

Dissertation Awards

10-26-2021

The Fruits of Our Labor: Essays on Work and Its Impacts

Natalia Emanuel

Harvard University, nemanuel@princeton.edu

2021

W.E. UPJOHN INSTITUTE DISSERTATION AWARD SUMMARIES

WINNER

Anna Stansbury
Harvard University

HONORABLE MENTIONS

Zachary Bleemer
University of California, Berkeley

Natalia Emanuel
Harvard University

The Fruits of Our Labor: Essays on Work and Its Impacts

Natalia Emanuel

These essays explore the modern workplace and how it impacts low-wage workers. The research focuses on the various ways low-wage work interacts with constraints and obligations inside and outside of work, including the effects of wage-setting decisions on productivity and retention among warehouse workers and customer service representatives, reasons underpinning the gender-pay gap among unionized bus and train operators, whether work obligations may be at fault for failure to appear at court dates, and why firms may be loath to offer remote work possibilities despite their effect on worker productivity. My work utilizes a number of different methods, including natural experiments, randomized controlled trials, and surveys. Each paper has leveraged data specific to the firm or organization I study.

Together, these essays aim to understand various different forces that are at play when firms are making key personnel decisions—such as how much to pay and the location of work—and when workers are optimizing, given the decisions that firms have made.

Essay 1

The Payoffs of Higher Pay: Elasticities of Productivity and Labor Supply with Respect to Wages

(with Emma Harrington)

Firm wage-setting decisions must balance the benefits to the firm of higher pay—lower turnover, higher worker effort, and enhanced recruitment—against the direct costs of higher compensation. This essay provides new evidence on the returns to the firm of higher pay. We estimate elasticities of productivity, turnover, and recruitment among warehouse workers, and customer service employees at a Fortune 500 retailer. We leverage idiosyncrasies in the firm's pay-setting policies to estimate the response of productivity, turnover, and recruitment to pay using three complementary empirical strategies.

First, we use a large, abrupt jump in pay to look at the difference in turnover and productivity before and after. In one of the warehouses at the firm, the pay went from \$16.20 to \$18.00/hour with no warning. The firm had long been concerned about this warehouse, but due to administrative issues in the compensation department, it had not followed through with adjusting pay. As such, the *timing* of the pay

bump was as good as random. Moreover, no other changes in how the warehouse was run occurred at the same time, and accordingly we see no change in the demographic composition of the warehouse.

Moving in step with the pay increase at the warehouse, productivity jumps and turnover drops among the warehouse workers. Workers initially handled about 4.9 boxes per hour, but when their pay increased, they began handling another 0.33 boxes per hour. This would imply that a 1 percent increase in pay is associated with about a 1 percent increase in productivity. The increase in this productivity is both because they are spending more of their time in the warehouse actually moving boxes, and also because they are moving boxes more swiftly. Likewise, turnover drops around the pay change. Beforehand, over 13 percent of the warehouse's workers left each month. After the pay change, the turnover decreased by 2.5 percentage points. A 1 percent increase in pay in this warehouse was associated with a 3 percent decrease in turnover.¹

We then consider the same two responses to pay—productivity and turnover—among another group of workers using a different empirical approach, finding similar results. Nationally, the firm has sticky wages for their customer service workers. Over time, the constant pay leads to exogenous variation in the value of wages relative to workers' local outside options. By comparing changes in relative pay in various cities over the course of the year to the changes in the productivity and turnover in those cities, we are able to estimate the effect of changes in relative wages on workers' behavior. The key assumptions are that the wages at our firm have negligible effects on the outside option. We are comfortable with this assumption because in any of the studied cities, this firm hires less than 0.5 percent of the people employed in customer service who work in that city. Additionally, we assume that other local changes that would affect productivity are orthogonal to changes in the wage. Our estimate would be biased toward zero if there were measurement error in wage growth, which would lead to attenuation bias if wage growth caused by a local productivity shock or inflows of more productive workers increased productivity in high wage-growth areas.

