

DISCUSSION PAPER SERIES

IZA DP No. 14325

**The Return to Hours Worked within and
across Occupations: Implications for the
Gender Wage Gap**

Jeffrey T. Denning
Brian A. Jacob
Lars J. Lefgren
Christian vom Lehn

APRIL 2021

DISCUSSION PAPER SERIES

IZA DP No. 14325

The Return to Hours Worked within and across Occupations: Implications for the Gender Wage Gap

Jeffrey T. Denning

Brigham Young University, NBER and IZA

Brian A. Jacob

University of Michigan and NBER

Lars J. Lefgren

Brigham Young University, NBER and IZA

Christian vom Lehn

Brigham Young University and IZA

APRIL 2021

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

The Return to Hours Worked within and across Occupations: Implications for the Gender Wage Gap*

Prior research suggests that gender differences in hours worked play an important role in the gender pay gap. Yet common estimates of the wage returns to hours worked are close to zero, implying that hours differences cannot account much for the gender wage gap, even though men work more hours than women on average. However, while the wage returns to hours worked *within* occupation are small, we document that the wage returns to average hours worked *across* occupations are large. We develop a conceptual framework that reconciles these facts. We show that, under some assumptions, gender differences in hours worked can account for a substantial portion of the gender wage gap and that increases in the returns to hours worked over the past four decades slowed progress in reducing the gender pay gap.

JEL Classification: J16, J22, J31, J33

Keywords: hours, wages, occupation, gender wage gap

Corresponding author:

Christian vom Lehn
Brigham Young University
Department of Economics
435 CTB
Provo, UT 84602
USA

E-mail: cvomlehn@byu.edu

* The authors would like to thank Joe Altonji, Claudia Goldin, Larry Katz, Larry Kahn, Jessica Pan, and Jesse Rothstein for useful comments. They also would like to thank Sam Higbee and Shawn Martin for excellent research assistance. The authors are grateful for research support from the College of Family Home and Social Science of Brigham Young University.

Employed women work, on average, three to five hours fewer than employed men, even controlling for demographics and occupation.¹ A recent literature has argued these gender differences in hours worked and preferences for job flexibility are crucial to understanding the gender earnings gap. However, in order for hours worked to have a significant bearing on the gender *wage* gap, it must be the case that the wage return to hours worked is meaningful. While prior research has documented significant heterogeneity in the income elasticity of hours worked across occupations, the *average* estimated income elasticity is close to unity. This implies that the *wage-hours* elasticity is close to zero. In this case, the gender differences in hours worked would be incapable of explaining, in an accounting sense, a substantial fraction of the gender wage gap.²

Motivated by this surprising finding, we carefully re-examine the relationship between weekly hours worked and wages. Using individual data from the ACS, we contrast the estimated returns to hours worked when estimated at the individual and the occupation level. To reconcile these differences in returns to hours within and across occupations, we construct a simple model of worker compensation and job choice to illustrate how worker preferences/constraints over hours worked could generate large returns to hours worked across occupations and yet small returns to hours within occupations. We then apply these insights to the gender wage gap and illustrate how the returns to hours worked across occupations can account for a substantial portion of the gender wage gap. We further illustrate how increases in the returns to hours worked across occupations over time has slowed the closing of the gender wage gap.

¹ Based on 2016 ACS data.

² It is still possible for hours worked to explain gender wage differences through other channels (e.g., by increasing the likelihood of promotion in the future).

Related Literature

Our paper naturally builds on a recent literature emphasizing that gender differences in hours worked and preferences for job flexibility are crucial to understanding the gender pay gap, including Bertrand et al. (2010), Gicheva (2013), Cha and Weeden (2014), Goldin (2014), Goldin and Katz (2016), and Cortes and Pan (2019).³ Much of this literature has documented important heterogeneity in the returns to hours worked across occupations, such as Bertrand et al. (2010), Goldin (2014), and Goldin and Katz (2016), who show much higher returns to working longer hours in law and business occupations than occupations such as pharmacists, and correspondingly larger gender wage gaps.⁴ In contrast to this existing literature, we focus on the *average* level of the gender wage gap and returns to hours worked for *all* workers and not the variation in gaps and returns across occupations.⁵

Prior work, including research by Goldin (2014), Cortes and Pan (2019) and Bick, Blandin, and Rogerson (2019) find that, on average, weekly hours worked are not strongly associated with higher wages *within* occupations.⁶ Our paper finds similar results at the

³ More generally, Blau and Kahn (2017) offer an excellent overview of the evidence on the size and explanations for the gender wage gap in the cross-section and over time.

⁴ Job characteristics beyond hours are important in thinking about the gender wage gap. For example, a job with relatively low hours but rigid requirements about scheduling may be disproportionately unappealing to women. We abstract from this in the current paper as is common in this literature (Goldin (2014), Cortes and Pan (2019), Cha and Weeden (2014)). Men and women may also differ on actual experience which is an important determinant of wages (Blau and Kahn 2013). Cubas et al. (2017) demonstrate that workers who work at peak times of the day are paid a higher wage, which may affect occupational choice and the gender earnings gap.

⁵ What drives the returns to longer hours across occupations is an interesting and important research question in itself. Several explanations have been suggested in prior work. One explanation involves compensating differentials – i.e., individuals receive additional remuneration for the disamenity of long hours. Most other explanations stem from some feature of the production function in various occupations. As Goldin (2014) suggests, another explanation is that workers in some jobs are less substitutable for each other. It may also be that worker effort in high hours occupations are less easily observable to supervisors, so these jobs pay efficiency wages to ensure optimal effort. Or perhaps some long hour jobs also require scarce skills which demand higher wages (e.g., leadership in the case of CEOs). Finally, there may be increasing returns to hours spent in production—essentially returns to experience but on a shorter time frame.

⁶ See, for example, Figure 3 of Goldin (2014) and Table 2 of Cortes and Pan (2019). However, some occupations have non-negative wage returns; see Goldin (2014) for discussion.

individual level but documents a much larger relationship between hours and wages at the occupational level.

Our emphasis on hours variation across occupations and jobs is consistent with observations in the current literature on the gender wage gap. Cortes and Pan (2019) conjecture that relaxing hours constraints faced by women “enabled [them] to enter a different job... within the same occupation” and generated a significant closing of the gender pay gap. Further, Wasserman (2015) documents that when the total hours worked at medical residencies were capped at 80 hours a week, this had a significant impact on which medical specialties women chose to enter, all within the occupation of “doctor.” Based on our findings, we agree with the assessment of Cortes and Pan (2019) that “shifts in women’s position in the males’ earnings distribution are unlikely to occur if women simply worked more hours and stayed in the same job.”

Our results are consistent with recent work that demonstrates relaxing institutional constraints that limit women’s ability to work longer hours leads to an increase in hours worked and wages (Cortes and Pan 2019, Duchini and Van Effenterre 2019). It is also consistent with the findings of Hirsch (2005) who finds that the wage losses associated with part time work largely disappear when one controls for industry, occupation, and task measures.

Our analysis also relates to the literature on labor supply. Typically, studies of labor supply assume that workers can choose their hours in response to wages. However, in many settings workers have limited discretion over hours worked which complicates interpretation of the relationship between wages and hours worked.⁷ Several studies address this, and thus relate most closely to ours. Altonji and Paxson (1988, 1992) focus on changes in hours and wages

⁷ In some atypical settings it is quite plausible that workers can adjust hours such as stadium vendors (Oettinger 1999) or taxi cab drivers (Farber 2015).

among job switchers, arguing that job switching is a primary channel for workers to adjust their hours of work and associated compensation. More recently, Chetty et al. (2011) study labor supply decisions in a framework where firms set hours constraints on workers and thus changes in hours worked induced by wage changes for individuals typically come when switching jobs. Our analysis confirms that choices across jobs and occupations are crucial to individuals' labor supply decisions.

Within versus Across Occupation Estimates of the Wage-Hours Relationship

We begin our empirical analysis by estimating the relationship between hours worked and wages at both the individual and occupation level using data on hours worked and income from the Current Population Survey, the Census, and the American Community Survey (ACS). We focus on the civilian non-institutionalized population of workers between ages 25 and 55, dropping observations with missing occupation data as well as individuals who are self-employed. Our primary measure of weekly hours worked is given by an individual's reported usual weekly hours worked for the prior year. Hourly wages are computed by dividing total reported wage and salary income for the prior year by the product of usual hours worked and weeks worked in the previous year. Where incomes are top coded, we set them to be 1.5 times above the top coded level, and we drop observations with reported wages below half the minimum wage. We also drop observations with imputed usual hours, weeks worked, income or occupation.

Our analysis establishes two key facts: (1) the relationship between hours worked and wages at the individual level is consistently small; and (2) the relationship between hours worked and wages at the occupation level is substantively larger. Figure 1 shows the relationship

between raw wages and hours at both the individual (left panel) and occupational level (right panel), using data from the 2016 ACS. To illustrate the relationship between individual hours worked and wages, we divide individuals across percentiles of the logarithm of number of reported usual weekly hours. For occupations, we compute the average log hours and average log wages for occupations at the three digit level. The average log wages and average log hours for each individual bin or occupation are plotted as open circles on the figure, with the size of the circles corresponding to the size of the bin or occupation, as determined by survey weights and the number of workers in each bin/occupation. The line is the best fit line from a bivariate regression.

For individual level hours and wages, the relationship is not particularly strong, implying that a 10 percent increase in usual hours worked is associated with 2.4 percent higher hourly wages. In contrast, the relationship between occupational averages for hours and wages is much stronger, implying that a 10 percent increase in hours worked is associated with a 23.5 percent higher hourly wage. Thus, increases in hours across occupations are associated with an order of magnitude larger increase in wages than increases in hours by individuals.⁸

Within Occupation

We now show that the weak relationship between hours and wage at the individual level is robust across various samples and specifications. Table 1 reports estimates from a standard

⁸ The results in Figure 1 show that the variation in average hours worked across occupations is notably lower than the overall variance in hours worked across workers. In terms of occupational variation, approximately 80 percent of the variation in average hours can be explained by the fraction of workers in that occupation who work part-time. We discuss robustness of our findings to focusing just on full-time workers in Appendix D. In terms of specific occupations, the occupations with the highest average hours are senior management occupations (including CEOs), lawyers, doctors and EMTs, workers in oil/gas mining (drillers, engineers, miners), and high skilled sales occupations (particularly financial sales). Erosa et al. (2017) also analyze within occupation variation in hours and its relationship to average hours in an occupation.

