Asymmetric Information in Labor Contracts: Evidence from an Online Experiment

Daniel Herbst
University of Arizona

Please click here for most recent version

PRELIMINARY AND INCOMPLETE: PLEASE DO NOT CITE OR DISTRIBUTE

2024-09-06

Abstract

For short-term workers facing uncertain output, hourly wage contracts provide implicit insurance compared to self-employment or piece-rate pay. But like any insurance product, these contracts are prone to market distortions through moral hazard and adverse selection. Using a model of wage contracts under asymmetric information, I show how these distortions can be identified as potential outcomes in a marginal-treatment-effects framework. I apply this framework to a field experiment in which data-entry workers are offered a choice between a randomized hourly wage and a standardized piece rate. Using experimental wage offers as an instrument for hourly wage take-up identifies moral hazard, while comparisons between piece-rate workers who declined different wage offers identifies adverse selection.

1 Introduction

The proliferation of "gig work" has transformed how millions of workers are paid (Garin et al., 2023; Collins et al., 2019; Katz and Krueger, 2019; Abraham et al., 2017; Jackson et al., 2017). Rather than clocking their hours, a growing workforce of short-term freelancers are compensated by the number of miles driven, pages written, or tasks completed. Compared to traditional, time-based labor contracts, this type of self employment can be risky and unpredictable—an hourly worker knows what they will earn from a day's work, whereas a gig worker's earnings often depend on uncertain factors like weather or traffic. But if workers value the insurance provided by hourly wages, why are so many short-term labor markets dominated by such risky compensation?

In this paper, I investigate how two distortionary forces—moral hazard and adverse selection—can lead to an inefficiently low provision of hourly work. On the one hand, if hourly wages reduce effort through moral hazard (Lazear, 2000), employers may opt for freelance hiring or piece-rate pay to increase productivity. On the other hand, if less productive workers adversely select into time-based wages, firms might forego hourly employment contracts to avoid attracting the wrong type of worker. Distinguishing the relative magnitudes of these two effects has important policy implications. For example, government mandates or subsidies that promote hourly wages might mitigate welfare loss from adverse selection, but are unlikely to induce the added worker effort necessary to alleviate moral hazard.

I formalize this intuition with a model of short-term labor markets under asymmetric information. Using this framework, I show how the provision of hourly employment contracts is determined by two curves: a worker's reservation wage—the minimum they will accept for an hour of labor—and the average value of output among workers with comparatively lower reservation wages. An hourly worker can be profitably hired only if their average value exceeds their reservation wage. Using hourly wage *offers* as an instrumental variable, I show how this model has a direct correspondence to the marginal-treatment effects framework (Björklund and Moffitt, 1987; Heckman and Vytlacil, 1999, 2005, 2007). Similar

to the health-insurance context (Kowalski, 2023b), this correspondence reveals how both moral hazard and adverse selection can be expressed as functions of "marginal values"—the potential output values of workers with a given reservation wage.

Motivated by my theoretical framework, I conduct a two-stage field experiment designed to separately identify moral hazard and adverse selection into hourly wages. First, I offer workers a choice between a randomized hourly wage and a standardized piece rate in exchange for performing a data-entry task. Then, after workers choose a payment option but before they begin the task, I increase hourly wages for a randomized subset of those who accepted hourly offers, bringing them to parity with higher hourly wage offers. Using the different initial wage offers as an instrument for accepting the same hourly contract allows me to identify the moral hazard effect of time-based compensation independently of any behavioral response to higher effective wages. Meanwhile, comparisons across piece-rate workers who decline different hourly wage offers identifies selection—both groups work under the same form of compensation but faced different ex-ante menus of contracts. Importantly, experimental wage offers include contracts that may not be profitable to a real-world employer, avoiding the "under-the-lamppost" problem inherent to many empirical studies of information asymmetries (Einav et al., 2010b).

This study relates to several streams of existing research. Building upon the seminal work of Akerlof (1970), several studies have applied the theory of adverse selection to labor markets, showing how the self-sorting of workers by unobserved productivity can lead to inefficient hiring, compensation, or other labor contract provisions (Weiss, 1980; Malcomson, 1981; Jovanovic, 1982; Greenwald, 1986; Lazear, 1986; Gibbons and Katz, 1991; Levine, 1991; Kugler and Saint-Paul, 2004; Moen and Rosen, 2005). While empirical analyses of these selection phenomena are less common, several papers document differential sorting and/or treatment effects of job characteristics like compensation schemes (Shearer, 1996; Lazear, 2000; Angrist et al., 2021; Shearer, 2004; Kantarevic and Kralj, 2016) or remote work (Emanuel and Harrington, 2024). These and other studies use observational data to estimate selection and incentive effects in both labor markets and elsewhere (Einav

et al., 2010a; Chiappori and Salanie, 2000; Hendren, 2017), but few have used experimental methods to form these estimates. Karlan and Zinman (2009), which randomizes contract offerings on microfinance loans in South Africa, is a notable exception. Using a design similar to the second stage of my experiment, they isolate selection on unobservables by comparing borrowers who faced different menus of options but ultimately faced the same contract terms. They find strong evidence of moral hazard and weaker evidence of adverse selection.