Using the sticky-pay empirical strategy, we find that when customer service workers receive pay that is \$1 more than their outside option, their productivity increases and they are less likely to leave the firm, just as we had found with the warehouse workers. Customer service workers on average handled 25 calls per day but handled 1.9 calls, or 7.5 percent more, when paid more. So, a 1 percent increase in pay is associated with a 1.2 percent increase in productivity. Interestingly, this does not come at the expense of customer satisfaction reviews: there does not appear to be a quantity-quality trade-off. Rather, workers seem to spend less time on paperwork and a few more minutes per day working. Turnover among customer service workers decreases when

pay is higher than the outside option. Among customer service workers at this firm, 4.5 percent leave per month. But where pay is higher than the local outside option, there is a 1.2 percentage point reduction in turnover. This implies that a 1 percent increase in pay is associated with a 4.5 percent reduction in turnover.

Recruitment also responds to pay. To estimate how, we use the fact that the firm sets wages nationally for its remote customer service workers. Regardless of location, the firm offers a wage of \$14/hour, which creates heterogeneity in wages relative to the outside option. We estimate how the pay relative to workers' outside options relates to the number of workers recruited, controlling for local controls, including the local unemployment rate and the number of people working on that occupation locally. The key assumption we make is that the covariance of the relative wage with features of the locality other than the pool of available workers and unemployment rate does not also correlate with differential recruitment. Since recruitment is done nationally and the firm is indifferent about the location of its remote customer service workers, we believe this to be a safe assumption. We find that for each dollar of pay, recruitment numbers increase by nearly 23 percent. Since we are measuring the number of people hired, we are capturing a measure of how many *qualified* candidates a dollar recruits, which is the relevant metric for a firm.

Finally, we estimate gender-specific responses to higher pay to understand what our model suggests about occupation-level gender pay gaps. Male customer service workers at our firm are almost 10 percent more likely to leave their jobs than are women. The gender difference in turnover is consistent with women facing smaller pools of outside options; for example, due to less willingness to commute. Moreover, men adjust their likelihood of leaving the job more in response to pay than do women: when paid a dollar more, men reduce their turnover by 40 percent. The female labor supply response is not statistically significant. While we have no evidence about the source of this difference, it is consistent with men doing more on-the-job-search than women. Assuming a constant production function and that there are no other gender differences, the difference in the labor supply responses would suggest that for every dollar a male worker earns, a female worker would earn \$0.94.

In contrast, female customer service workers at our firm have *larger* productivity responses to pay. When paid a dollar more relative to the outside option, female customer service workers increase the number of calls handled by 9 percent relative to the male increase of 5 percent. Given the firm's cost per call, the female increase in productivity represents a savings of \$1.91 per hour for a dollar increase in wages. For male workers, the firm just breaks even. That women increase their productivity more in the face of higher pay than do men suggests a force that pushes in

the opposite direction of the existing gender pay gap. The different productivity response is consistent with a number of potential explanations, none of which we can identify from our administrative data,² including women having more firm-loyalty, believing that they have fewer outside options, and believing that they face discrimination.

As a whole, this essay aims to understand the incentives firms face when setting pay in regard to productivity, recruitment, and turnover. However, the essay leaves many unanswered questions about the gender pay gap. Sources of the within-occupation gender pay gap are addressed in the next paper in my dissertation.

Essay 2

Why Do Women Earn Less Than Men? Evidence from Bus and Train Operators

(with Val Bolotnyy)

At the Massachusetts Bay Transit Authority (MBTA), the unionized organization that runs Boston-area buses and trains, there exists a gender pay gap: for each dollar that the typical male employee takes home, the typical female employee earns \$0.89.