Mincer wage regression using individual level data. Specifically, we estimate variants of the following model:

$$(1) \quad \ln w_{io} = \beta_0 + \beta_1 \ln h_{io} + X_i \beta_2 + \delta_o + \varepsilon_{io}$$

where i indexes individuals and o indexes measured 3 digit occupation. The variables w and h represent wages and hours respectively, and X_i is a vector of individual covariates including binary indicators for race (Black, Hispanic, Asian and other race) and gender, indicators for educational attainment (less than high school, high school only, some college, BA degree, masters degree or higher), and a quartic in age. δ_o represents a set of occupation fixed effects. All regressions are weighted with the person weights provided by ACS, and robust standard errors are clustered by three-digit occupation.

Column 1 of Table 1 presents estimates from the bivariate regression of log wage on log hours (as shown in Figure 1). The inclusion of race, gender, age and educational attainment in column 2 further reduces the coefficient on log hours.

To this point, we have not controlled for occupation fixed effects and consequently these point estimates represent some weighted average of the returns to hours within and across occupations. Including occupation fixed effects in column 3, we isolate the return to hours within occupations which causes the estimated return to flip signs. Within occupation, a 10 percent increase in hours is associated with 1.2 percent *lower* hourly wages.

Columns 4-5 show that this same basic pattern holds true for men and women. In Appendix D, we show that the same general relationship holds for those with and without a BA degree, when controlling for major industry group, and in samples limited to full-time workers, defined as those working at least 35 hours per week.

Table 2 shows that these relationships are robust to a variety of alternative specifications using individual level data from the March CPS or Outgoing Rotation Groups of the CPS, pooled across years 2012-2017. Columns 1-3 in Table 2 mirror the results from Table 1 that use ACS data. As in the ACS data, the bivariate relationship is positive and significant (elasticity of .195) but falls to zero with the inclusion of basic controls and becomes negative with the inclusion of occupation fixed effects. In Columns 4 and 5, we use the individuals with repeated observations in the outgoing rotation groups of the CPS, which allows us to control for person fixed effects.⁹ Regardless of whether additional demographic and occupation controls are included or not, we observe a negative relationship between hours and wages, suggesting that the relationship is not being driven by unobserved time-invariant heterogeneity.

Given that wages are constructed as the ratio of income and hours worked, measurement error in the hours worked variable would introduce a negative bias in our estimate of the impact of hours worked.¹⁰ To address this measurement error and resulting division bias, the specification in column 6 instruments for a worker's reported usual hours this year with the usual hours the individual reported in the prior March. Although the coefficient increases from -.14 to .15, it remains small. Alternatively, column 7 instruments for a worker's reported usual hours this year using the number of hours the worker reported actually working in the prior week. The coefficient is -.001 and statistically insignificant.

Table 3 uses the CPS's Outgoing Rotation Groups to examines the hours-wage relationship separately for salaried versus hourly workers. Given that, in salaried jobs, the

⁹ One distinction from the ASEC and ACS/Census samples is that wages for wage earners in the ORG sample are collected directly rather than computed, and wages of salary earners are computed by dividing weekly earnings by usual hours worked.

¹⁰ Note that this bias is not specific to our use of wages on the left hand side of our specifications. If we were to just use income instead of wages, as in Goldin (2014), there would be no division bias, but the coefficient would still be biased by measurement error leading to attenuation. In fact, the coefficient would mechanically be the coefficient we obtain from the wage regression plus one.

marginal effect of additional hours worked on the wage is negative and that hourly wage workers may earn higher wages for overtime work, we would expect a stronger relationship for hourly wage workers relative to salaried workers. The wage-hours relationship is indeed slightly larger for hourly wage workers, though remains small for both groups. With the inclusion of occupation fixed effects, the coefficient on log hours is .17 for wage earners. This rises slightly to .23 when we instrument log hours using lagged hours. In contrast, the analogous specification for salary workers produces a negative relationship between log hours and wages of -.22. Instrumenting for log hours with lagged hours increases the coefficient to .01, which is statistically insignificant.

Across Occupation

We now turn to estimates of the wage-hours relationship that rely on variation across occupations. Table 4 reports estimates from wage-hour regressions at the three-digit occupation level. Specifically, we estimate variants of the following model:

$$(2) \quad \bar{w}_o = \beta_o + \beta_1 \bar{h}_o + \beta_2 tasks_o + \varepsilon_o$$

where \bar{w}_o (\bar{h}_o) indicate average of log wages (hours) in occupation o .

In contrast to the small estimated relationship between hours and wages at the individual level, the relationship between hours and wages at the occupational level is consistently large. As seen in Figure 1, the bivariate relationship between occupational wages and hours implies that a 10 percent increase in hours worked is associated with a 23.5 percent higher hourly wage. This relationship remains large when we residualize hours and wages to account for differences in the characteristics of workers in each occupation. To construct these residuals, we regress log wage (hours) on the same demographics used at the individual level (shown in column 3 of Table 1) as well as a full set of occupation fixed effects. We use the coefficients on the occupation fixed

effects as measures of average residualized log wage and log hours. Column 2 shows the results of the bivariate regression of residualized log wage on residualized log hours. The relationship between hours and wages is still very large, with an elasticity of roughly 2. In Appendix D, we show that these large elasticities persist when we focus on a subset of occupations (high or low skill, male or female dominated, dropping the largest occupations), when we focus on particular sub-samples of workers (men and women separately, full-time workers), and when we control for additional characteristics at the individual or occupational level (major industry group, fraction women in each occupation, 1 digit or 2 digit occupations).

While we control for observable measures of individual skill in Column 2, it is possible that unobserved heterogeneity may account for the observed occupational-level relationship. For example, it may be that high skill workers find working longer hours less unpleasant, and the large relationship between hours and wages across occupations simply represents worker sorting on the basis of this heterogeneity. To address this concern, we again use outgoing rotation groups of the CPS to residualize both hours and wages on person fixed effects.

Column 5 of Table 4 shows that using CPS data yields substantively similar results to what we found in the ACS. Once we control for demographics and education, a 10 percent increase in occupational hours is associated with an 18 percent increase in occupational wages. Controlling for individual fixed effects, a 10 percent increase in occupational average hours is associated with a similarly sized 17 percent increase in occupational wages. These results suggest that differences in selection into occupation on the basis of individual time-invariant skills, preferences, or constraints do not drive the strong relationship between wages and hours at the occupation level.

Our estimates of the relationship between hours and wages across occupations may also reflect occupation characteristics. For example, it may be that certain occupational tasks are associated with longer hours and that these tasks are compensated with higher wages. Following the literature on occupational tasks, we control for five occupation-level task measures: social skills as defined in Deming (2017), abstract analytical, manual and routine tasks as defined in Acemoglu & Autor (2011), and competitiveness as defined in Cortes and Pan (2018).¹¹ When we control for these tasks in column 3-4 of Table 4, the return to average hours decreases substantially. This suggests that a substantial portion of the returns to hours worked across occupations stems from higher hours being associated with certain highly compensated job tasks.¹² However, even when controlling for these tasks, the relationship between hours and wages remains quite large. In column 4, for example, the estimates suggest that a 10 percent increase in average hours worked is associated with 9.3 percent higher hourly wages. In the following section, we discuss how in some cases, the bivariate relationship between occupational wages and hours may be most relevant for understanding the returns to preferences regarding hours worked.

While we acknowledge that the estimates we discuss above still may be subject to omitted variable bias, it is worth noting that the magnitudes of our estimates are consistent with recent quasi-experimental evidence aimed at a causal understanding of the role of hours worked in determining wage differences between men and women. For example, Duchini and Van

¹¹ We include abstract analytical, manual, routine and social tasks constructed from the ONET 4.0. For more detail on how the task measures are constructed, see Appendix B. There are seven occupations for which task measures are unavailable, which are dropped in these regressions.

¹² High task intensity of any of these tasks is generally predictive of higher hours worked in an occupation; this may be due to high task intensity implying higher specialization and lower worker substitutability. However, in terms of predicting higher occupational wages, abstract tasks have the largest impact, followed by competitive and social tasks. Thus, the occupational wage-hours relationship is most sensitive to the inclusion of these three tasks.

Effenterre (2019) show that changes in school schedules in France that reduced time commitments to child care for mothers raised average hours worked by 1% and average wages by 1.5%, implying an elasticity of 1.5, which lies in the middle of our cross-occupation estimates of the returns to hours worked.

Conceptual Framework

In order to understand possible reasons for the difference in the returns to hours at the individual and occupation level, and to understand what these differences mean for the gender wage gap, it is important to make two theoretical points. The first is that when employers contract with workers over an expected workload instead of (or in addition to) the number of realized hours an individual works, the empirical relationship between average wages and average hours at the occupation level may be very different than the corresponding relationship at the individual level. We hypothesize that the occupation-level relationship may proxy for the types of trade-offs workers face when selecting across jobs *within* an occupation as well as when selecting jobs *across* occupations.

The second point is that when workers choose jobs on the basis of expected hours worked (possibly because of constraints that preclude selecting jobs with high hours), if those preferences are separable from preferences over the tasks of the job, the consequence for their wages is dictated by the bivariate relationship between hours and wages. For example, it may be the case that jobs requiring high levels of analytical reasoning pay high wages but typically require a high hours commitment as well. In this case, a worker who is constrained to select a

job with low hours will miss out on both the direct effect of high hours on wages as well as the indirect effect of high hours operating through access to jobs requiring analytical reasoning.

