

Online Appendix for:
“Price Floors and Employer Preferences:
Evidence from a Minimum Wage Experiment”

By *John Horton*

January 2025

A1. Randomization

Table A1 shows the means for a host of pre-treatment job opening outcomes for both the control and MW4. We can see that these 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 characters of text, and “prior spend” is the cumulative amount of money paid by the buyer on wages prior to the experiment.

A2. Workers sorting across openings

One test of sorting is whether applicant counts differ by experimental cell. Figure A1 shows the effects on the number of organic job applications per opening, by treatment cell. The sample is restricted to job openings that received at least one application.

In the population, the counts are very slightly negative. In the sub-populations where would expect larger results, the estimates are less precise rather than larger, with the MW2 cell in LowPREDWAGE at zero. This lack of a “dose-response” relationship suggests any difference in application counts is likely just due to sampling variation. It seems that workers neither avoided nor sought out job openings with imposed minimum wages.

A3. Firms seeking to avoid the minimum wage on the platform

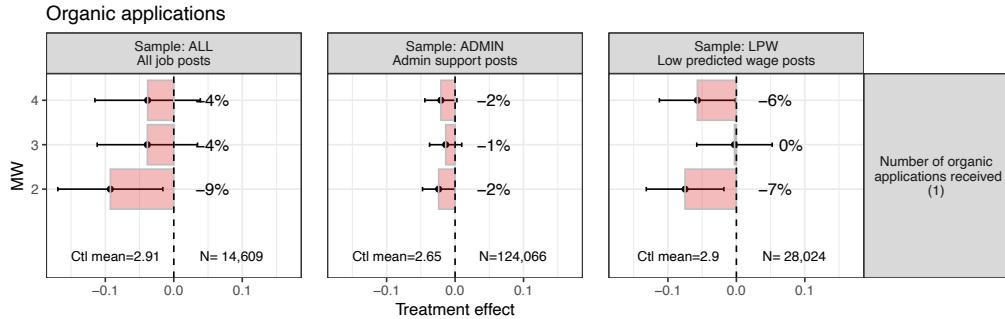
An internal validity concern is that employers might post additional jobs to avoid the minimum wage (which would not actually be effective, given the design of the experiment). This re-posting hypothesis would tend to overstate the extensive margin reductions in hiring. Given that these effects on the extensive margin

Table A1—: 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
<i>Observation Counts</i>				
	9,725	91,781		
<i>Type of work</i>				
Technical (1 if yes, 0 otherwise)	0.426 (0.005)	0.422 (0.002)	0.004 (0.005)	0.471
<i>Type of work—(more detailed)</i>				
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
<i>Vacancy attributes</i>				
New employer?	0.784 (0.004)	0.782 (0.001)	0.002 (0.004)	0.690
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.

Figure A1. : Effects of minimum wage treatment on the number of organic applications per job opening

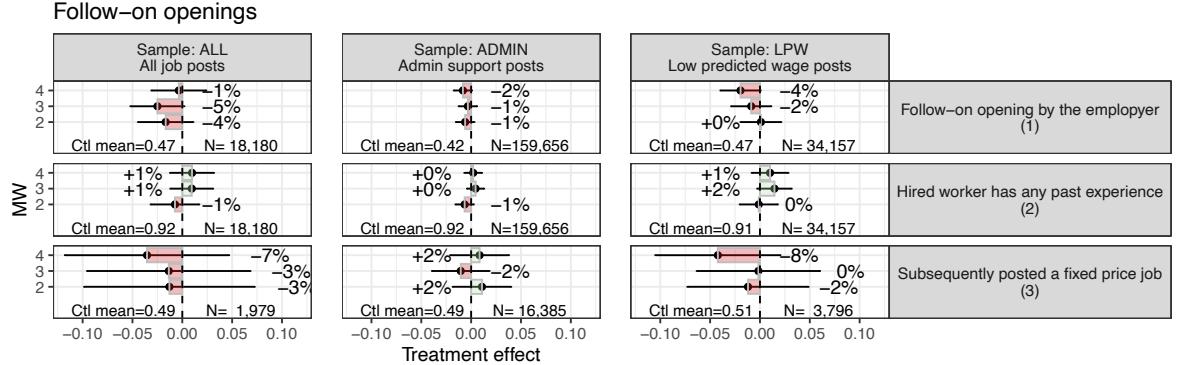


Notes: This figure reports estimates of Equation 1, where the outcome is the log number of organic applications the job opening received. See Section II for details on the experimental design and sample definitions.