My study design also builds upon existing experimental work using online freelancer platforms. DellaVigna and Pope (2018) and DellaVigna and Pope (2022) estimate the effects of both monetary and non-monetary incentives in an online typing task. Pallais (2014) and Pallais and Sands (2016) demonstrate the signaling benefit of entry-level hiring and employer references using online experiments. My paper complements such findings, as the unraveling phenomenon I investigate could help explain why Spence (1973)-style information treatments provide a net social benefit.

Relative to existing work on adverse selection in labor markets, my study offers several distinct advantages. First, my experiment would reveal workers' decisions over contracts that are unavailable to them in the real world, allowing me to quantify welfare losses in unraveled markets where efficient wage contracts cannot be observed. Second, the ability to randomize wage offers allows me to reliably separate selection-on-unobservables from the potential treatment effects of different compensation schemes. Finally, my empirical framework maps the insurance literature to a wage-setting environment and shows how the welfare impact of information asymmetries can be identified through marginal-treatment effects estimation.

The rest of this paper proceeds as follows: In Section 2, I develop a model of hourly wage contracts under asymmetric information. In Section 3, I describe my experiment and underlying empirical strategy. In Section 4, I discuss results. Section 5 concludes.

2 Model of Asymmetric Information in Wage Contracts

In this section, I present model of short-term labor markets under asymmetric information. The model borrows from Einav et al. (2010a) and Herbst and Hendren (2021), who develop models of asymmetric information in health insurance markets and college financing markets, respectively. I then show how the parameters of this model can be mapped into marginal-treatment effects framework, motivating my experimental design in Section 3.

Consider a perfectly competitive labor market in which risk-neutral firms face a population of observably equivalent workers.[1] Each worker, i, can produce some level of hourly output, $q_i = f(\zeta_i, e_i, \nu_i)$, which is a function of unobserved worker characteristics (ζ_i), individual effort (e_i), and random noise (ν_i). Firms can buy worker output at a constant market price of p per unit.[2] Alternatively, they can offer a flat, up-front wage of their own choosing, w, in exchange for a claim on a worker's hourly output, q_i.[3]

For individual i, I define the reservation wage, $\overline{w}_i$, as the minimum w at which they would accept an hourly contract. The relative supply of hourly workers is given by

$$S(w) \equiv \Pr(\overline{w}_i < w). \tag{1}$$

Assuming strict monotonicity ($S(w) > S(w')$ for all $w > w'$), I index workers by a type parameter, $\theta_i \in [0, 1]$, equal to the share of the worker population willing to accept a lower wage than worker i's reservation wage, $\theta_i \equiv S(\overline{w}_i)$. I can then rewrite a worker's reservation wage as a function their type, $\overline{w}_i = \overline{w}(\theta_i)$, where

$$\overline{w}(\theta) \equiv S^{-1}(\theta). \tag{2}$$

[1] I focus on perfect competition because it serves as a useful benchmark for welfare calculation. It is straightforward to adapt the model to alternative market structures, including those in which employers hold monopsony power.

[2] Directly purchasing a worker's product of labor can be thought of as either piece-rate employment or hiring a self-employed contractor.

[3] While a worker's output, q_i, can vary between hourly and piece-rate contracts (e.g., through moral hazard effects), I assume an hourly worker's output does not vary with the wage, w (i.e., no wage effects). I relax this assumption in Section 2.2.

Facing this population of observably identical workers with unknown types, employers set wages to maximize profits. I define the *marginal value* of type θ as

$$MV(\theta) \equiv E[Y_i|\theta_i = \theta], \tag{3}$$

where $Y_i = pq_i$, the incremental value of output q_i produced under an hourly contract.[4] Absent any behavioral response to the hourly contract, $MV(\theta)$ equals type θ's expected earnings under the market piece rate, p. If θ were risk averse, we would expect their reservation wage to fall below this "actuarially fair" wage (i.e., $\overline{w}(\theta) < MV(\theta)$). In other words, they would accept lower expected earnings in exchange for the implicit insurance provided by hourly wages relative to piece rates. In this case, a fully informed employer could profit from offering an hourly wage of $w = \overline{w}(\theta)$ exclusively to type θ.