Expected Versus Realized Hours

Consider a simple conceptual framework of employee compensation. The key feature of our framework is distinguishing between the expected hours of a job and the realized hours an individual may work at any point in time in a particular job. We assume that a job, j , is an employment contract offered by a firm and is associated with a compensation level, $c(h_j^e, h_j)$, which is a function of ex ante expected hours worked prior to starting the job, h_j^e , and the actually realized hours worked, h_j . These contracts are determined in hedonic equilibrium by firms offering contracts consistent with profit maximization and workers choosing jobs consistent with utility maximization. Importantly, we allow the return to expected hours to differ from realized hours.

It is helpful to consider a few illustrative examples of common compensation schedules that have this form. Consider first the case of salaried workers. In this case, workers agree to a level of compensation based on an ex ante expected workload, but the salary does *not* adjust in the short run to the actual realized hours worked. This compensation is determined in a simple hedonic equilibrium between workers and firms. Realized hours worked may vary from expected hours due to variation in skill across workers, idiosyncratic expectations of a demanding or lenient supervisor, or other reasons.¹³ In the context of our framework, this implies that the

¹³ Our notion of realized hours worked is different conceptually than the measure of “actual” hours worked used in Table 2 to instrument measures of “usual” hours worked in the CPS. In the CPS, usual hours worked corresponds to typical weekly hours of that particular worker once employed; actual hours worked corresponds to the actual hours worked in the prior week. Our notion of expected hours worked corresponds to a worker’s expectation of the work hours prior to beginning the job and our notion of realized hours worked corresponds to the hours worked once a worker is employed. Hence, realized hours in our framework may correspond more closely to usual hours worked as

compensation function has the following form: $c(h_j^e, h_j) = c(h_j^e)$. This leads to a potentially large return to expected hours across jobs if $c'(h_j^e)$ is large, even as the individual's personal compensation to realized hours worked is zero. The implication for the implicit hourly wage, $\frac{c(h_j^e)}{h_j}$, is that the effect of a marginal hour worked within a job is to *reduce* the wage.

While many jobs are salaried, other jobs pay wages at an hourly rate. In such jobs, it seems natural to model the compensation function in the following way: $c(h_j) = wh_j$, where w is the hourly wage. In such jobs, earnings scale linearly with hours worked, and the wage return to the marginal extra hour worked is zero. Even in these cases, however, the firm may place restrictions on the range of hours offered to the worker. As a consequence, it may not be possible to obtain a high wage hourly job without committing to a high expected hours of work, h_j^e . In this case, the menu of compensation across jobs may take the following form: $c(h_j^e, h_j) = w(h_j^e)h_j$. For an individual working an extra hour within a given job, the wage is fixed at $w(h_j^e)$. However, moving across jobs, the wage may increase in expected hours if in equilibrium $w'(h_j^e) > 0$. Hence, in hourly jobs the wage return of an increase in expected hours across jobs may be quite high, even as the wage is constant within a job.

A natural exception to this is the case where workers are paid a higher hourly rate when they are required to work overtime. This would suggest a higher relationship between realized hours worked and wages, which is consistent with what we observe when splitting our results into hourly wage workers and salaried workers in Table 3. However, if the possibility of

measured in the ACS and CPS. The ACS and CPS do not contain a similarly close analog to our concept of expected hours, instead we proxy for this using average hours worked by occupation.

overtime is only available in jobs with already high expected hours, this still implies that expected hours play a potentially substantial role in a worker's compensation.

Both of these examples suggest that the wage return to expected hours of work across jobs may be substantially larger than the wage return of realized hours worked. Our descriptive analysis of the relationship between hours and wages is consistent with this conjecture. We believe that occupation average hours represents a good proxy for expected hours of a job since it is averaged across many workers, and thus interpret the returns to hours across occupations as representing a useful approximation of the return to expected hours, h_j^e . In contrast, estimates of the return of an individual's actual hours worked, controlling for occupation fixed effects, proxy for the return to realized hours.

Hours Choices and Compensation

Our second theoretical point is aimed at understanding why the returns to hours across occupations is so large and how we should think about the role of other job characteristics in this relationship. Our key insight is that hours considerations can have a direct effect on wage through the effect of hours on wages (holding job tasks fixed) as well as an indirect effect operating through the correlation of job task mix and expected hours. This generates ambiguity regarding whether one should control for task mix when estimating the return to hours worked.

To illustrate this intuition, we develop a simple model, which is included in Appendix A. In this model, individuals select jobs for a variety of reasons, including tasks, hours, and compensation. Moreover, wages are determined by the hours worked and tasks. Our framework explains how these forces interact to affect job choice.

In our model, an individual, facing a binding constraint over hours worked or preferring to work fewer hours, will experience low utility from *all* high hours occupations. This individual would tend to avoid such occupations, even if they are associated with lucrative or enjoyable tasks. As a consequence, constraints or preferences that lead to low hours reduce wages both due to the direct relationship between hours and wages as well as the indirect effect operating through the correlation between hours worked and job tasks.

To illustrate this intuition empirically, we separate occupations by how much they use abstract tasks, a job feature associated with high wages and high hours. Figure 2 shows the occupation average hours for the occupations in the top and bottom quartile of occupation-level abstract task measures (using the 2016 ACS). There are a few things to note with Figure 2. First, occupations with high abstract task content have high hours. Second, very few occupations with high abstract task content work fewer than 40 hours a week on average. This suggests that if a worker wants to work in a high abstract task occupation, they will generally need to work more hours than a worker in a low abstract task occupation. Hence, if a binding hours constraint induces a worker to work in an occupation with lower abstract task intensity than desired absent the constraint, controlling for abstract tasks will understate the total (direct and indirect) effect of hours on wages. Essentially, workers may need to commit to longer hours to access high abstract task jobs.

For some workers, the driver of occupation choice may be preferences or constraints over tasks or other job characteristic. Indeed, there is evidence that discrimination, attachment to the labor force, and aversion to competition may be key factors in the gender-based sorting across jobs and occupations (see for example Blau and Kahn, 2017). This fundamentally changes the appropriate way to price any resulting hours differences between workers.

Suppose that, across workers, there are no differences in preferences or constraints over hours. However, differences in preferences or constraints over tasks lead to differences in hours worked because expected hours are correlated with tasks. In this case, our model would suggest that the resulting differences in hours should not be priced at all because differences in job tasks (and not hours) are the fundamental driver of wage differences between the two types of individuals.

To this point, we have laid out two extreme cases – one in which hours drive all wage differences and one in which hours have no role in generating wage differences. Of course, there are intermediate cases. Suppose we use an alternative utility function for workers' so that they select jobs lexicographically on the basis of tasks and then choose hours. In this case, it is appropriate to price hours holding constant job characteristics. This is because differences in constraints to or preference for hours worked will lead to differences only across jobs with similar characteristics. Hence, a shift to a higher hours job will not be associated with any additional wage premium due to correlated changes in job tasks. However, a shift to higher hours will still be compensated, within job task categories.

As we move to considering the gender wage gap, whether one should or should not condition on job characteristics when pricing the gender hours gap depends on the precise model used. If the primary driver of the difference in job and occupational decisions between men and women is the difference in time constraints or preferences over long hours and workers exhibit a willingness to compromise on job tasks, then our theoretical framework suggests that the *unconditional* relationship between hours and wages across occupations may be appropriate. If, however, hours considerations are not the primary driver of occupational differences but rather one of several, it may be appropriate to price the gender hours gap controlling for occupational

task mix. If preferences or constraints over hours do not drive differences in occupational choice at all, then the gender hours gap should not be used to explain the gender wage gap.

Our framework should help researchers interested in understanding the impact of an intervention that relaxes women's constraints over hours such as subsidized childcare. We would predict a substantial wage increase operating through the direct effect of expected hours. If workers are flexible with regards to other job characteristics, wages would also increase due to an indirect effect operating through correlated tasks. Because we are uncertain regarding worker flexibility across tasks, we present analyses of the gender wage gap using occupational returns to hours both in which we control and do not control for occupation characteristics besides hours.

Application to the Gender Wage Gap

It is well documented that, on average, female workers work fewer hours than working men. In 1980, the raw gender difference was 6.6 hours. Even after controlling for standard demographics and occupation fixed effects, the gender difference was still 4.9 hours. While work hours for men and women have converged over time, the gap has not closed. In the 2016 ACS, for example, employed women work an average of 39.0 hours compared with 43.5 for men, a difference of 4.5 hours. After accounting for race, age and education and three-digit occupation fixed effects, we find that women still work 3.2 fewer hours than men.¹⁴

Recent research emphasizes the importance of these differences in hours worked for the gender wage gap. Our analysis of the gender wage gap works within the following regression framework, similar to earlier regressions focusing on the returns to hours:

$$(3) \quad \ln w_i = \pi + \alpha female_i + \beta \ln h_i + X\Gamma + \varepsilon_i$$

¹⁴ If we limit our analysis to full-time full-year workers, women work 1.8 fewer hours than men in 2016.

Note that the gender gap in hours is priced at rate β and α represents the residual wage gap. Given the positive hours gap between men and women, as β rises, the residual wage gap will decline. Indeed, if β becomes large enough, it is possible that the residual wage gap could be negative, even with a substantial unconditional gender wage gap. As we have demonstrated above, individual-level estimates of β show hours worked has a very small positive or even negative correlation with wages (see Table 1). Thus, a simple regression adjustment for hours worked will do little to close the residual gender wage gap.

In column 1 of Table 5, we present a baseline residual wage gap in which we do not account for hours or occupational choice at all. Regressing log wages on a female indicator variable as well as race, age, and education level, we obtain a residual wage gap of .25 log points. We now consider alternative specifications in which we make different assumptions regarding how hours should be taken into account.

