional upon a hire.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than $5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 4. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 7: Effects of the minimum wage on the log past wage of the hired worker

	Dependent variable:		
	the log past wage of the hired worker		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	0.041*	0.131***	0.153***
	(0.018)	(0.032)	(0.033)
MW3	−0.015	0.074*	0.076*
	(0.018)	(0.031)	(0.031)
MW2	0.034	0.073*	0.055
	(0.018)	(0.032)	(0.032)
Constant	2.224***	1.362***	1.285***
	(0.005)	(0.009)	(0.010)
Observations	46,030	12,487	11,180
R^2	0.0002	0.002	0.002
Panel B			
Minimum Wage, $\underline{w}$	0.006	0.031***	0.033***
	(0.004)	(0.006)	(0.006)
Constant	2.224***	1.362***	1.284***
	(0.005)	(0.009)	(0.009)
Observations	46,030	12,487	11,180
R^2	0.0001	0.002	0.002

Notes: The dependent variable is log past wage of the hired worker.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than $5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 5. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$.

Table 8: Effects of the minimum wage on log number of organic applications received

	Dependent variable:		
	log number of organic applications received		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	−0.062	−0.464	−0.513
	(0.137)	(0.293)	(0.324)
MW3	−0.267	0.002	0.056
	(0.138)	(0.293)	(0.321)
MW2	−0.157	−0.227	−0.273
	(0.139)	(0.295)	(0.326)
Constant	16.135***	15.207***	16.815***
	(0.042)	(0.090)	(0.099)
Observations	108,876	25,691	21,874
R^2	0.00004	0.0001	0.0001
Panel B			
Minimum Wage, $\underline{w}$	−0.045	−0.080	−0.084
	(0.028)	(0.059)	(0.065)
Constant	16.128***	15.210***	16.818***
	(0.042)	(0.089)	(0.098)
Observations	108,876	25,691	21,874
R^2	0.00002	0.0001	0.0001

Notes: The dependent variable is the log number of organic applications received.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than $5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 10. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 9: Effects of the minimum wage on a follow-on opening by the employer

	Dependent variable:		
	a follow-on opening by the employer		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	−0.008	−0.013	−0.019
	(0.005)	(0.009)	(0.010)
MW3	−0.003	0.001	−0.009
	(0.005)	(0.009)	(0.010)
MW2	−0.006	0.001	0.001
	(0.005)	(0.009)	(0.010)
Constant	0.421***	0.454***	0.474***
	(0.001)	(0.003)	(0.003)
Observations	159,656	39,620	34,157
R^2	0.00003	0.00005	0.0001
Panel B			
Minimum Wage, $\underline{w}$	−0.002*	−0.002	−0.004
	(0.001)	(0.002)	(0.002)
Constant	0.420***	0.455***	0.474***
	(0.001)	(0.003)	(0.003)
Observations	159,656	39,620	34,157
R^2	0.00003	0.00002	0.0001

Notes: The dependent variable is whether the employer posted a second job opening after his or her first opening.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than $5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 11. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$.

Table 10: Effects of the minimum wage on markup in the wage bid of the hired worker

	Dependent variable:		
	markup in the wage bid of the hired worker		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	0.109***	0.253***	0.284***
	(0.009)	(0.019)	(0.022)
MW3	0.054***	0.138***	0.168***
	(0.009)	(0.019)	(0.021)
MW2	0.032***	0.063***	0.107***
	(0.009)	(0.019)	(0.021)
Constant	−0.053***	−0.134***	−0.171***
	(0.003)	(0.006)	(0.006)
Observations	38,922	11,531	10,288
R^2	0.004	0.018	0.022
Panel B			
Minimum Wage, $\underline{w}$	0.023***	0.054***	0.064***
	(0.002)	(0.004)	(0.004)
Constant	−0.054***	−0.137***	−0.172***
	(0.003)	(0.006)	(0.006)
Observations	38,922	11,531	10,288
R^2	0.004	0.017	0.021

Notes: The dependent variable is the markup in the wage bid of the hired worker.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than \$5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 15. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$.

Table 11: Effects of the minimum wage on log past earnings of the hired worker

	Dependent variable:		
	log past earnings of the hired worker		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	0.100**	0.261***	0.273***
	(0.035)	(0.071)	(0.077)
MW3	−0.011	0.171*	0.182*
	(0.034)	(0.069)	(0.074)
MW2	0.048	0.201**	0.215**
	(0.035)	(0.071)	(0.076)
Constant	7.806***	7.483***	7.341***
	(0.011)	(0.021)	(0.023)
Observations	50,322	13,016	11,679
R^2	0.0002	0.002	0.002
Panel B			
Minimum Wage, $\underline{w}$	0.016*	0.067***	0.070***
	(0.007)	(0.014)	(0.015)
Constant	7.805***	7.485***	7.344***
	(0.010)	(0.021)	(0.022)
Observations	50,322	13,016	11,679
R^2	0.0001	0.002	0.002

Notes: The dependent variable is the log past earnings of the hired worker.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than \$5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 17. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$.