However, if employers cannot observe types, they cannot prevent borrowers with $\theta_i \neq \theta$ from opting into a contract offered at wage w. In this case, the hourly position would be accepted by all types θ_i such that $\overline{w}(\theta_i) \leq w(\theta)$. So instead of obtaining type θ's marginal value, $MV(\theta)$, the employer would obtain their *average value*, defined as

$$AV(\theta) \equiv E[Y_i|\theta_i \leq \theta]. \tag{4}$$

The average value, $AV(\theta)$, of type θ is given by the average value of output produced among all types $\theta_i < \theta$. When we account for this selection into contracts, the employer's profits are given by

$$\Pi(w) = S(w)(AV(\theta^w) - w), \tag{5}$$

where $\theta^w \equiv S(w)$, the worker type with reservation wage equal to w.

I assume that at least one worker's marginal value exceeds their reservation wage, ($\overline{w}(\underline{\theta}) > MV(\underline{\theta})$ for some $\underline{\theta} > 0$). This assumption will hold unless all workers are risk-

[4] Y_i reflects the market value of q_i units of output, or, equivalently, the amount the firm saves by not buying hourly worker i's output at the piece rate. This measure of incremental value is analogous to the incremental cost of insurance defined in Einav et al. (2010a).

loving, risk-neutral, or hold over-optimistic beliefs about their productivity. I further assume that $MV(\theta)$ crosses the supply curve at most once (if $\overline{w}(\overline{\theta}) < MV(\overline{\theta})$ for some $\overline{\theta}$, then $\overline{w}(\theta) < MV(\theta)$ for all $\theta > \overline{\theta}$). With these simplifying assumptions in hand, the equilibrium condition for the share of workers under hourly contracts, θ^{EQ}, is given by

$$\overline{w}(\theta^{EQ}) = AV(\theta^{EQ}). \tag{6}$$

In equilibrium, firms offer wage contracts up the point where the marginal worker's reservation wage, $\overline{w}(\theta^{EQ})$, is exactly equal to the average value of their hourly employees' output, $AV(\theta^{EQ})$.

Figure 1 illustrates the welfare impacts of adverse selection for an example population. An efficient allocation of contracts would lead to hourly employment for all types $\theta \leq \theta^{EF}$, as these workers would accept wages at or below their marginal values ($\overline{w}(\theta) \leq MV(\theta)$). But while type θ^{EF}'s reservation wage (red line) is equal to their marginal value (blue line), an employer offering an hourly wage of $w = \overline{w}(\theta^{EF})$ would only recoup the average value (green line) among everyone accepting the offer (i.e., all $\theta \leq \theta^{EF}$). The employer could lower their wage offer, but that would drive those with the highest productivity out of the market, further reducing the contract's average value. This process continues across all types for whom $\overline{w}(\theta) > AV(\theta)$, so that the equilibrium share of workers under hourly contracts is θ^{EQ}, where $\overline{w}(\theta^{EQ}) = AV(\theta^{EQ})$. In this stylized example, roughly one-third of the population—$\theta \in (\theta_{EQ}, \theta_{EF})$—cannot obtain hourly employment despite a willingness to work for less than their expected earnings under the market piece rate. The result is a welfare loss corresponding to the area of the region shaded in pink, which is equal to

$$DWL = \int_{\theta_{EQ}}^{\theta_{EF}} \left(MV(\theta) - \overline{w}(\theta)\right) d\theta. \tag{7}$$

In summary, private knowledge of productivity can create a gap between the marginal and average values of labor, preventing Pareto-improving exchanges of hourly wage contracts—workers are paid by the hour if and only if their reservation wage is no higher than the

average value of those with lower reservation wages. This information asymmetry reduces total welfare below what it would be under a full-information benchmark.

2.1 Incorporating Moral Hazard

Note that the model above allows for moral hazard effects, even if those effects are not explicitly discussed. To see how, consider worker i's potential output values under counterfactual contracts. Specifically, let Y_{1i} denote i's output value if they work under the hourly wage, and let Y_{0i} denote their output value if they work under the piece rate. The moral hazard effect for worker i is given by their individual treatment effect of the hourly wage, $MH_i \equiv Y_{1i} - Y_{0i}$.[5] However, since piece-rate workers sell their output at a constant price per unit, their productivity per hour has no affect on firm profits. So while firms care about a worker's output under the hourly contract (Y_{1i}), they don't care how this output compares to the piece-rate counterfactual (Y_{0i}). As a result, $AV(\theta)$ and $MV(\theta)$ are defined conditional on accepting the hourly contract, and thus depend only on output under hourly wages, Y_{1i}. The profit condition (5) and welfare calculation (7) are therefore inclusive of any incentive effects.