The existence of a gender pay gap in this workplace is curious since many of the traditional explanations for the gender pay gap are moot in this context. The union-negotiated contract specifies pay rates and tenure-based promotion schedules with the explicit intent of eliminating any managerial bias in these domains. The structures of the contract also mean that any gender-based differences in workers' willingness to compete or negotiate are rendered irrelevant in this workplace. Finally, since we are looking only at one occupation, we are not picking up on occupational sorting, one prominent explanation for economy-wide gender pay gaps.

In this essay we explore the roots of the gender pay gap among unionized bus and train operators. We leverage the fact that seniority is the sole determinant of workplace opportunities such as overtime at the MBTA to understand the operators' value of time and other amenities. Conditional on seniority, male and female operators face the same choice sets of schedules, routes, vacation days, and overtime hours, among other amenities. Nevertheless, the earnings gap persists even when we condition on seniority.

Mechanically, the gender pay gap can be explained by the fact that male operators take on 1.5 more overtime hours (83 percent) per week and take off 1.3 fewer unpaid hours off work (49 percent) per week than their female counterparts. This is indicative that female operators value their time outside of work more than do male operators, consistent with women handling more household and child care duties than men.

The difference in male and female overtime is driven by overtime opportunities that arrive on short notice and therefore demand that operators are flexible about when they work. When overtime is scheduled the day before or the day of the necessary shift, male operators work almost twice as many of those hours as female operators. In contrast, when overtime hours are scheduled three months in advance, male operators sign up for only 7 percent more of them than female operators.

Exacerbating the disparity in overtime acceptance rates, male operators strategically substitute regular hours for higher-paying overtime hours using the Family Medical Leave Act (FMLA), a federal statute that allows workers with a doctor's note to take unpaid time off of work to attend to their own medical needs or those of a family member. At the MBTA, FMLA has been nicknamed the "Friday-Monday Leave Act" for the way that operators have used it to avoid undesirable shifts. Both male and female operators take more FMLA hours when faced with undesirable shifts (e.g., a weekend or holiday shift). However, male operators also work enough overtime hours in weeks with an undesirable shift that they effectively trade off hours paid at the regular wage for overtime hours paid at 1.5 times their wage. Female operators also work more overtime hours in weeks with undesirable shifts and therefore more unpaid time off, but do not completely make up the pay lost to FMLA leave. Together, the last-minute overtime and the usage of FMLA mechanically account for the pay gap.

We explore how family arrangements relate to the differences in propensity to accept overtime. The gap between the acceptance of overtime opportunities of male and female operators is larger if the operators have dependents (6.8 percentage points) than if they do not (5.7 percentage points). Male acceptance rates, meanwhile, are similar for the two groups (38.2 percent for male operators with dependents, 41.1 percent for male operators without dependents). Though dependents generate this wedge in acceptance rates among married and unmarried operators, the wedge is largest among married operators. Married men with dependents accept overtime opportunities 27.1 percent of the time, while married women with dependents accept them 19.6 percent of the time. For unmarried men with dependents the acceptance rate is 40.3 percent, compared to 33.6 percent for unmarried women with dependents. These results are consistent with male operators doing more child care through their pocketbooks, and with female operators doing more child care through time spent outside of work. Differences in caretaking approaches and responsibilities thus appear to be a significant reason why female operators work less overtime than male operators.

We next seek to understand more about male and female scheduling preferences, focusing on scheduling conventionality and controllability. While both male and female operators avoid unconventional shifts, female

operators do so more. The most senior operators, who pick their schedules first, have much lower incidence of unconventional shifts—weekend shifts, shifts on holidays, and shifts split into two blocks ("split shifts")—relative to operators who choose their schedules later. While 95 percent of the least senior operators get stuck with a weekend shift on their schedules, only 28 percent (female operators) to 35 percent (male operators) of the most senior operators do. The same pattern holds true for holiday shifts and split shifts. While all operators avoid unconventional schedules, female operators avoid them more readily. Female operators are on average about 2.5 percentage points less likely to select a weekend shift than are male operators.