Table 12: Effects of the minimum wage on log profile rate of the hired worker

	Dependent variable:		
	log profile rate of the hired worker		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	0.028	0.123***	0.141***
	(0.015)	(0.027)	(0.026)
MW3	0.008	0.096***	0.065*
	(0.015)	(0.026)	(0.026)
MW2	0.032*	0.086**	0.077**
	(0.015)	(0.027)	(0.026)
Constant	2.176***	1.534***	1.430***
	(0.005)	(0.008)	(0.008)
Observations	47,289	12,869	11,506
R^2	0.0002	0.003	0.003
Panel B			
Minimum Wage, $\underline{w}$	0.007*	0.032***	0.031***
	(0.003)	(0.005)	(0.005)
Constant	2.176***	1.535***	1.430***
	(0.004)	(0.008)	(0.008)
Observations	47,289	12,869	11,506
R^2	0.0001	0.003	0.003

Notes: The dependent variable is the log profile rate of the hired worker.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than $5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 16. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 13: Effects of the minimum wage on total wage bill

	Dependent variable:		
	total wage bill		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	−0.541	10.425	−5.926
	(6.449)	(9.068)	(9.423)
MW3	−14.638*	−10.396	−21.767*
	(6.472)	(9.093)	(9.385)
MW2	−2.272	−6.156	1.904
	(6.534)	(9.207)	(9.511)
Constant	148.582***	121.081***	123.283***
	(1.990)	(2.780)	(2.885)
Observations	159,656	39,620	34,157
R^2	0.00003	0.0001	0.0002
Panel B			
Minimum Wage, $\underline{w}$	−1.685	0.131	−2.979
	(1.295)	(1.819)	(1.886)
Constant	148.402***	120.523***	123.342***
	(1.969)	(2.752)	(2.855)
Observations	159,656	39,620	34,157
R^2	0.00001	0.00000	0.0001

Notes: The dependent variable is the total wage bill.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than $5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 12. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$.

Table 14: Effects of the minimum wage on log total wage bill

	Dependent variable:		
	log total wage bill		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	0.075**	0.154**	0.131*
	(0.027)	(0.056)	(0.062)
MW3	0.003	−0.006	0.016
	(0.027)	(0.055)	(0.060)
MW2	0.018	0.005	0.047
	(0.027)	(0.056)	(0.060)
Constant	4.802***	4.584***	4.491***
	(0.008)	(0.017)	(0.018)
Observations	55,262	13,996	12,660
R^2	0.0001	0.001	0.0004
Panel B			
Minimum Wage, $\underline{w}$	0.012*	0.022	0.023
	(0.005)	(0.011)	(0.012)
Constant	4.801***	4.581***	4.490***
	(0.008)	(0.016)	(0.018)
Observations	55,262	13,996	12,660
R^2	0.0001	0.0003	0.0003

Notes: The dependent variable is the log total wage bill.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than \$5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 13. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$.

Table 15: Effects of the minimum wage on an indicator for any past on-platform experience

	Dependent variable:		
	an indicator for any past on-platform experience		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	0.002	0.012	0.010
	(0.004)	(0.009)	(0.010)
MW3	0.004	0.015	0.015
	(0.004)	(0.008)	(0.009)
MW2	−0.007	0.001	−0.001
	(0.004)	(0.009)	(0.009)
Constant	0.917***	0.918***	0.914***
	(0.001)	(0.003)	(0.003)
Observations	54,865	14,149	12,749
R^2	0.0001	0.0003	0.0003
Panel B			
Minimum Wage, $\underline{w}$	0.0004	0.003	0.003
	(0.001)	(0.002)	(0.002)
Constant	0.917***	0.918***	0.914***
	(0.001)	(0.003)	(0.003)
Observations	54,865	14,149	12,749
R^2	0.00000	0.0003	0.0002

Notes: The dependent variable is whether the hired worker had any past on-platform experience.In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than \$5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 18. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : * * *$.

Table 16: Effects of the minimum wage on hours-worked

	Dependent variable:		
	hours-worked		
	ALL	ADMIN	LPW
	(1)	(2)	(3)
Panel A			
MW4	-2.791^{***}	-8.758^{***}	-12.534^{***}
	(0.702)	(2.192)	(2.592)
MW3	-3.152^{***}	-8.650^{***}	-12.141^{***}
	(0.704)	(2.198)	(2.581)
MW2	-2.164^{**}	-6.965^{**}	-6.187^{*}
	(0.711)	(2.226)	(2.616)
Constant	16.293^{***}	29.990^{***}	33.943^{***}
	(0.216)	(0.672)	(0.793)
Observations	159,656	39,620	34,157
R^2	0.0002	0.001	0.001
Panel B			
Minimum Wage, $\underline{w}$	-0.847^{***}	-2.544^{***}	-3.409^{***}
	(0.141)	(0.440)	(0.519)
Constant	16.256^{***}	29.878^{***}	33.926^{***}
	(0.214)	(0.665)	(0.785)
Observations	159,656	39,620	34,157
R^2	0.0002	0.001	0.001

Notes: The dependent variable is the hours worked (including 0)In Panel A, the independent variables are indicators for each experimental group, with the control group excluded; in Panel B, the independent variable is the imposed minimum wage (with 0 for the control). ADMIN are job openings posted in the administrative category. LPW are job openings predicted to have hourly wages less than \$5/hour based on a model fit with historical data. A plot of the data in this table can be found in Figure 14. Significance indicators: $p \leq 0.05 : *$, $.p \leq 0.01 : **$, and $p \leq .001 : ***$.