While not strictly necessary to calculate welfare loss, explicitly modeling and estimating moral hazard effects can nonetheless be useful, especially if firms have ways of mitigating the incentive response to hourly wage contracts. For example, a firm might combine hourly wages with a smaller piece-rate portion to ensure workers have some "skin in the game," similar to restaurant tipping or sales commissions. This type of compensation would likely attenuate the disincentive effects of hourly pay but do little to prevent adverse selection—low-productivity workers would still prefer the partial insurance of mixed compensation compared to a pure piece rate. This scenario can easily incorporated into my framework—it simply requires reframing the model as a market for supplemental hourly wages on top of a preexisting piece rate. However, to identify the model under this counterfactual, I must explicitly separate selection from treatment effects under the "pure" hourly wage offers in

[5]Strictly speaking, $Y_{1i} - Y_{0i}$ captures worker i's overall output response to the hourly wage contract. This response could result from behavioral phenomena not traditionally classified as "moral hazard."

my experiment.

To separate the incentive response of the contract from selection on underlying unobservables, I must consider the value of worker output across types under hourly wage and both with *and without* the hourly wage contract. To do so, I split Equations (3) and (4) into two pairs of curves. First, I define marginal values of a type θ as the conditional means of potential output value with (Y_{1i}) and without (Y_{0i}) the hourly wage:

$$MV_1(\theta) \equiv E[Y_{1i}|\theta_i = \theta] \tag{8}$$

$$MV_0(\theta) \equiv E[Y_{0i}|\theta_i = \theta] \tag{9}$$

Note that $MV_1(\theta)$ is simply a relabeling of $MV(\theta)$ from Equation (3)—it captures the expected output value under hourly wage $w = S^{-1}(\theta)$ for the worker who is indifferent between accepting or declining the offer. $MV_0(\theta)$, on the other hand, captures the expected output value of that same worker if they had instead rejected wage offer w and remained on the piece rate.[6] The difference between these two marginal value curves identifies the moral hazard effect for a given type:

$$MH(\theta) \equiv MV_1(\theta) - MV_0(\theta). \tag{10}$$

Similarly, the average value curve can be split into two counterfactuals:

$$AV_1(\theta) \equiv E[Y_{1i}|\theta_i \leq \theta] \tag{11}$$

$$AV_0(\theta) \equiv E[Y_{0i}|\theta_i \leq \theta]. \tag{12}$$

$AV_1(\theta)$ is equivalent to $AV(\theta)$ from Equation (4); it equals the average value of output among hourly-pay workers with lower reservation wages than type $\overline{w}(\theta)$. $AV_1(\theta)$, on the other hand, equals the average value among those same workers if they had instead worked

[6] In a loose sense, these two curves can be thought of as bounds. If firms have some way of mitigating moral hazard, the true marginal value curve would lie somewhere between $MV_0(\theta)$ and $MV_1(\theta)$.

under a piece rate.

Correspondence to Marginal Treatment Effects With the introduction of potential output under different compensation schemes, this framework has a useful mapping to the causal-inference literature. Treating the wage offer w as a continuous instrument, marginal values among workers of a given type are equivalent to mean potential outcomes for a given "resistance to treatment," where "treatment" is take-up of the hourly contract, and "resistance" is quantile reservation wage, $\theta_i \equiv S(\overline{w}_i)$. The moral hazard effect in Equation 10 is therefore equivalent to the marginal treatment effect of the hourly contract, $MTE(\theta) \equiv E[Y_{1i} - Y_{0i}|\theta_i = \theta]$ (Björklund and Moffitt, 1987; Heckman and Vytlacil, 1999, 2005, 2007). It measures the average effect of treatment (hourly contract take-up) among those whose resistance to treatment (quantile reservation wage, $\theta_i \equiv S(\overline{w}_i)$) is equal to a given propensity score (share of hourly workers, $\theta = S(w)$).[7] This correspondence is illustrated in Figure 2.