In column 2, we add a simple control for the log of realized hours worked. In this specification we do not control for occupation fixed effects as occupational choice may be driven by differences in hours. This model implicitly assumes that there is no difference in the returns between realized and expected hours. The coefficient on log hours worked is quite small at .05, and reflects the return to some combination of realized and expected hours (along with other occupation-specific factors that might be correlated with expected hours and wage). As a consequence, it does little to shrink the residual wage gap. This clearly illustrates the difficulty in explaining the gender wage gap by a simple regression adjustment for hours worked.

In column 3, we add occupation fixed effects but eliminate the control for hours worked. This specification implicitly controls for differences between men and women in the expected hours associated with their occupational choices at the level of the 3-digit occupation codes we

observe. The inclusion of occupation fixed effects substantially reduces the residual gender wage gap to .16.¹⁵ This substantial reduction is due to the fact that the fixed effects control for not only the expected occupational hours, but also any other occupational differences, such as job tasks, that drive compensation.

In column 4, we again include occupation fixed effects but also control for realized hours worked. This regression prices the hour-wage relationship at the individual level. In this model, the coefficient on log hours should be interpreted as the return to realized hours worked driven by idiosyncratic factors such as supervisor leniency, unobserved speed at completing tasks, etc. We estimate a residual wage gap of .17, slightly *larger* than what we find in the prior specification (column 3) in which we do not control for hours. This is due to the fact that realized hours have a negative price in the regression. As a consequence, the fact that women work fewer hours within occupations means that they should have higher wages than men.

However, at the other extreme, one can instead assume that the hours differences between men and women, both across and within occupations, is driven entirely by gender differences in preferences/constraints over hours worked, manifest in choosing jobs with different levels of expected hours worked. Under this assumption, researchers would wish to price the gender difference in hours at the expected hours wage rate. One tractable way to implement this thought experiment is to instrument individuals' realized hours with the occupation level average hours, omitting the individual's own contribution to this average.¹⁶ Note that we are not using the instrumental variables approach here to address omitted variable bias, as is often the case. The

¹⁵ Additionally, we have experimented with further including interaction terms between hours worked and occupation fixed effects, which would allow for an occupation-specific wage-hours elasticity. This doesn't significantly change the point estimate on the residual gender wage gap.

¹⁶ We do not control for occupational fixed effects, since they are collinear with the occupational average hours. This method is similar to that used in Angrist (1991).

IV approach here is simply a convenient way to use the across occupation variation to identify the wage-hours relationship in the context of an individual level wage regression.¹⁷

Column 5 of Table 5 shows these results. Consistent with our expectation, the coefficient on hours worked is 1.7, quite similar to the aggregate cross-occupation return to average hours worked presented earlier. In this specification, the residual gender wage gap virtually disappears. In this extreme case, these results would suggest that gender-correlated choices over hours worked could account for the lion's share of the difference in wages between men and women.

However, given that other considerations may also play an important role in the gender differences in job choice and the significant correlation between occupational hours and occupational-level task measures, we also consider how the gender wage gap is impacted when we hold fixed the set of occupational tasks. Depending on preferences, if we do not control for occupation task measures, we may misrepresent the nature of occupational choices by women and overstate the role of hours worked. By using a measure of the returns to hours across occupations that holds fixed job tasks, we adopt a more conservative measure of the relationship between hours and wages.

In column 6 of Table 5, we show our results when we control for all occupational tasks. Again, we instrument hours worked with average occupational hours (omitting the reference individual). However, we control in the second stage for the set of occupational tasks discussed earlier. This has the effect of lowering the return to hours worked to 0.94. Relative to the prior specification, controlling for task mix expands the residual gender wage gap to 0.09. The gap,

¹⁷ Note here we are implicitly assuming that the expected hours wage rate is the same across measured occupations as well as across jobs within a measured occupation. In reality, it is possible that the return to expected hours may differ across jobs. But because three digit occupation is the finest level of job we can observe, we are unable to assess how this return may differ across jobs within occupations. For the current exercise, we believe the assumption of a constant return to expected hours as a reasonable first-order approximation of the truth that will allow us to gain insight to the role of hours in explaining the gender gap.

however, remains substantially narrower than cases in which hours are priced at the realized hours rate implicit in the individual level regressions.¹⁸ Thus, even with a conservative pricing of the occupational relationship between hours and wages, we still find an important role for hours differences in accounting for the gender wage gap.

Collectively, our results show that in spite of the low observed returns to hours worked within occupations, gaps in hours worked between men and women may yet represent an important component of the gender wage gap. Although it is not possible to provide a definitive pricing of the gender hours gap given available data, our findings are consistent with recent work showing men and women have different willingness to pay for flexible work arrangements (e.g. Mas and Pallais 2017, Wiswall and Zafar 2017). Further, our framework shows that modelling choices about hours worked meaningfully affect the conclusions about hours' contribution to the gender wage gap.

Gender Wage Gap over Time

We now consider how changes in the returns to hours worked over time have impacted the evolution of the gender wage gap. Since the 1980s, the difference between women's log wages and men's log wages has shrunk considerably (blue line in Figure 3), and at the same time the gender gap in log hours has also been shrinking, albeit to a lesser extent. And yet, over this same time period, the returns to hours worked has increased substantially, magnifying the impact of the gap in hours on the gap in wages.

¹⁸ The number of observations is slightly lower in column 6 compared with columns 1 through 5. This is because a small number of occupations do not have ONET task information. If we replicate the specifications in columns 1-5 on the sample used in column 6, the results are qualitatively the same.

Figure 4 plots the estimated wage-hours elasticities from both within and across occupation models in 1980, 1990, 2000 and 2016.¹⁹ Looking at the individual level wage-hours estimates, we see the relationship has increased modestly between 1980 and 2016, but is always negative and small in absolute magnitude. Examining the relationship between occupation level wages and hours, we see a much more dramatic change. In 1980, a 10 percent increase in occupation average hours was associated with a 9.1 percent increase in average wages. By 2016, the relationship had doubled so that a 10 percent increase in average hours was associated with a 19.5 percent increase in average wages.²⁰ Even when controlling for occupational tasks in Figure 5 (using the same tasks as in Table 4), we see that the coefficient on log hours remains large and positive across all time periods and is still trending upward over time. With controls, the coefficient on hours is 0.44 in 1980 and rises to 0.92 by the year 2016.

Taken together, these facts suggest that wage convergence between men and women would have been larger if the wage premium for hours worked had remained the same. To quantify how the return to hours worked influenced evolution of the gender gap, we conduct a simple counterfactual exercise. We calculate what the evolution of the gender gap would have been if the returns to hours worked had remained at 1980 levels. This exercise examines the “unconditional” gender wage gap—that is not conditioning on hours but controlling for other characteristics such as age and education. We provide more detail regarding this procedure in Appendix C.

¹⁹ These estimates come from the coefficient on log hours in each year corresponding to the regressions with individual controls and occupation fixed effects (Table 1 Column 3) and the residualized occupation regressions (Table 4 Column 2).

²⁰ We have calculated analogous measures for the CPS, looking at 5-year moving averages of the coefficients to take into account the smaller sample size. We also examined both OLS and IV estimates of the return to individual hours worked. For the IV specifications, we instrument usual hours worked in the reference year with usual hours worked in the prior year to account for measurement error. All these specifications generate a similar pattern.

Again, we are faced with the question of what price should be used to relate the hours gap to the wage gap period by period. In Table 6 we present these results for four different prices for log hours worked. In each of the cases, row 1 shows the evolution of the “total” wage gap—that is, the residual wage gap plus the wage gap that can be explained by the gender gap in hours. Row 2 considers the counterfactual evolution of the gap assuming that the residual gap evolved according to the observed time series but the prices on hours are held constant at 1980 levels. Row 3 shows the percentage difference between Row 1 and Row 2—how different the gender wage gap would have been if returns to hours remained constant at 1980 levels.

This exercise suggests that the changing returns to hours over have had a significant countervailing effect on closing the gender wage gap. When we price the hours gap at the rate implied by the return to realized hours (top panel), the gender wage gap would have been nearly 20 percent smaller in the year 2016 had the returns to hours worked not increased over the time period; adding occupation fixed effects (second panel), the gap would be roughly 23 percent smaller. In the bottom two panels of Table 6, pricing the hours gap at the rate implied by the cross-occupation wage-hours schedule, the counterfactual wage gap would have been even lower had the hours premium remained at 1980 levels -- 43 percent lower using the unconditional relationship between hours and wages and 26 percent lower when controlling for tasks.

Conclusion

In this paper, we examine the wage returns to hours worked, and the extent to which hours differences can explain the gender pay gap. We demonstrate that, while the hours worked by an individual has only a weak relationship with wages, the average hours within an occupation is much more strongly related to wages. This relationship holds, controlling for individual characteristics and occupation tasks. Indeed, a 10 percent increase in occupational

hours worked is associated with between a 10 and 20 percent increase in wages depending on the specification.

We then revisit the gender wage gap. Because hours have such a weak relationship with wages, they explain very little of the gender wage gap when priced using individual data. We demonstrate that the extent to which the gender wage gap is mediated by hours worked depends crucially on assumptions regarding the origins of hours worked differences between men and women. If we assume that the gender difference in hours worked should be priced according to the relationship we observe across occupations, then hours worked can account for a large portion of the unexplained gender wage gap. If we price hours according to the cross-occupation schedule and control for the occupational task mix, hours differences still account for approximately half of the gender wage gap. Critically, these approaches require different assumptions about the origins of differences in hours worked across gender.

We also examine the potential role of hours differences in the evolution of the gender wage gap over time. Consistent with Cha and Weeden (2014), the increase in returns to hours worked imply that the counterfactual wage gap would have been substantially narrower had the price of hours not increased over time. Significantly, if we price hours worked according to the occupation-level relationship, the counterfactual gender wage gaps would have been between 26 and 43 percent smaller. This suggests that the increase in the return to occupational hours worked made it increasingly difficult for women to achieve wage parity with men.