are already small, there is not much “room” for this kind of adjustment by firms. There is also no evidence for this phenomenon: Figure A2 shows the effects on whether the employer posted another job within the period covered by the sample. The effects on follow-on openings are negative. If firms were re-posting because they thought they received an idiosyncratically bad draw, this effect should be

positive. Furthermore, the number posting fixed price jobs—which would not be subject to a minimum wage—does not increase, which also cuts against the notion of alerted employers trying to avoid the minimum wage.

Figure A2. : Effects of a minimum wage of whether the employer posted another opening



Notes: This figure reports estimates of Equation 1 where the outcome is whether the employer posted another job opening.

In the population, each active treatment cell has a negative coefficient, though none are significant. In the sub-populations ADMIN and LowPREDWAGE, 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 is not consistent with large numbers of employers thinking they received a bad “draw” of applicants or could avoid the policy by re-posting. It is consistent with them simply believing that prices were higher and thus posting more jobs was less attractive (previewing the finding that jobs posted in ADMIN declined post-imposition).¹⁵

A4. 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 I 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

¹⁵Although one might be tempted to perform a difference-in-differences analysis with the first and second jobs posts to assess the effects of exposure to a minimum wage, there are concerns about this being a selected sample (even if the counts of second jobs by treatment are discernibly different). The kinds of employers that post a second job after exposure to a minimum wage might be quite different. Furthermore, the actual platform-wide-imposition does not raise this concern and so I focus my attention on the platform-wide rollout to assess the longer-term effects of the minimum wage.

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 on an alternative online labor market whose collection of job titles is available. The resultant fitted model is

$$(A1) \quad \text{Pr}(\text{Title Match on Alt. Platform}) = \underbrace{0.0034}_{0.0069} \cdot \mathbf{1}\{\underline{w} = 4\} + \underbrace{0.155}_{0.0018}$$

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.

A5. What did employers know?

Changed beliefs are a nuisance for our purposes, as they would not happen with a market-wide imposition of a minimum wage. Although we do not know what employers actually believed, we can explore if their behaviors are consistent with either alerted or persuaded beliefs.

This could be because the higher bids accurately reflect the cost of hiring a worker for the job, or because the employer was made aware they were part of an experiment or believed the worker to be more productive than they otherwise would have. These changes in beliefs may have affected the employer’s behavior, such as posting another job to avoid their treatment group or waiting to post another job after the experiment ended.

It is difficult to determine whether beliefs changed and if behaviors changed as a result, but the number of “alerted” employers may have been reduced by the design of the platform interface. During the post-experiment phase, there was no evidence that employers tried to quickly fill jobs before the minimum wage deadline, and market outcomes only changed once the minimum wage was implemented. The likelihood of an employer being “alerted” or “persuaded” may depend on their experience with the platform, with experienced employers being more likely to be “alerted” and new employers being more likely to be “persuaded.”

Some treated employers with received atypically high wage bids for their jobs. Instead of simply taking these wage bids as given, some employers might have been (a) “alerted” they were in an experiment or (b) “persuaded” that some applying workers were more productive because they were proposing higher wages. Changed beliefs are a nuisance for our purposes, as they would not happen with a market-wide imposition of a minimum wage. Although we do not know what employers actually believed, we can explore if their behaviors are consistent with either alerted or persuaded beliefs.

A “persuaded” employer might infer a worker bidding more is more productive.

This kind of inference is consistent with [Wolinsky \(1983\)](#) in which price reveals quality, despite quality being imperfectly observed by some buyers. If we adapt this logic to the experiment, some workers are forced to bid more than is rational for them, but they do not face the full lost-business downside of this bidding up because employers (incorrectly) update their beliefs about the bidding worker's productivity. In contrast, an "alerted" employer knows there is some external reason causing the higher bids, perhaps believing they are in an experiment. In response, they might post another job hoping to avoid their cell, or wait to post another job after the experiment was over.