2.2 Wage Effects

The framework above allows workers' expected output to vary between hourly versus piece-rate compensation. It does not, however, allow that output to vary with the wage level under an hourly contract. In other words, it ignores any potential wage effects that higher hourly compensation might have on worker output. I can, however, incorporate wage effects into the model by allowing each worker's potential output under the hourly contract to vary with the wage (i.e., $Y_{1i} = Y_{1i}(w)$). With this added dimension to potential outcomes, I rewrite $AV_1(\theta)$ as the average value of output among lower types *at θ's reservation wage*:

$$AV_1^E(\theta) \equiv E\left[Y_{1i}(\overline{w}(\theta))|\theta_i \leq \theta\right]. \tag{13}$$

Assuming wage effects are weakly positive and non-decreasing in θ, the equilibrium condition is given by $\overline{w}(\theta^{EQ}) = AV_1^E\left(\theta^{EQ}\right)$. In this case, firms pay an hourly wage equal to the

[7] Kowalski (2023b) makes an analogous connection between the Einav et al. (2010a) model and marginal treatment effects in the health insurance context.

average value of accepting workers' output *under that wage*, $AV^E(\theta^{EQ})$. Relative to the benchmark model, positive wage effects will therefore push the average value curve upwards and increase the share of hourly contracts under asymmetric information.

Note, however, that the efficient equilibrium—the one that would exist in a full-information counterfactual—is also complicated by the presence of wage effects. A fully-informed firm may benefit from paying a worker above their reservation wage if their expected increase in output exceeds the wage premium (i.e., if $E\left[Y_{1i}(w) - Y_{1i}(\overline{w_i})|\theta_i = \theta\right] > w - \overline{w}(\theta)$ for some w).[8] I thus rewrite $MV_1(\theta)$ as the marginal value of type θ's output *at their profit-maximizing wage*, so

$$MV_1^E(\theta) \equiv E\left[Y_{1i}(w^*(\theta))|\theta_i = \theta\right], \tag{14}$$

where

$$w^*(\theta) \equiv \underset{w}{\text{argmax}}\, E\left[Y_{1i}(w) - w|\theta_i = \theta\right]. \tag{15}$$

Note that allowing for wage effects means I can no longer interpret Equation 10 as the marginal treatment effect of hourly-contract take-up—if the wage level influences worker output independently of the hourly compensation structure, the wage-offer instrument no longer satisfies the exclusion restriction. The randomized wage raises in my experimental design eliminate this concern. By equalizing the paid wages of low-offer accepters with those of high-offer accepters, these surprise wage increases isolate variation in *offered* wages conditional on a given *effective* wage. I can therefore identify the marginal treatment effect of being paid a given hourly wage among those indifferent to a particular wage offer. I discuss this instrument validity and estimation of wage effects in Section 3.3.

[8] I avoid the term "efficiency wages," which refers to a class of models explaining unemployment as a general-equilibrium consequence of firms' strategic wage-setting behavior (Weiss, 2014; Krueger and Summers, 1988; Yellen, 1984). In many efficiency-wage models, above-market wages are driven not by causal effects of wages on productivity, but by worker selection, firms' monitoring ability, or turnover costs (Salop, 1979; Weiss, 1980). Other explanations, like unemployment avoidance (Shapiro and Stiglitz, 1984) or employee loyalty (Akerlof, 1982) are less applicable to short-term labor markets.

3 Experimental Design

In this section, I describe my experimental design and empirical strategy. The goal of my experiment is to two-fold: First, I aim to identify the incentive effects of hourly wage contracts on worker performance (moral hazard). Second, I want to identify how workers with different unobserved productive potentials self-select into these contracts (adverse selection). Separately identifying these forces poses an empirical challenge—differences in realized output between workers who opted into a given wage offer reflect both the ex-ante productivity differences between those self-selected groups and the causal effect of the different wage offers they chose.

To overcome this challenge, my experimental design offers workers a choice between a randomized hourly wage and a standardized piece rate. Comparing realized output between individuals who faced different hourly wage offers but ultimately chose the common piece rate identifies adverse selection—both groups ultimately face the same compensation scheme but made decisions under different alternative options. So, if workers choose contracts based on their privately known productivities, those who decline more generous hourly payments should perform better than those foregoing more modest wages. At the same time, because I observe worker output under both contract choices in each treatment group, a standard two-stage-least-squares estimation allows me to separately identify treatment effects of hourly wages among those who accept the offer.