While this article illustrates the importance of thinking carefully about the returns to hours worked and how such differences should inform analysis of issues like the gender pay gap, we recognize that much further research is needed. It is still not clear which factors drive the observed differences in hours worked across gender, or the relative importance of various

factors. It would also be useful for researchers to examine further why some occupations have higher hours on average, and why the returns to hours have increased over time. We leave these questions to future research.

References

- Acemoglu, Daron, and David Autor. "Skills, tasks and technologies: Implications for employment and earnings." *Handbook of Labor Economics*. Vol. 4. Elsevier, 2011. 1043-1171.
- Altonji, J.G. and Paxson, C.H., 1988. Labor supply preferences, hours constraints, and hours-wage trade-offs. *Journal of labor economics*, 6(2), pp.254-276.
- Altonji, Joseph G., and Christina H. Paxson. "Labor supply, hours constraints, and job mobility." *Journal of Human Resources* 27, no. 2 (1992): 256-279.
- Angrist, J. D. (1991). Grouped-data estimation and testing in simple labor-supply models. *Journal of Econometrics*, 47(2-3), 243-266.
- Autor, David H., Larry F. Katz, and Melissa S. Kearney. "Trends in U.S. Wage Inequality: Revising the Revisionists." *The Review of Economics and Statistics* 90.2 (2008): 300-323.
- Autor, David H., Frank Levy, and Richard J. Murnane. "The skill content of recent technological change: An empirical exploration." *The Quarterly Journal of Economics* 118.4 (2003): 1279-1333.
- Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz. "Dynamics of the gender gap for young professionals in the financial and corporate sectors." *American Economic Journal: Applied Economics* 2, no. 3 (2010): 228-55.
- Bick, Alexander, Adam Blandin, and Richard Rogerson (2019). "Hours and Wages," Working Paper.
- Blau, Francine D., and Lawrence M. Kahn. "The gender wage gap: Extent, trends, and explanations." *Journal of Economic Literature* 55.3 (2017): 789-865.
- Blau, F. D., & Kahn, L. M. (2013). The feasibility and importance of adding measures of actual experience to cross-sectional data collection. *Journal of Labor Economics*, 31(S1), S17-S58.
- Cha, Youngjoo and Kim A. Weeden (2014). "Overwork and the Slow Convergence in the Gender Gap in Wages." *American Sociological Review*. 79(3): 457-484.
- Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri. "Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records." *The quarterly journal of economics* 126, no. 2 (2011): 749-804.
- Cortes, P., & Pan, J. (2018). Occupation and gender. *The Oxford Handbook of Women and the Economy*, 425.
- Cortes, Patricia and Jessica Pan. (2019). "When Time Binds: Returns to Working Long Hours and the Gender Wage Gap among the Highly Skilled." *Journal of Labor Economics*. 37(2): 351-398.

- Cortes, Patricia and Jessica Pan. 2016b. "Prevalence of Long Hours and Skilled Women's Job Choices." IZA DP No. 10225.
- Cubas, G., Juhn, C., & Silos, P. (2017). *Coordinated Work Schedules and the Gender Wage Gap*. Mimeo, University of Houston.
- Deming, David J. "The growing importance of social skills in the labor market." *The Quarterly Journal of Economics* 132.4 (2017): 1593-1640.
- Dorn, David. "Essays on Inequality, Spatial Interaction, and the Demand for Skills." Dissertation University of St. Gallen no. 3613, September 2009.
- Duchini, E., & Van Effenterre, C. (2019). Do women want to work more or more regularly? Evidence from a natural experiment. Working paper.
- Erosa, A., Fuster, L., Kambourov, G., & Rogerson, R. (2017). Hours, occupations, and gender differences in labor market outcomes. *NBER Working Paper*, (w23636).
- Farber, H. S. (2015). Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics*, 130(4), 1975-2026.
- Gicheva, D. (2013). Working long hours and early career outcomes in the high-end labor market. *Journal of Labor Economics*, 31(4), 785-824.
- Goldin, Claudia, and Lawrence F. Katz. "A most egalitarian profession: pharmacy and the evolution of a family-friendly occupation." *Journal of Labor Economics* 34.3 (2016): 705-746.
- Goldin, Claudia (2014), "A Grand Gender Convergence: Its Last Chapter." *American Economic Review*, 104(4), 1091-1119.
- Hirsch, Barry T. (2005), "Why Do Part-Time Workers Earn Less? The Role of Worker and Job Skills." *Industrial and Labor Relations Review*, 58 (4), 525–51.
- Kuhn, Peter and Fernando Lozano (2008). "The Expanding Workweek? Understanding Trends in Long Work Hours among U.S. Men, 1979-2006." *Journal of Labor Economics*. 26(2): 311-343.
- Mas, A., & Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, 107(12), 3722-59.
- Oettinger, G. S. (1999). An empirical analysis of the daily labor supply of stadium vendors. *Journal of political Economy*, 107(2), 360-392.
- Wasserman, Melanie. 2015. "Hours Constraints, Occupational Choice and Fertility: Evidence from Medical Residents." MIT Working Paper.
- Wiswall, M., & Zafar, B. (2017). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics*, 133(1), 457-507.

Table 1 - OLS Regressions of Ln(Wage) on Ln(Hours), 2016 ACS Individual-Level

Log Wage	(1) all	(2) all	(3) all	(4) male	(5) female
Log Hours	0.244*** (0.046)	0.051** (0.026)	-0.117*** (0.002)	-0.149*** (0.004)	-0.108*** (0.003)
Demo. Controls	No	Yes	Yes	Yes	Yes
Occ FE	No	No	Yes	Yes	Yes
N	633,927	633,927	633,927	320,729	313,198

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Columns (1)-(2) use robust standard errors. Columns (3)-(5) include occupation fixed effects. Observations are weighted using perwt. Sample includes prime-age workers aged 25-55. Demographic controls include a quartic in age, indicators for race/ethnicity (black, Hispanic, Asian, other), sex, and indicators for educational attainment (less than HS degree, HS degree, some college, bachelor's degree, masters degree or more). Hours are derived from usual hours worked per week. Hourly wage is calculated as total reported wage and salary income for the prior year divided by the product of usual hours worked and weeks worked in the previous year. We trim wages that are below half of the federal minimum wage and inflate wages to 2012 dollars using the Bureau of Economic Analysis (BEA) National Income and Product Accounts Personal Consumption Expenditures.

Table 2 - OLS Regressions of Ln(Wage) on Ln(Hours), CPS Individual-Level

	All workers						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log Wage	OLS	OLS	OLS	OLS	OLS	IV: Lagged	IV: Actual
Log Hours	0.195*** (0.039)	0.003 (0.020)	-0.144*** (0.003)	-0.09*** (0.006)	-0.110*** (0.006)	0.146*** (0.039)	-0.001 (0.024)
Demo. Controls	No	Yes	Yes	No	Yes	Yes	Yes
Occupation FE	No	No	Yes	No	Yes	Yes	Yes
Individual FE	No	No	No	Yes	Yes	No	No
CPS Survey	March	March	March	ORG	ORG	March	March
N	317,544	317,544	317,544	383,562	383,562	65,815	309,502

All columns use standard errors clustered at the occupation level. Columns (3), (6), (7) and (10) include occupation fixed effects. Columns (4) and (5) use data from the Outgoing Rotation Groups in all months; all other columns use the March Supplement to the CPS. Column (6) instruments usual hours worked with usual hours worked, reported in the previous March. Column (7) instruments usual hours worked with actual hours worked the week previous to the survey. Observations are weighted using asewct. Sample includes prime-age workers aged 25-55. For details on the measurement of hours, wages, and control variables, see notes to Table 1.

Table 3 - Regressions of Log(Wage) on Log(hours), CPS Individual-Level Data, Wage vs. Salary Workers

	Wage Earners					Salary Workers				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	OLS	OLS	OLS	IV: Lagged	IV: Actual	OLS	OLS	OLS	IV: Lagged	IV: Actual
Log Hours	0.304*** (0.035)	0.273*** (0.023)	0.167*** (0.002)	0.227*** (0.022)	0.177*** (0.014)	0.063 (0.040)	-0.110*** (0.026)	-0.220*** (0.004)	0.010 (0.046)	-0.189*** (0.025)
Demo.										
Controls	No	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Occ FE	No	No	Yes	Yes	Yes	No	No	Yes	Yes	Yes
N	382,353	382,353	382,353	96,325	371,227	315,154	315,154	315,154	89,347	304,644

All columns use standard errors clustered at the occupation level. Columns (3)-(5) and (8)-(10) include occupation fixed effects. Columns (4) and (8) instruments usual hours worked with usual hours worked, reported in the previous March. Columns (5) and (10) instruments usual hours worked with actual hours worked the week previous to the survey. Observations are weighted using earnwt. Sample includes prime-age workers aged 25-55. For details on the measurement of hours, wages, and control variables, see notes to Table 1.

Table 4 - OLS Regressions of Ln(Wage) on Ln(Hours), 2016 ACS at the Occupational Level

	(1)	(2)	(3)	(4)	(5)	(6)
	Average Log Wage					
Average Log Hours	2.352*** (0.243)	1.952*** (0.188)	1.156*** (0.204)	0.925*** (0.198)	1.804*** (0.189)	1.669*** (0.331)
Residualized Against Demographics	No	Yes	No	Yes	Yes	Yes
Residualized Against Person Fixed Effects	No	No	No	No	No	Yes
Task Controls	No	No	Yes	Yes	No	No
Sample	ACS	ACS	ACS	ACS	CPS	CPS
N	474	474	467	467	531	522

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Columns (1)-(6) use robust standard errors. Columns (1) and (3) regress occupation average of log wage on occupation average of log hours. Occupations are weighted by the occupation total of individual weights. Columns (2), (4), (5), and (6) regress occupation average of residualized log wage on occupation average of residualized log hours. Residuals are constructed by regressing log wage (hours) on demographic controls and a full set of occupation fixed effects. We use the coefficients on the occupation fixed effects as measures of average residualized log wage and log hours. Demographic controls used in the residualization include black, hispanic, asian, other race, less than HS degree, HS degree only, some college, bachelor's degree, and a masters degree or more and a quartic in age. Person fixed effects are included using Outgoing Rotation Group samples from the CPS which include multiple observations per worker. Hours are usual hours worked per week. For details on the measurement of hours and wages, see notes to Table 1. Task data is constructed from the ONET 4.0 aggregated to the occ1990dd level. There are five occ1990dd occupations missing task data in the ONET 4.0 which yield missing task data for seven occs. Task controls include social skills, as defined in Deming (2017), and abstract analytical, manual and routine, as in Acemoglu & Autor (2011), and competitiveness as in Cortes and Pan (2018). See appendix B for an explanation how of these composites were created. Task measures are standardized to be mean zero and standard deviation one across the occ1990dd distribution.