The number of "alerted" employers was likely reduced by the design of the platform interface. At the time the experiment was run, applicants were not ranked by wage bids, but rather by arrival time. Wages bids were also simply listed next to an applicant, rather than visualized in some way that would make a spike at \$2, \$3 or \$4 salient. Such a spike in the distribution of wage bids might not be noticeable even if visualized, as applicants were instructed to continue to bid up until they reached the floor rather than simply being told the floor. Applicants were paginated, with just 10 applicants showing at a time, making it difficult for the employer to make an inference about the change in the pool at a glance.

Inconsistent with many being "alerted," treated employers did not try to avoid the minimum wage by posting another job, posting another job but switching to a fixed price job, or posting another job on alternative platform (Appendix [A.A3](#) and Appendix [A.A4](#)). Furthermore, during the post-experiment phase, when the upcoming minimum wage was widely announced, there is no evidence that employers tried to quickly fill jobs before the deadline (discussed at length in Section [V](#)). Market outcomes only changed once the minimum wage was actually implemented.

We might expect that whether an employer had "alerted" or "persuaded" beliefs would in turn depend on their experience with the platform. An experienced employer would be more likely to be "alerted" while a new employer unfamiliar with rates would be like to "persuaded." If an experienced employer had more accurate beliefs about worker productivity, then the [Wolinsky \(1983\)](#) logic suggests they would be less likely to hire over-priced applicants unlike more credulous rookies. Despite the plausibility of this argument, I find no evidence that the treatment effects on hiring depended on whether the employer had prior experience on the platform (Appendix [B.B5](#)). Furthermore, there is no evidence that the labor-labor substitution results depend on employer experience, which is a direct test of the [Wolinsky \(1983\)](#)/persuaded conjecture.

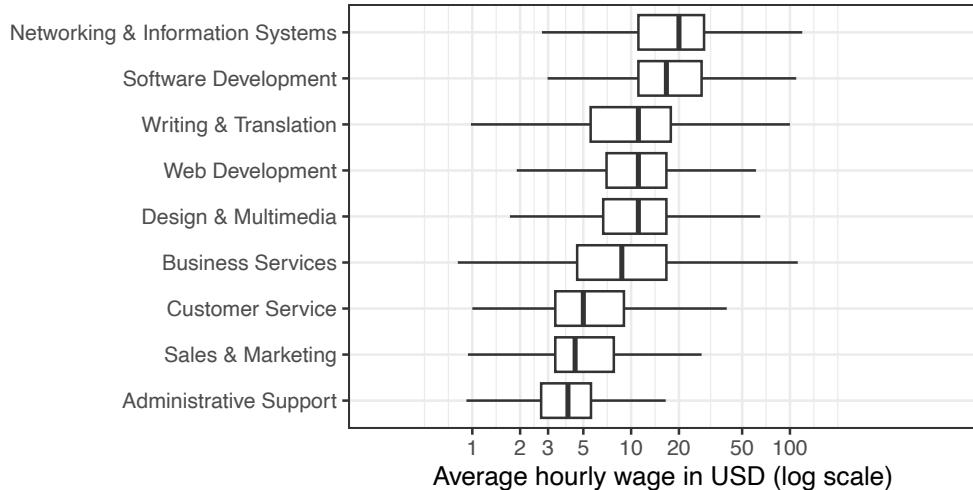
When the minimum wage was imposed platform-wide, everyone had "alerted" beliefs and no one would have "persuaded" beliefs. Despite this very different informational environment, the results on the composition shifts in hiring are not just directionally the same, but also similar in magnitude. We also see the same amount of labor-labor substitution, which is contrary to the notion that

many employers were “persuaded” during the experimental phase. This might be surprising, but consider that bids themselves might not be very informative. Anyone is free to propose any number as wage bid and the only cost to doing so is the cost of making an offer, which is close to zero, and there are many reasons why workers might submit uncompetitive bids. As such, even if we think the wage offers are credible as a willingness to sell—the worker would really work at that price—there is little reason to think it reflects a credible willingness to buy, which is what would be informative about productivity.

A6. Constructing LPW group

Figure A3 shows boxplots for the log wages for each on-platform category of work in the control group.¹⁶

Figure A3. : Wages by category of work in the control group



Notes: This figure shows the distribution of hourly wages 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.