To formalize this intuition, consider a potential outcomes framework in which some worker i chooses between two mutually exclusive contracts—a piece rate (p) and an hourly wage (w). As in Section 2, let Y_{1i} denote i's output if they work under the hourly wage, and let Y_{0i} denote their output if they work under the piece rate. Given these potential outcomes, worker i's observed output, Y_i, is given by

$$Y_i = D_i Y_{1i} + (1 - D_i) Y_{0i}, \tag{16}$$

where D_i is a binary indicator for whether i chooses the hourly wage. A comparison of

realized outputs between hourly ($D_i = 1$) and piece-rate ($D_i = 0$) workers would yield the following:

$$E\left[Y_i|D_i=1\right]-E\left[Y_i|D_i=0\right] = \underbrace{E\left[Y_{1i}-Y_{0i}|D_i=1\right]}_{\text{Treatment on the Treated } (TOT)} + \underbrace{E\left[Y_{0i}|D_i=1\right]-E\left[Y_{0i}|D_i=0\right]}_{\text{Selection into Treatment } (\Delta\mu_0)}. \tag{17}$$

This difference is the sum of two components. The first is the treatment-on-the-treated effect, TOT, which equals the average effect of hourly pay among those who accept wage offer w over the piece rate. The second is the selection into treatment, $\Delta\mu_0$, which equals the average difference in potential outcomes under the piece rate between those choosing hourly pay ($D_i = 1$) and those choosing the piece rate ($D_i = 0$). These components are difficult to separate because piece-rate outcomes among hourly workers ($Y_{0i}|D_i = 1$) are always unobserved.

Now suppose that, rather than face a uniform menu of piece rate and hourly wage offers, workers are randomly assigned to one of two offer conditions, $Z_i \in \{0, 1\}$. Only workers assigned to $Z_i = 1$ are offered the choice between the piece rate and hourly wage, while workers assigned to $Z_i = 0$ are paid the piece rate with no alternative. Comparing worker output across these two treatment-offer groups and scaling by the hourly-wage take-up rate yields the classic TOT estimator from Wald (1940):

$$\begin{aligned} TOT &\equiv E\left[Y_{1i}-Y_{0i}|D_i=1\right] \\ &= \frac{E\left[Y_i|Z_i=1\right]-E\left[Y_i|Z_i=0\right]}{\pi}, \end{aligned} \tag{18}$$

where $\pi \equiv Pr\left(D_i = 1|Z_i = 1\right)$, the probability an individual accepts the hourly contract conditional on an offer.

In the context of this paper, however, the selection component $\Delta\mu_0$ from Equation (17) is equally as important as treatment effects. I can identify this component by simply comparing output between piece-rate workers in the control group ($Z_i = 0$) and piece-rate

workers in the hourly-offer group ($Z_i = 1$) who declined the hourly wage offer:

$$\begin{aligned} \Delta\mu_0 &\equiv (E[Y_{0i}|D_i = 1, Z_i = 1] - E[Y_{0i}|D_i = 0, Z_i = 1]) \\ &= \frac{E[Y_i|Z_i = 0] - E[Y_i|D_i = 0, Z_i = 1]}{\pi}, \end{aligned} \tag{19}$$

where equality follows from randomized assignment.[9] Figure 3 provides a graphical illustration of the intuition from Equation (19). The control group, by construction, is subject to the standardized piece rate, while the treatment-offer group is offered an hourly wage as an alternative. Selection is identified by comparing the control group ($D_i = 0, Z_i = 0$) to those in the treatment group ($D_i = 0, Z_i = 1$) who chose to remain on the piece rate. This selection-on-unobservables estimator captures the average difference in potential untreated outcomes for "compliers" versus "never-takers" (Black et al., 2022; Kowalski, 2023a; Huber, 2013).

3.1 Identifying Wage Effects with Multiple Treatment Offers

The example above simulates a simplified version of my experimental design with a binary treatment assignment, $Z_i \in \{0, 1\}$. In practice, however, my experiment features several experimental groups facing different hourly wage offers. Multiple experimental offers not only allows me to identify non-linearities in treatment and selection into hourly wages, but also enables me to identify wage effects of different payment levels, conditional on an hourly pay structure.[10] I identify these wage effects by adding an additional dimension of randomization in the spirit of Karlan and Zinman (2009). After workers choose their compensation option, but before they begin the task, I randomly increase hourly wages for a random subset of those accepting lower wage offers, bringing their them to parity

[9] Randomized assignment implies $E[Y_{0i}|Z_i = 1] = E[Y_{0i}|Z_i = 0] = E[Y_i|Z_i = 0]$, so $E[Y_{0i}|D_i = 1, Z_i = 1] = \frac{E[Y_i|Z_i=0]-(1-\pi)E[Y_i|D_i=0,Z_i=1]}{\pi}$. Equation (19) can also be derived by subtracting the Wald estimator (18) from the difference in hourly versus piece-rate outcomes in the treatment-offer group, $E[Y_i|D_i = 1, Z_i = 1] - E[Y_i|D_i = 0, Z_i = 1]$.