Female operators value time outside of work and schedule controllability more than do male operators, especially when they have dependents. Female operators with dependents are considerably less likely than male operators with dependents to accept an overtime opportunity. When it comes to overtime hours worked, unmarried female operators with dependents work only 6 percent fewer of them when they are preplanned three months in advance, but about 60 percent fewer of them when they are offered on short notice. Unmarried women with dependents also take the largest amount of unpaid time off with FMLA, making them the lowest earners in our setting.

Lastly, we study the impact of two policy changes at the MBTA, each of which made it harder for operators to swap regular hours for overtime hours. The first policy change made it more difficult for operators to obtain FMLA certification, to use FMLA for anything other than a medical issue, and to take unpaid time off at a moment's notice. The second policy change, a year later, redefined overtime hours from any hours worked in excess of 8 in a given day to any hours worked in excess of 40 in a given week.

These policies simultaneously reduced the gender earnings gap and hurt workers. The gender earnings gap shrank from 12 percent before the FMLA policy change to 9 percent between the first and second policies' implementation and to 6 percent following the second policy's arrival. Yet, in addition to reducing the gap, these policies also reduced schedule controllability. Those who took more unpaid time off via FMLA before the policy changes now took more unexcused leave instead, indicating that these operators still desired control over their schedules. Operators began procuring this control at a higher cost, since unexcused leave can result in suspensions and discharge from work (unlike FMLA leave). Because female workers have greater revealed preference for schedule controllability, these policies—particularly the first—affected female operators more negatively than they did male operators.

Our results suggest two potentially Pareto-improving strategies that could be implemented in this and similar settings. First, if operators are allowed to exchange or transfer shifts, unexpected absenteeism could be reduced.

Reducing absenteeism would decrease both unpaid time off and resultant last-minute overtime opportunities, both of which fuel the earnings gap. Service provision would also improve if absenteeism drops and operator productivity consequently increases. Second, expanding the number of operators whose job is specifically to cover for others' absences would also likely decrease the earnings disparity, overtime expenses, and inconsistent service.

This paper uses the unique setting of unionized bus and train operators to explore how seemingly neutral firm policies can result in unequal outcomes. While the essay focuses on this singular occupation, the lessons learned are broadly applicable. Many occupations have last-minute scheduling—whether it is the server at JambaJuice, which assigns shifts based on the upcoming weather, or the lawyer who is asked to do a dinner tomorrow night with a client. In these occupations, employees who have obligations outside the workforce may have to forego valuable opportunities in the same way as bus and train operators. Similarly, many occupations have relatively little flexibility, prompting workers to try to inject flexibility into their schedules in ways that can be quite costly to them. Since many of the features that prompt the pay gap among bus and train operators show up in other occupations as well, this paper highlights specific sources of the gender pay gap that may be generalizable outside this particular setting.

Essay 3

Behavioral Biases and Legal Compliance

(with Helen Ho)

When considering attendance at court hearings, traffic, municipal, and misdemeanor courts have a problem: as many as 40 percent of defendants fail to appear (FTA). FTAs are costly for the court since it leads to unused personnel time. Statutorily, they are also costly to defendants: they can result in warrants for the defendant's arrest, fines, and fees.

Court officials as well as defendants who did show up believed that work makes attending a court date particularly challenging. Work obligations can make spending four hours in court plus travel time quite costly. Nevertheless, if the hearing occurs while the defendant is not working, the inability to bring children into the courtroom may be a challenge. If they are charged with a large fine or fee, liquidity constraints may make attendance unattractive. To address these time and resource constraints, courts offer some accommodations, including rescheduling and payment plan options. However, accessing these court accommodations is notoriously difficult to people who don't know, for example, that they must file a request for a "one-time continuance" if they seek to reschedule. It is

also possible, however, that defendants are forgetful or disorganized and simply need the arraignment brought to their attention.