The training data was 100,000 pre-experiment job openings in which a hire was made. The outcome was the log hourly wage for the hired worker. The candidate predictors included the category of work, skills required, the anticipated duration, and the job opening title.¹⁷ To estimate the model, I used the *glmnet* package

¹⁶The sample is restricted to wages above 25 cents per hour in 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.

¹⁷For textual predictors, I used the *RTextTools* package, developed by [Jurka et al. \(2012\)](#) to create a

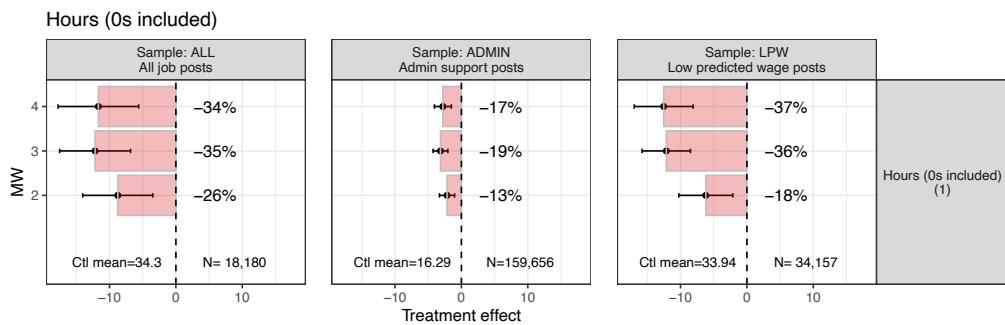
developed by Friedman et al. (2009), using LASSO for regularization and variable selection (Tibshirani, 1996). The optimal tuning parameters were selected via cross validation. Using the fitted model, I made predictions for every job opening in the experiment, and then selected those predicted to pay less than \$5/hour.

ADDITIONAL EXPERIMENTAL OUTCOMES

B1. Hours-worked (0s included)

Figure B1 shows the effect of the intervention on hours-worked, with 0s included for unfilled jobs. It shows the same large decrease in hours-worked found when conditioning on hiring.

Figure B1. : Effect on hours-worked, with 0s included for unfilled jobs



Notes: This figure reports estimates of Equation 1 for hours-worked, with 0s included for job openings where a hire was not made.

B2. Hours-worked quantile regression

Perhaps hours-worked fell because jobs that would have many hours went unfilled. However, 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 8% decrease in hours-worked, making a pure selection explanation implausible. Even if we take the 2% MW2 reduction in ADMIN, if we assume worse-case missingness for ADMIN (i.e., the 2% unfilled would be the 6% that would have had the most hours), the log point reduction in hours-worked would be 0.08 log points, not the 0.28 we observe). Furthermore, a quantile regressions of hours-worked show that reductions occurred throughout the distribution of hours-worked and that the difference between treatment and control was not simply that treated jobs were “missing” jobs with many hours.

document term matrix.

Table B1 reports a collection of quantile regressions of hours-worked on the treatment indicators. It shows that hours-worked decreased at points throughout the distribution.

Table B1—: Quantile regressions of hours-worked (0s included) in ADMIN category

	Dependent variable:				
	Hours-worked				
	99th	95th	90th	80th	70th
	(1)	(2)	(3)	(4)	(5)
MW4	−65.611 (15.317)	−12.792 (2.884)	−7.083 (1.270)	−2.278 (0.397)	−0.833 (0.155)
MW3	583.000 (29.491)	157.667 (6.224)	72.500 (2.406)	21.833 (0.821)	7.500 (0.329)
Observations	18,180	18,180	18,180	18,180	18,180

Notes: The outcome in this table is hours-worked, in levels. Each column is quantile regression, with the quantile labeled.