Table 5 - Estimates of the Gender Wage Gap, 2016 ACS

Log Wage	(1)	(2)	(3)	(4)	(5)	(6)
Female	-0.252*** (0.023)	-0.245*** (0.022)	-0.157*** (0.007)	-0.167*** (0.007)	-0.018 (0.023)	-0.086*** (0.020)
Log Hours		0.051** (0.026)		-0.117*** (0.020)	1.748*** (0.158)	0.940*** (0.192)
Demo. Controls	Yes	Yes	Yes	Yes	Yes	Yes
Occ. FE	No	No	Yes	Yes	N/A	N/A
Occ. IV	No	No	No	No	Yes	Yes
Task Controls	No	No	No	No	No	Yes
N	633,927	633,927	633,927	633,927	633,927	628,344


Columns (3) and (4) include occupation fixed effects. Columns (5) and (6) are IV regression with individual log hours instrumented with occupation average log hours, thus occupation fixed effects cannot be included. Occupation average log hours are calculated using a leave-out mean by excluding the individual's hours. Standard errors are clustered at the occ level. For details on measurement of hours, wages, tasks, and control variables, see notes to Tables 1 and 4. Observations are weighted using perwt. Sample includes prime-age workers aged 25-55.

Table 6 – Evolution of the Gender Wage Gap

	1980	1990	2000	2016
Hours priced according to OLS regression of individual log wages on individual log hours				
(1) Resid. Gap plus Hours Gap	0.433	0.320	0.259	0.252
(2) Resid. Gap plus CF Hours Gap	0.433	0.275	0.217	0.205
(3) % Diff (2) & (3)		14.2%	16.4%	18.4%
Hours priced according to OLS regression of individual log wages on individual log hours with occupation fixed effects				
(1) Resid. Gap plus Hours Gap	0.317	0.235	0.172	0.152
(2) Resid. Gap plus CF Hours Gap	0.317	0.193	0.136	0.117
(3) % Diff (2) & (3)		17.9%	20.7%	23.4%
Hours priced according to IV regression of individual log wages on individual log hours using occupation hours as instrument				
(1) Resid. Gap plus Hours Gap	0.434	0.320	0.257	0.240
(2) Resid. Gap plus CF Hours Gap	0.434	0.258	0.173	0.137
(3) % Diff (2) & (3)		19.4%	32.8%	43.0%
Hours priced according to IV regression of individual log wages on individual log hours using occupation hours as instrument, controlling for occupation tasks				
(1) Resid. Gap plus Hours Gap	0.426	0.311	0.234	0.205
(2) Resid. Gap plus CF Hours Gap	0.426	0.264	0.193	0.152
(3) % Diff (2) & (3)		15.2%	17.4%	26.2%


Observations are weighted using perwt. Sample includes prime-age workers aged 25-55. The residual wage gap plus hours gap (“Resid. Gap plus Hours Gap”) is defined as $\alpha^t + \beta^t gap_h^t$, where the gap in hours (gap_h^t) corresponds to the difference between male and female log hours in year t. The residual gap plus counterfactual gap (“Resid. Gap plus CF Hours Gap”) is defined as $\alpha^t + \beta^{1980} gap_h^t$. In the first panel, β^{1980} is the coefficient on log hours from a regression of log wage on a female dummy and a set of demographic controls. The second panel adds in occupation controls to derive the coefficient on log hours, β^{1980} . In the third panel β^{1980} is the coefficient from log hours in a IV regression of log wage on female, demographic controls and individual log hours instrumented with occupation average log hours. Occupation average log hours are calculated using a leave-out mean by excluding the individual's hours. In the fourth panel, we add occupation tasks controls to derive the coefficient on log hours, β^{1980} . In each panel, α^t is the coefficient on female and β^t is the coefficient on log hours from the corresponding regression in year t. For details on measurement of hours, wages, tasks, and control variables, see notes to Tables 1 and 4.

Figure 1




Note: Occupations collapsed to 3 digit occ level. Weights are sum of perwt in each occupation. Individuals split across 1000 percentiles of lnhours. Because of the concentrated distribution of the hours worked distribution, this generates 57 distinct bins. Weights are the sum of person weights in each bin.

Figure 2



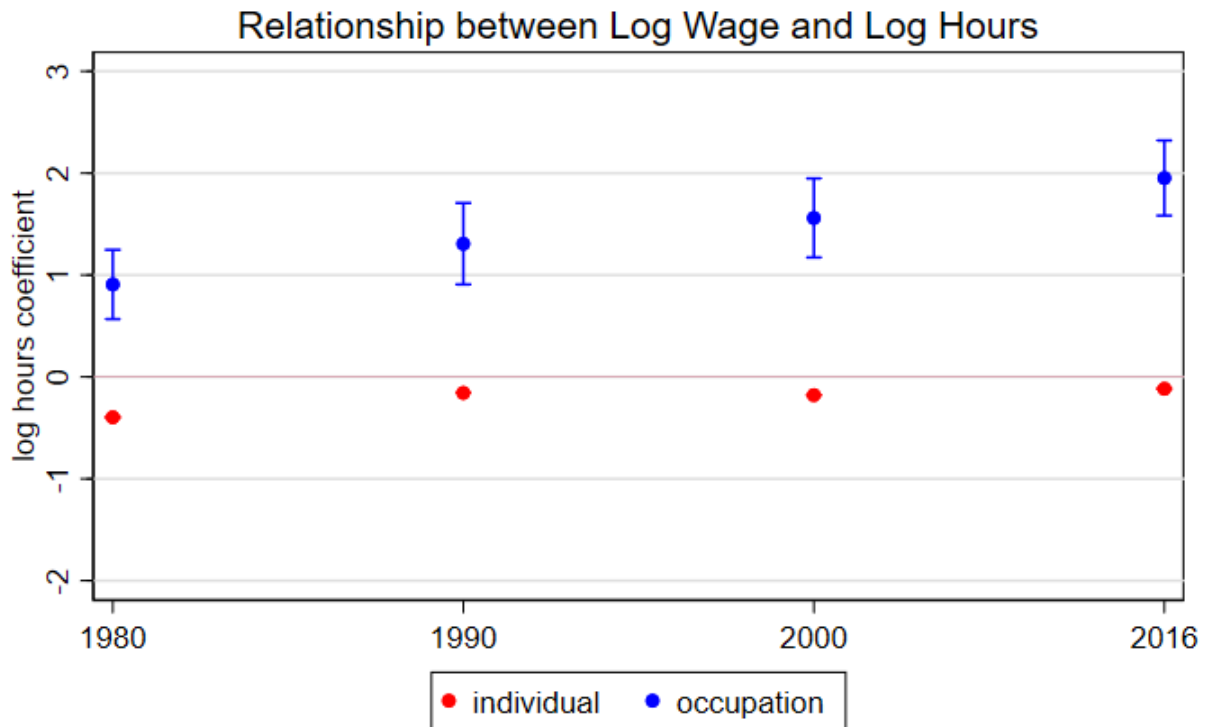
Note: This figure plots the hours worked for occupations in the top and bottom quartile of abstract task from the 2016 ACS.

Figure 3




Note: Plotted is the log difference between female and male average log hours and average log wages.

Figure 4



Note: Presented are point estimates and 95% confidence intervals. Individual regression includes controls for gender, education, race and age and occupation dummies. Occupation level regresses average residualized log wage on average residualized log hours. Log wages and log hours are residualized on controls for education, race, age and female and occupation dummies. The residuals are the predicted occupation dummies. Occupations aggregated to occ level.

Figure 5



Note: Presented are point estimates and 95%. Occupation average residualized log wage is regressed on occupation average residualized log hours. Log wages and log hours are residualized on controls for education, race, age and female and occupation dummies. The residuals are the predicted occupation dummies. Social, manual, routine, competitive, and abstract analytical tasks are included. Tasks are time invariant and come from the ONET 4.0.

Appendix A

Consider a simple model of job choice. Suppose that a group of individuals, perhaps women with a particular skill level, has the following utility function defined over occupations:

$$(1) U_{ij} = \beta_0 + \beta_1 \ln(w_j) + \beta_2 \ln(h_j) + \beta_3 T_j + \epsilon_{ij}$$

Where i indexes the group of individuals, j indexes a job, w_j is the wage of the job, h_j is the hours of the job, T_j represents some other job characteristic (such as tasks of the occupation), and ϵ_{ij} is a type 1 extreme value random utility component. Although we model preferences as a function of hours worked, these can reflect constraints as well as preferences. Consequently, it is most useful to think of this simple utility function as a reduced form representation reflecting both preferences and constraints. Given that we are modeling job choice, h_j is the most naturally interpreted as the expected hours of the job, h_j^e . Since we only have one measure of hours in this section, we suppress the superscript in this section.