To understand the potential causes of an FTA and examine low-cost, scalable ways to address the problem, we collaborated with a large court system to run a randomized controlled trial (RCT). In this trial, we provided defendants with no additional information (the treatment arm), with informational nudge reminders that helped remind defendants about the arraignment and possible court accommodations, or with informational nudges alongside the offer of personalized assistance.

If constraints such as work and child care obligations prevent attendance at court, we would expect that the defendants randomized to receive either informational nudges or personalized assistance would have greater uptake of court accommodations and greater attendance at their arraignments as well. If difficulty navigating the court system were part of the cause, then we would expect to see additional boosts in usage of court accommodations and reductions in FTA rates.

In our RCT, we found that personalized assistance increased uptake of court accommodations such as rescheduling and payment plans. However, both informational nudges and personalized assistance were equally effective in reducing FTA rates, taking them from a baseline of 21 percent FTA down by 39 percent. Moreover, there was no difference between the two arms in terms of the amount paid out by the defendants.

Essay 4

“Working” Remotely? Selection, Treatment, and the Market for Remote Work

(with Emma Harrington)

Prior to the COVID-19 pandemic, only 6 percent of U.S. jobs were remote. This is curious because many jobs are “remoteable,” many workers report wanting to work remotely, and prior studies suggest there is a positive treatment effect of working remotely. During the COVID-19 pandemic, as many as 45 percent of workers reported they would like to remain remote even after the pandemic. Mas and Pallais (2017) found that call-center workers were willing to accept 8 percent lower wages in order to be remote. Moreover, Bloom et al. (2015) found remote work actually increased productivity. Nevertheless, even at what appears to be the tail end of the pandemic, only 9 percent of job postings are for fully remote positions.

The puzzle thus arises: Why is there so little remote work despite the fact that such a work setup seems to enhance worker happiness and productivity? In this essay, the final

paper completed during my PhD,³ we argue that adverse selection is at fault. Worse workers are inclined to accept remote jobs, thereby raising the cost of remote positions to the firm. To map the dynamics around remote work, we leverage data from customer service representatives at a Fortune 500 retail firm.

We develop a model that ties promotion penalties to the selection of workers who choose remote jobs. In the model, remote work reduces the probability that firms learn about workers' abilities. Latently low-ability workers consequently sort into remote jobs to hide their ability while latently high-ability workers sort into on-site jobs to reveal their high ability. The resulting adverse selection into remote work raises its average cost above its marginal cost, causing remote work to be underprovided.

This model fits several indicators of informational frictions present in the firm. Managers appear to be less certain about remote workers' productivity: managers' evaluations are less predictive of the future performance of remote workers than on-site workers. Accordingly, workers at the retailer who chose remote jobs had about half of the promotion chances as those who chose on-site jobs. This is consistent with Bloom et al.'s (2015) finding that remote work negatively impacts promotion.

We test the model's predictions empirically, analyzing productivity differences between workers who are hired into comparable remote and on-site jobs. We use the number of calls the call-center worker handles as the primary metric of productivity, but check also that there is no quantity-quality trade-off by measuring customer satisfaction ratings. Two natural experiments at the retailer allow us to separately identify the treatment and selection effects of remote work.

First, the retailer introduced a program in 2018 that allowed on-site workers to apply to remote call-center jobs. The initiative changed new hires' offers. Considering the group of people who ultimately ended up working remotely, we compare the productivity of those who thought that they would likely always work on-site to those who were offered a job where they knew they would go remote. This comparison captures the selection effect of offering remote work. To estimate the treatment effect of remote work, we can look at the productivity of those workers who started on-site and also worked remotely. For workers who thought that they would always work on-site and were surprised by the opportunity to go remote, this comparison represents only the treatment effect.

Second, COVID-19 caused all of the call-center workers to work remotely. Starting on April 6, 2020, all on-site call-center workers were sent to work from home. Comparing these workers' productivity captures both the treatment effect of remote work as well as the effects of COVID-19. However, since July 2018 the retailer has hired workers directly into remote call-center jobs, creating a control group. Using a difference-in-differences approach, we can

cancel out the common shocks of COVID-19 on worker productivity. Moreover, by comparing the prepandemic productivity of those who were initially on-site to those who were initially remote, we capture the selection effect of remote work.