B3. Any past experience of the hired worker

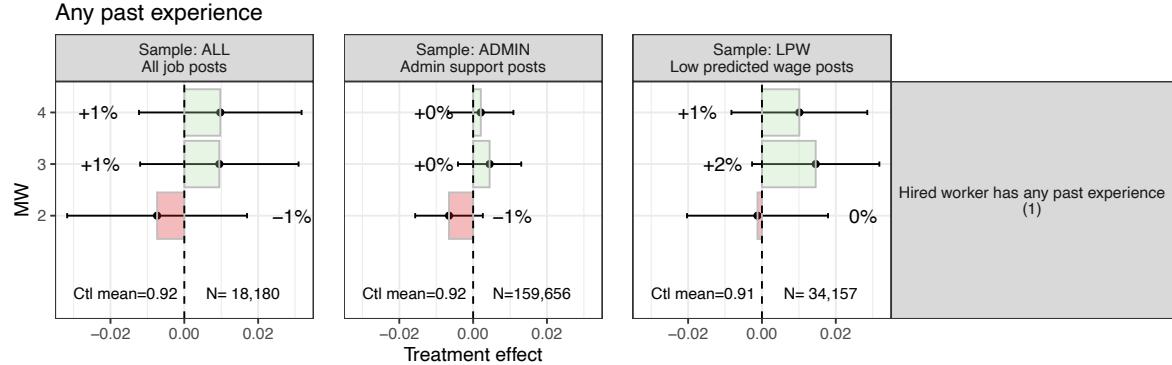
The outcome variable in the analysis shown in Figure B2 is an indicator for whether the hired worker had any on-platform work experience at the time they were hired. There is no strong evidence that the treatment affected the probability that the hired worker had prior experience.

B4. Subjective evaluations post-contract

When a firm or a worker ends a contract, both sides are asked to rate the other side on a five-star scale. Although five stars are, by far, the most common rating, there is some variation in ratings. Generally, worker ratings of firms are more favorable than firm ratings of workers. It is beyond the scope of this paper to model what this feedback actually means—particularly since the given feedback clearly has some strategic implications—it is at least plausible that it proxies for a party’s surplus from a relationship.

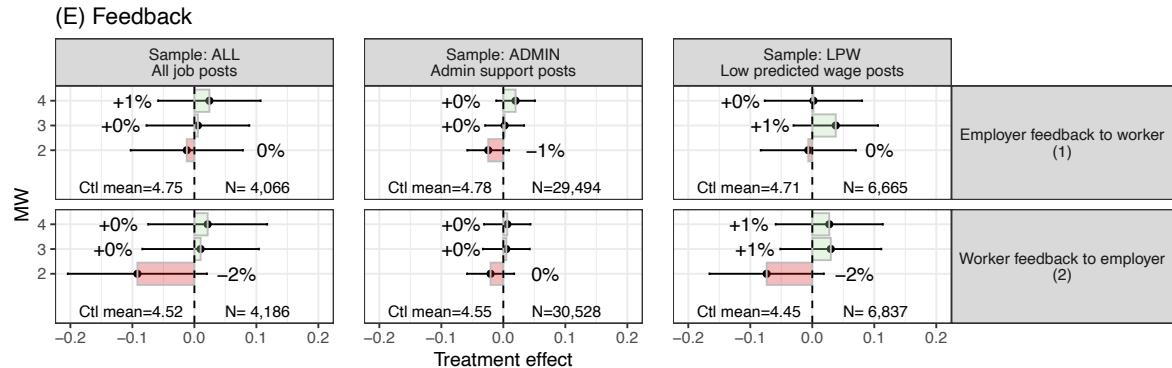
The top panel of Figure B3 shows the difference in average feedback by firms in the active treatment cells compared to those in the control group. The bottom panel of Figure B3 shows the difference in average feedback by workers in the active treatment cells compared to those in the control group. There is no strong evidence of any difference in average feedback of either type.

Figure B2. : Effects of imposing job-specific minimum wages on the probability the hired worker had any on-platform experience



Notes: This figure reports estimates of Equation 1 where the outcome is an indicator for whether the hired worker had any on-platform experience.

Figure B3. : Effects of imposing job-specific minimum wages on job outcomes relative to the control group



Notes: This figure reports estimates of Equation 1 for the two feedback measures: a worker on the employer and employer on the worker.

B5. Treatment effects on hiring by employer experience

A reasonable conjecture is that employers with on-platform experience in treatment cells might have made different inferences about the wage increases than those employers new to the platform. However, there is no evidence this is the case: employer prior experience is positively correlated with the probability of hiring, but there is no evidence that treatment effects differ by employer experience. Nor is there a detectable difference in the experience level of hired workers,

by experience level.