[10] Separating these two forces is important to satisfy the exclusion restriction. If higher hourly wages induce greater productivity through increased motivation or satisfaction, my estimates would be biased, as wage offers would correlate with potential outcomes under hourly wages ($w_i \not\perp Y_{1i}$).

with higher treatment-offer groups. This design creates random variation in *offered* wages among workers of a given *effective* wage, allowing me to separate potential wage effects from moral hazard and adverse selection: Wage effects can be identified by comparing groups that accepted the same wage offers but face different contracted wages, moral hazard can be identified by instrumenting for hourly take-up with offered wage holding the effective wage constant, and selection on wage offers can be estimated by comparing groups that faced different ex-ante wage offers but ultimately worked under the same contract wages.

Figure 4 illustrates the variety of treatment and selection margins I can investigate with multiple wage offers in a two-stage design with multiple treatment arms. The top row of boxes represents individuals in each of the three experimental groups who remain on the piece rate. Because all three groups face the same ex-post payment terms but different ex-ante wage offers, comparisons between them isolates worker selection on underlying productivity under the piece rate. The bottom two boxes represent workers who opted into low and high hourly wages, respectively. In the second stage of the experiment, a random subset of those accepting the low hourly wage are promised an additional top-up compensation before they begin working on the task. This surprise top-up equalizes their effective wage with that of the high-offer group, allowing me to separate wage effects (diagonal arrow) from underlying differences in productivity under the high hourly wage.[11]

3.2 Estimating Marginal Values

The example above illustrates how to estimate average treatment and selection among those induced into hourly pay by an experimental wage offer. However, my model in Section 2 demonstrates how the welfare effects of asymmetric information rely on the distribution of *marginal* treatment and selection across a range of wages. In this section, I show how treating multiple experimental wage offers as continuous instrument in a marginal-treatment-effects framework identifies these marginal values.

To see how a continuous wage-offer instrument identifies marginal outcomes from my

[11] A formal proof of this result is provided in Appendix B

model, first note that the supply curve in equation (1) can be straightforwardly identified as the share of accepters across wage offers:

$$S(w) \equiv \Pr\left(\overline{w}_i < w\right) = \Pr(D_i = 1 | w_i = w). \tag{20}$$

I use a logit regression to estimate $S(w)$. Predicted values from this regression allow me to which allows me estimate marginal and average values for a given type, $\theta \equiv S(\overline{w})$, under both hourly and piece-rate contracts. The (potential) average value curve under the hourly wage can be identified using average output value among those who accept the hourly wage offer that induces θ-share of workers into the hourly contract:

$$AV_1\left(\theta_w\right) \equiv E\left[Y_{1i} | \theta_i \leq \theta\right] = E\left[Y_i | S(w_i) = \theta, D_i = 1\right]. \tag{21}$$

Likewise, the (potential) average value curve under the piece rate can be identified using average output value among workers *declining* the wage offer that is accepted by θ-share of workers:

$$AV_0\left(\theta_w\right) \equiv E\left[Y_{0i} | \theta_i \leq \theta\right] = \frac{E\left[Y_i | S(w_i) = 0\right] - (1-\theta) E[Y_i | S(w_i) = \theta, D_i = 0]}{\theta}, \tag{22}$$

where $E\left[Y_i | S(w_i) = 0\right]$ is the average output value of workers in the control group, who all work under the piece rate.

Finally, marginal values can be identified by separately differentiating take-up weighted conditional means for decliners and accepters of each offer:

$$MV_1(\theta) = \frac{\partial\left(E\left[Y_{1i} | \theta_i \leq \theta\right]\theta\right)}{\partial\theta} = \frac{\partial\left(E\left[Y_{1i} | S(w_i) = \theta, D_i = 1\right]\theta\right)}{\partial\theta} \tag{23}$$

$$MV_0(\theta) = -\frac{\partial\left(E\left[Y_{0i} | \theta_i > \theta\right](1-\theta)\right)}{\partial\theta} = -\frac{\partial\left(E\left[Y_i | S(w_i) = \theta, D_i = 0\right](1-\theta)\right)}{\partial\theta}. \tag{24}$$

Equations (21) through (24) depend on $E[Y_i | S(w_i) = \theta, D_i = 1]$, $E[Y_i | S(w_i) = \theta, D_i = 0]$, and their derivatives with respect to θ. Following Carneiro et al. (2011), I estimate

these objects using local polynomial regression of Y_i on $S(w_i)$ with a bandwidth of 0.2.[12] Using this semi-parametric method allows me to flexibly estimate marginal and average value curves, which is critical to accurately estimate welfare consequences of asymmetric information. Standard errors are calculated using bootstrap replications.