Now suppose that there are J different jobs with different combinations of w , h , and T which in equilibrium exhibit the following relationship:

$$(2) \ln(w_j) = \pi_0 + \pi_1 \ln(h_j) + \pi_2 T_j$$

Given this framework, assuming that individuals choose their job to maximize utility, by virtue of the type I extreme value random utility component, the share of individuals within the group that choose job j will be given by

$$(3) \gamma_j = \frac{e^{\beta_0 + \beta_1 \ln(w_j) + \beta_2 \ln(h_j) + \beta_3 T_j}}{\sum_{j=1}^J e^{\beta_0 + \beta_1 \ln(w_j) + \beta_2 \ln(h_j) + \beta_3 T_j}}$$

and the average wages and hours of the group will be given by:

$$(4) E[\ln(w)] = \sum_{j=1}^J \gamma_j \ln(w_j)$$

$$(5) E[\ln(h)] = \sum_{j=1}^J \gamma_j \ln(h_j)$$

With this simple framework, we can better conceptualize what it means for there to be a *causal* relationship between occupational hours and wages. One natural thought is to conceptualize the causal effect of hours as π_1 from equation (2). This is the price of hours holding tasks and other job characteristics fixed. However, in order for this parameter to reflect the tradeoffs that individuals make between hours and wages, it presupposes that every combination of tasks and hours is available. If instead, there is a menu of jobs available and there is a correlation between occupational hours and tasks, then in order to find an occupation with fewer hours, an individual may also have to compromise on the set of tasks associated with the chosen job.

Alternatively, we can conceptualize the trade-off between hours and wages is to see how equilibrium hours and wages change as β_2 , the preference (or constraint) over hours from equation (1), changes. We perform this analysis by differentiating equations (4) and (5) with respect to β_2 :

$$(6) \frac{\partial E[\ln(w)]}{\partial \beta_2} = \sum_{j=1}^J \omega_j [\ln(h_j) - \overline{\ln(h)}][\ln(w_j) - \overline{\ln(w)}] = cov[\ln(h_j), \ln(w_j)]$$

$$(7) \frac{\partial E[\ln(h)]}{\partial \beta_2} = var[\ln(h_j)]$$

where $\overline{\ln(h)}$ and $\overline{\ln(w)}$ are the share weighted averages of the log of hours and wages across occupations. Note the effect of a change in preferences/constraints over hours on wages depends solely on the covariance between hours and wages. Or in other words, if we construct

the implicit wage-hours relationship from changes in preferences/constraints over hours, we obtain:

$$(8) \frac{\frac{\partial E[\ln(w)]}{\partial \beta_2}}{\frac{\partial E[\ln(h)]}{\partial \beta_2}} = \frac{cov[\ln(h_j), \ln(w_j)]}{var[\ln(h_j)]} = \left(\pi_1 + \pi_2 \frac{cov[\ln(h), T]}{var[\ln(h)]} \right)$$

where the last part of the expression is obtained by substituting in equation (2). Note, the expression in (8) is simply the regression coefficient from an unconditional regression of wages on hours.

These equations imply that as individuals either prefer fewer hours or are constrained to work fewer hours, the consequence for their wages is represented by both the direct effect of occupational hours on wages as well as the indirect effect operating through the correlation of hours with tasks. This applies to other characteristics as well. For example, if an individual wishes to work at a job with high levels of social interaction, this will imply a change in wages due both to increased social interaction and the change in correlated tasks and hours.²¹

The implications of this model are specific to the form of the utility function we have chosen. Alternative utility functions could generate different implications regarding the implicit relationship between hours and wages. For example, suppose that workers choose jobs lexicographically on the basis of tasks and then choose hours. In this case, differences in preferences or constraints over hours may generate variation in hours worked, but will not generate differences in the task mix, which remains fixed.

²¹ Note that this finding does not depend on the specific structure of the compensation function. Equation (2) is not necessary but is useful for explaining the intuition of the model.

Online Appendix B

Creating Task Measures

Task measures are constructed from the raw data in the ONET 4.0. We first standardize the raw level variables to be mean zero and standard deviation one across the 900 ONET-SOC occupation codes. We then collapse to 677 soc2000 codes by taking the simple average across ONET-SOC codes associated with a single soc2000 code.¹ The composites are created as the average of the included variables (see details below) and are standardized.

For analysis that uses data from 1980, 1990, 2000 and 2016 we merge on task data using David Dorn's *occ1990dd* classification system. The *occ1990dd* system consists of 330 codes that provide a balanced panel of occupations covering the 1980, 1990 and 2000 Censuses and the 2005 ACS. For the purpose of our analysis, we extend the coverage to the 2016 ACS by creating a crosswalk from the codes used in the 2016 ACS to the *occ1990dd* system. We start with the composite task measures at the soc2000 level and merge on soc2000 weights. We create soc2000 weights by pooling data from the 2005, 2006 and 2007 Occupational Employment Statistics (OES) survey.² We then collapse task measures to the occ2000 and standardize, yielding composite task measures for 445 occ2000 codes. Lastly, we use the occ2000 to occ1990dd from Dorn (2009) and the sum of soc2000 weights for each occ2000 code to collapse task measures to the occ1990dd level. The final dataset merged onto data for 1980, 1990, 2000 and 2016 consist of task data for 325 occupations standardized to be mean zero and standard deviation one. Thus, there are five

¹ For example, onetsoccode 11-1011.01 and 11-1011.02 are collapsed into soc2000 code 11-1011.

² Specifically, we follow the procedure used by Autor & Acemoglu (2011) to create soc2000 weights from the 2005, 2006 and 2007. Weights are calculated as the mean of employment across the three survey waves for each soc code.

occ1990dd occupations for which we are unable to obtain task data; these map into seven occupations in the 2016 ACS coding system.³

We use four composite task measures in our analysis taken previously from the literature. Each measure is constructed as the average of the included variables. For each composite the variable names are given in italics, the variable type in parenthesis and the variable question text in quotations.

1. Social Skills (Deming 2017):

- *Coordination*: (skill) “Adjusting actions in relation to others' actions.”
- *Negotiation*: (skill) “Bringing others together and trying to reconcile differences.”
- *Persuasion*: (skill) “Persuading others to change their minds or behavior.”
- *Social Perceptiveness*: (skill) “Being aware of others' reactions and understanding why they react as they do.”

2. Abstract Analytical (Acemoglu & Autor 2011):

- *Interpreting the Meaning of Information for Others*: (work activity) “Translating or explaining what information means and how it can be used.”
- *Thinking Creatively*: (work activity) “Developing, designing, or creating new applications, ideas, relationships, systems, or products, including artistic contributions.”

³ The *occ1990dd* occupations for which we cannot construct task data are occupations 76, 346, 37, 349 and 415.

- *Analyzing Data or Information*: (work activity) “Identifying the underlying principles, reasons, or facts of information by breaking down information or data into separate parts.”

3. Manual (Acemoglu & Autor 2011):

- *Spend Time Using Your Hands to Handle, Control, or Feel Objects, Tools, or Controls*: (context) “How much does this job require using your hands to handle, control, or feel objects, tools or controls?”
- *Manual Dexterity*: (ability) “The ability to quickly move your hand, your hand together with your arm, or your two hands to grasp, manipulate, or assemble objects.”
- *Operating Vehicles, Mechanized Devices, or Equipment*: (work activity) “Running, maneuvering, navigating, or driving vehicles or mechanized equipment, such as forklifts, passenger vehicles, aircraft, or water craft.”
- *Spatial Orientation*: (ability) “The ability to know your location in relation to the environment or to know where other objects are in relation to you.”

4. Routine (Acemoglu & Autor 2011)⁴:

- *Controlling Machines and Processes*: (work activity) “Using either control mechanisms or direct physical activity to operate machines or processes (not including computers or vehicles).”
- *Spend Time Making Repetitive Motions*: (context) “How much does this job require making repetitive motions?”

⁴ The measure of routine used in Acemoglu & Autor 2011 also included the variable *Structured versus Unstructured Work* (“To what extent is this job structured for the worker, rather than allowing the worker to determine tasks, priorities, and goals?”) but the variable is unavailable in the ONET 4.0 and responses were not gathered until subsequent installations of the ONET.

- *Pace Determined by Speed of Equipment*: (context) “How important is it to this job that the pace is determined by the speed of equipment or machinery? (This does not refer to keeping busy at all times on this job.)”
 - *Importance of Being Exact or Accurate*: (context) “How important is being very exact or highly accurate in performing this job?”
 - *Importance of Repeating Same Tasks*: (context) “How important is repeating the same physical activities (e.g., key entry) or mental activities (e.g., checking entries in a ledger) over and over, without stopping, to performing this job?”
5. Competitiveness (Cortes and Pan 2018)⁵:
- *Competitiveness*: (context) “To what extent does this job require the worker to compete or to be aware of competitive pressures?”

⁵ This measure is unavailable in the ONET 4.0 and is instead constructed using data from the ONET 14.0.

Online Appendix C

Counterfactual Evolution of the Gender Wage Gap

To construct counterfactual measures of the gender wage gap where the returns to hours did not rise over time, we start by estimating individual level wage regressions separately by year for 1980, 1990, 2000 and 2016. We do so both estimating the relation between wages and hours worked via OLS and also via IV in which we instrument individual hours with occupational average hours. This allows us to consider how the counterfactual wage gap would have evolved had the return to hours not increased. Our primary empirical specification is given by:

$$\ln w_{it} = \pi^t + \alpha^t female_i + \beta_1^t \ln h_{it} + \gamma^t demos_{it} + \varepsilon_{it}$$

Because changing demographics and the coefficients on those demographics are not the focus of the analysis, while we control for them in our analysis, we ignore them for the purposes of presentation. Taking means of both sides separately for men and women, we can write the gender wage gap (conditional upon demographic covariates) as

$$gap_w^t = \alpha^t + \beta_1^t gap_h^t$$

This expression allows us to consider how the gender wage gap would have increased if the residual wage gap evolved according to the observed time series but the prices on hours remained constant at 1980 levels. To consider how the gender wage gap would have increased if the residual wage gap evolved according to the observed time series but the prices on hours remained constant at 1980 levels, we allow α to vary by period t ($\alpha = \alpha^t$), but fix β_1^t at the 1980 level ($\beta_1 = \beta_1^{1980}$). Put another way, we allow the unexplained part of the gender wage gap to change period by period but fix the returns to hours worked at 1980 levels.