Table B2 reports these results. In Columns (1) and (2), the outcome is whether the employer hired. In Column (1), the regression is an estimate of Equation 1 but augmented with the employer experience. In Column (2), employer experience is interacted with the treatment indicators. In Columns (3) and (4), the outcome is the log cumulative earnings of the hired worker.

Table B2—: Effects of the employer prior experience

	Dependent variable:			
	Anyone hired?		Hired worker past earnings	
	(1)	(2)	(3)	(4)
MW4	-0.007 (0.003)	-0.006 (0.003)	0.074 (0.024)	0.062 (0.033)
MW3	0.246 (0.007)	0.248 (0.008)	-0.167 (0.064)	-0.185 (0.072)
MW2		-0.002 (0.005)		0.026 (0.048)
Prior Experience	0.356 (0.005)	0.356 (0.005)	6.811 (0.046)	6.820 (0.049)
Observations	18,180	18,180	8,027	8,027
R ²	0.057	0.057	0.002	0.002

Notes: The outcome in Columns (1) and (2) are whether the employer hired anyone at all. In Columns (3) and (4), it is the cumulative prior earnings of the hired worker, if any. In addition to the treatment cell indicators, one of the regressors is a whether the employer had any on-platform experience (as measured by having hired in the past).

The confidence interval for the coefficients on every employer's prior experience/treatment cell interaction term comfortably includes zero. Furthermore, when I perform an LR test, with the simpler model nested within the model with interactions. I fail to reject the null hypothesis in both cases, and the log-likelihoods for the two models are nearly identical. This result is not driven by experienced employers simply being rare or exceedingly common—about 34% of employers in the experiment had hired workers on the platform in the past. If there were big differences in treatment effects by experience, we would likely detect them.

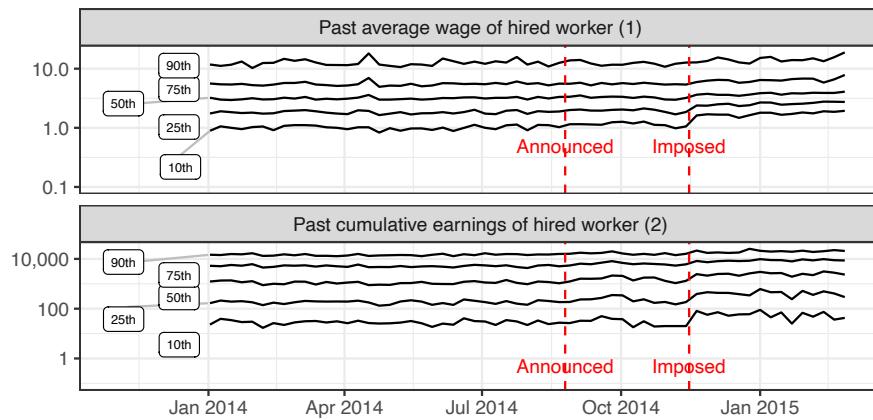
B6. Characteristics of hired workers in ADMIN post-imposition as event study

Recall that in the experiment, there was a strong shift towards hiring more experienced workers, but that this might not be borne out in equilibrium. To detect whether this shift also occurred following the platform-wide roll-out, Figure B4 plots the attributes of workers hired in ADMIN over time. The announcement and imposition are indicated by dashed vertical lines.

In the top row, labeled “Past average wage of hired worker (1),” the times series are the quantiles of the hired worker’s past average wage that week. Importantly, this average past wage is computed only using wages earned prior to the announcement of the minimum wage. There is no visual evidence of a change at any quantile post-announcement. However, post-imposition, there is a clear increase in the average past wages of the hired workers, at multiple quantiles. This pattern of more experienced workers being hired is repeated in the bottom panel for cumulative earnings, labeled “Past cumulative earnings of hired worker (2).”

Figure B4 suggests that the substitution towards more productive workers that occurred during the experiment also occurred post-imposition. However, a direct comparison of the observational and experimental magnitudes requires a different approach, which is the focus of the next section.

Figure B4. : Time-series of the attributes of hired workers, in ADMIN only



Notes: This figure plots weekly quantiles for the attributes of hired workers in the ADMIN category. The minimum wage policy announcement and imposition are indicated with vertical dashed lines.