These two natural experiments give relatively similar estimates of the treatment and selection effects of remote work. In the first, we estimate that being offered remote work reduces productivity by 21 percent. The selection effect as measured in the second natural experiment estimates an 18 percent reduction in productivity. In neither setting does the decreased number of calls handled mean that the remote call-center workers are compensating by achieving higher customer satisfaction. The treatment effect estimated in the first natural experiment marks a 7 percent increase in productivity on account of remote work. The increase in calls handled arises from them answering queries 3 percent faster and spending 4.5 percent more time on the phone rather than waiting, filing paperwork, and performing other adjacent tasks. The treatment effect as measured in the second natural experiment represents a 10 percent increase in productivity. Again, these increases in the calls handled do not come at the expense of customer satisfaction. Together, the two natural experiments provide evidence of a positive treatment effect, but a negative selection effect of remote work.

Finally, we quantify the distortion arising from the adverse selection that we document. We estimate workers' demand for remote work using the retailer's policy of paying all remote workers the same wage nationally. By paying all workers \$14/hour, regardless of location, the firm creates variation in the opportunity cost of taking the remote job because workers' local on-site alternative jobs vary in their wages. Using the estimated demand curve, we find that adverse selection likely reduces the share of call-center workers working remotely from 17 percent to 6 percent nationally, leading to losses of \$824 million annually just among the 3.2 million American call-center workers. Relative to other occupations, call-center jobs tend not to have tremendous promotion ladders; at our retailer, there are only two major promotion levels. Call-center work is also relatively easy to monitor from afar. As such, the losses may be even more substantial in other occupations where career concerns are more acute and where it is more challenging to monitor workers' productivity if they are not on-site.

Our analysis suggests that the pandemic will attenuate but not eliminate adverse selection into remote work. On the worker side, surveys suggest that the retailer's workers have learned more about their tastes during the lockdown, causing more high-ability workers to choose remote jobs. This reduction in the average cost of remote work would increase its prevalence by 1.1 percentage point. On the firm side, we find little evidence that the retailer's experience with remote work reduced the promotion penalty and the consequent incentive for workers to sort on ability.

Having leveraged data from call-center workers, a highly “remoteable” job, our paper unfortunately cannot speak to tasks that require much coordination or intense concentration. For jobs requiring such tasks, prior studies have found less benefit of remote work. As such, provision of remote work for these jobs could arise from both the treatment and selection effects.

Conclusion

Together, the essays completed in my doctoral studies address several issues of work for low-wage workers, including how firms set wages, the various forces undergirding the gender pay gap, how work obligations may impact legal obligations like attending a court arraignment, and why remote work appears to be underprovisioned in the marketplace. The essays use a range of tools—from natural experiments to surveys and RCTs—to address the pressing issues for low-wage workers at the moment.

Notes

1. I have noted the results of an analysis looking at an interrupted time series. But the results are robust to considering a difference-in-differences approach, using the firm’s other warehouses as the control group.
2. The partner firm was unwilling to let us run a survey on this topic.
3. This paper was completed during my PhD, under the advisement of my dissertation committee. It was not included in the dissertation submitted to Harvard because I delivered my baby prematurely and did not have time to format it according to the university’s specifications before the submission deadline.

THE W.E. UPJOHN INSTITUTE, a nonprofit research organization, was established on July 1, 1945. It is an activity of the W.E. Upjohn Unemployment Trustee Corporation, which was formed in 1932 to administer a fund set aside by Dr. W.E. Upjohn for the purpose of conducting “research into the causes and effects of unemployment and measures for the alleviation of unemployment.”

The observations and viewpoints expressed are the sole responsibility of the authors. They do not necessarily represent positions of the W.E. Upjohn Institute for Employment Research.