Online Appendix D: Additional Results

Here we present additional results regarding the relationship between hours and wages, both at the individual level and at the occupational level.

Table D.1 presents extensions of the core results presented in Table 1, focusing on individual level regressions of wages on hours. The extensions in Table D.1 are to: estimate this relationship for college educated and non-college educated workers separately, estimate this relationship for full-time workers only (full-time defined as working 35 or more hours a week) and to estimate this relationship when controlling for 16 major industry groups (as defined by the 2012 Census Industrial Classification system). In each case, we see that our findings are consistent with those in Table 1 – at the individual level, the relationship between wages and hours worked is weak, and significantly negative when occupational fixed effects are included.

Table D.2 presents extensions of the core results presented in Table 4, focusing on the relationship between average wages and average hours worked at the occupational level. In Table D.2 we present these extensions without controlling for occupational tasks; in Table D.3 we present each of these extensions controlling for occupational tasks. The extensions are to estimate the relationship between occupational hours and wages: for occupations below and above the median occupation's fraction of college-educated workers, for occupations below and above the median occupation's fraction of female workers, for men and women only, for full-time workers only, where we additionally residualize on fixed effects for major industry group, where we include controls for 1 or 2 digit occupation fixed effects, where we control for the fraction of workers who are female in each occupation, and where we omit the largest 5% of occupations.

We see that, in Tables D.2 and D.3, our results in these extensions are generally very similar to our baseline findings in Table 4. A few specific estimated coefficients merit comment.

- When limiting the sample to full-time workers only, without task controls, the coefficient on hours worked is substantially larger than our baseline, but with task controls, substantially smaller than our baseline (with tasks). However, both of these coefficients have large standard errors. This reflects the fact that once part-time workers are removed from the sample, variation in average hours worked across occupations shrinks substantially.⁶ As a result, we cannot reject that these point estimates are different from our baseline estimates.
- Absent task controls, inclusion of 1 digit and 2 digit occupation fixed effects do reduce the estimated relationship between occupational hours and wages when, though the coefficient magnitudes remain sizable. However, when task controls are added, the magnitudes of these coefficients are not unusually different from other estimates. This simply reflects the fact that these fixed effects are (unsurprisingly) highly correlated with task variation across occupations.
- We note that when we residualize hours and wages across industries, we do observe somewhat smaller coefficient magnitudes, particularly when controlling for occupational tasks. This suggests that some of the relationship between hours and wages at the occupational level can be explained by industry features, such as productivity, production technologies or social norms. Whether or not this estimate of the hours-wage relationship is appropriate for assessing the return to

⁶ The fraction of workers working full-time in an occupation can explain 80% of the variation in hours worked across occupations.

expected hours depends on how a worker views industry in their job choice. If conditional on occupation and hours worked, industry is irrelevant to the worker, then controlling for industry will understate the return to hours. However, if a worker's willingness to work longer hours depends on the products produced by the firm she works in, as measured by industry, then this coefficient may be more suitable.

- Although the relationship between hours and wages does not vary much when considering men and women separately, or when controlling for the fraction of workers in each occupation who are female, these specifications may be difficult to interpret if the aim is to estimate the expected returns to hours. The challenge with is that these specifications implicitly control for selection of workers into occupations. For example, if men and women sort into occupations differentially on the basis of hours worked, then we would expect the fraction of workers who are female to be correlated with hours worked. In this case, including this control, though it does not affect our results, would partially control for an outcome of a worker's job choice, which could inaccurately represent the returns to hours when making job decisions. However, a key input to job choices may be the likelihood of working with certain types of colleagues, in which case controlling for fraction female in an occupation could be important for estimating the returns to expected hours. As a result, such specifications may be difficult to interpret.

Table D.1 - Regressions of Log(Wage) on Log(hours), ACS Individual-Level Data, Additional Specifications

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	College Educated	Less than College	FT Only	Ind Controls	Ind Controls
Log Hours	-0.117*** (0.002)	-0.086*** (0.003)	-0.146*** (0.003)	-0.215*** (0.005)	-0.011 (0.022)	-0.131*** (0.040)
Demo. Controls	Yes	Yes	Yes	Yes	Yes	No
Occ FE	Yes	Yes	Yes	Yes	No	Yes
Ind FE	No	No	No	No	Yes	Yes
N	633,927	279,330	354,597	554,917	633,927	633,927

All columns use standard errors clustered at the occupation level. Columns (1)-(4) and (6) include occupation fixed effects. Column (2) is for a sample of workers who have at least a college degree; Column (3) is for a sample of workers with less than a college degree. Column (4) is for a sample of workers who report working full-time, defined as at least 35 hours a week. Columns (5) and (6) include fixed effects for 16 major industry groups, as defined by the 2012 Census Industrial Classification System. Observations are weighted using perwt. Sample includes prime-age workers aged 25-55. Demographic controls include a quartic in age, indicators for race/ethnicity (black, Hispanic, Asian, other), sex, and indicators for educational attainment (HS degree, some college, college degree or more). Hours are derived from usual hours worked per week. Hourly wage is calculated as total reported wage and salary income for the prior year divided by the product of usual hours worked and weeks worked in the previous year. We trim wages that are below half of the federal minimum wage and inflate wages to 2012 dollars using the Bureau of Economic Analysis (BEA) National Income and Product Accounts Personal Consumption Expenditures.

Table D.2 - OLS Regressions of Ln(Wage) on Ln(Hours), 2016 ACS at the Occupational Level, Additional Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	Baseline	High Skill Occ	Low Skill Occ	Male Dom. Occ	Female Dom. Occ	Male Only	Female Only	FT Only	Indus. Resid.	1 Dig. Occ FE	2 Dig. Occ FE	Control Frac. Female	No Largest Occs
Avg. Log Hours	1.952*** (0.188)	1.764*** (0.296)	1.380*** (0.151)	2.192*** (0.324)	1.921*** (0.226)	2.252*** (0.215)	1.628*** (0.150)	2.578*** (0.639)	1.795*** (0.205)	1.599*** (0.189)	1.349*** (0.193)	2.017*** (0.192)	1.862*** (0.200)
Demo. Resid.	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ind. Resid.	No	No	No	No	No	No	No	No	Yes	No	No	No	No
1 Digit Occ FE	No	No	No	No	No	No	No	No	No	Yes	No	No	No
2 Digit Occ FE	No	No	No	No	No	No	No	No	No	No	Yes	No	No
Frac. Female Task Controls	No	No	No	No	No	No	No	No	No	No	No	Yes	No
N	474	237	237	237	237	474	471	474	474	474	474	474	450

All columns use robust standard errors. Regressions of occupation average of residualized log wage on occupation average of residualized log hours. Residuals are constructed by regressing log wage (hours) on demographic controls and a full set of occupation fixed effects. We use the coefficients on the occupation fixed effects as measures of average residualized log wage and log hours. Demographic controls used in the residualization include black, hispanic, asian, other race, hs only, some college, ba plus and a quartic in age. Person fixed effects are included using Outgoing Rotation Group samples from the CPS which include multiple observations per worker. Hours are usual hours worked per week. Hourly wage is calculated as total reported wage and salary income for the prior year divided by the product of usual hours worked and weeks worked in the previous year. We trim wages that are below half of the federal minimum wage and inflate wages to 2012 dollars using the Bureau of Economic Analysis (BEA) National Income and Product Accounts Personal Consumption Expenditures. Occupations are weighted by the occupation total of individual weights. Column (1): baseline (as in Table 4, column 2). Columns (2) and (3): only occupations above (2) or below (3) the median occupation's fraction of workers with a college degree. Columns (4) and (5): only occupations above (4) or below (5) the median occupation's fraction of workers who are male. Columns (6) and (7): construct residualized average hours and wages separately for men (6) or women (7). Column (8): construct residualized hours and wages only for full-time workers. Column (9): include 16 major industry groups (as defined by 2012 Census codes) in the residualization. Column (10) and (11): Include as controls fixed effects for either 7 major occupation groups (10), which we term "1 digit occupations," or 25 detailed occupation groups (11), which we term "2 digit occupations." Column (12): Include as a control the fraction of workers who are female in each occupation. Column (13): Drop occupations whose size is above the 95th percentile for all occupations.

Table D.3 - OLS Regressions of Ln(Wage) on Ln(Hours), 2016 ACS at the Occupational Level, Additional Specifications with Task Controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
	Baseline	High Skill Occ	Low Skill Occ	Male Dom. Occ	Female Dom. Occ	Male Only	Female Only	FT Only	Ind. Resid.	1 Dig. Occ FE	2 Dig. Occ FE	Control Frac. Female	No Largest Occs
Avg. Log Hours	0.925*** (0.198)	1.011*** (0.209)	1.270*** (0.206)	0.887*** (0.246)	0.988*** (0.231)	1.123*** (0.226)	0.804*** (0.150)	0.530 (0.417)	0.602*** (0.197)	0.932*** (0.192)	0.828*** (0.169)	0.882*** (0.210)	0.999*** (0.158)
Demo. Resid.	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ind. Resid.	No	No	No	No	No	No	No	No	Yes	No	No	No	No
1 Digit Occ FE	No	No	No	No	No	No	No	No	No	Yes	No	No	No
2 Digit Occ FE	No	No	No	No	No	No	No	No	No	No	Yes	No	No
Frac. Female Task Controls	No	No	No	No	No	No	No	No	No	No	No	Yes	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	467	231	236	235	232	467	464	467	467	467	467	467	443

See notes to Table D.2. Task data is constructed from the ONET 4.0 aggregated to the occ1990dd level. There are five occ1990dd occupations missing task data in the ONET 4.0 which yield missing task data for seven occs. Task controls include social skills, as defined in Deming (2017), and abstract analytical, manual and routine, as in Acemoglu & Autor (2011), and competitiveness as in Cortes and Pan (2018). See appendix A for an explanation how of these composites were created. Task measures are standardized to be mean zero and standard deviation one across the occ1990dd distribution.