Welfare Impact Once I have estimated $\overline{w}(\theta)$, $MV(\theta)$, and $AV(\theta)$ curves, it is straightforward to calculate the welfare loss from Equation (7). First, I calculate equilibrium (θ^{EQ}) and efficient (θ^{EF}) shares of hourly wages using the intersection of $\overline{w}(\theta)$ with $AV(\theta)$ and $MV(\theta)$, respectively. Then, I calculate the cumulative difference in $\overline{w}(\theta)$ and $MV(\theta)$ over the region $\theta \in \left(\theta^{EQ}, \theta^{EF}\right)$. This calculation measures lost welfare as the implied "risk discount" workers would be willing to accept for the implicit insurance provided by hourly wages. I can then divide this welfare loss by the piece rate, $p = \$0.05$ to get the welfare loss per dollar of worker output.[13]

3.3 Estimating Wage Effects

The framework above uses experimental wage offers, w, as an instrument for take-up of an hourly contract, $D_i = 1$. However, as I mentioned in Section 2.2, if higher hourly wages induce greater productivity through increased motivation or satisfaction, my estimates would be biased; wage offers would correlate with potential outcomes under hourly wages ($w_i \not\perp\!\!\!\perp Y_{1i}$) violating the exclusion restriction.

I separate hourly incentive effects from wage effects by adding an additional dimension of randomization in the spirit of Karlan and Zinman (2009). After workers choose their compensation option, but before they begin the task, I randomly increase hourly wages for a random subset of those accepting lower wage offers, bringing their them to parity

[12] For those who accept the hourly contract ($D_i = 1$), I first residualize effective hourly wage from Y_i using double-residual regression methods (Robinson, 1988) to prevent potential wage effects from violating the exclusion restriction for the wage-offer instrument. These wage effects are identified using randomized top-up payments in the second round of my experiment, as described in Section 3.3.

[13] If workers' preferences exhibit constant relative risk aversion, their marginal values and reservation wages are proportional the per-unit price, p, of their labor product. Under this assumption, per-dollar estimates of welfare loss can therefore be extrapolated to piece rates other than the \$0.05-per-entry benchmark in my experiment. Proof in Appendix A.

with higher treatment-offer groups. This design creates random variation in *offered* wages among workers of a given *effective* wage, allowing me to separate potential wage effects from moral hazard and adverse selection: Wage effects can be identified by comparing groups that accepted the same wage offers but face different contracted wages, moral hazard can be identified by instrumenting for hourly take-up with offered wage holding the effective wage constant, and selection on wage offers can be estimated by comparing groups that faced different ex-ante wage offers but ultimately worked under the same contract wages.

Figure 4 illustrates the variety of treatment and selection margins I can investigate with multiple wage offers in a two-stage design with multiple treatment arms. The top row of boxes represents individuals in each of the three experimental groups who remain on the piece rate. Because all three groups face the same ex-post payment terms but different ex-ante wage offers, comparisons between them isolates worker selection on underlying productivity under the piece rate. The bottom two boxes represent workers who opted into low and high hourly wages, respectively. In the second stage of the experiment, a random subset of those accepting the low hourly wage are promised an additional top-up compensation before they begin working on the task. This surprise top-up equalizes their effective wage with that of the high-offer group, allowing me to separate wage effects (diagonal arrow) from underlying differences in productivity under the high hourly wage.[14]

3.4 Setting and Implementation

Participants in my experiment are recruited on Prolific, an online platform that allows clients to hire online workers for short-term tasks.[15] The job posting for my experiment offers participants a $1.00 reward for transcribing handwritten text into typed form for five minutes. Such transcription tasks are commonly requested on Prolific and other online platforms. The posting also informs participants they "can earn an additional $0.03 in bonus compensation for each correctly typed sentence." A screenshot of the job posting is

[14] A formal proof of this result is provided in Appendix B

[15] Douglas et al. (2023) finds that the Prolific platform compares favorably to Amazon Mechanical Turk ("MTurk") and other platforms across several dimensions of data quality.

provided in Figure 5A.

Workers could only see my experimental job posting if they met the following screening criteria: (1) located in the United States, (2) spoke fluent English, (3) successfully completed 10 or more previous tasks, and (4) earned an approval rate above 98 percent on previous tasks. These screening criteria accomplish two things: first, they allow me isolate "professional" online workers who regularly perform tasks to earn income, as opposed to casual users who may take the tasks less seriously. Second, they restrict the sample to workers who are observably equivalent to a hypothetical employer. The goal of my experiment is to identify selection on *private* information, so I want to remove any selection on screenable characteristics.

Workers who accept the job posting are taken to an external link to perform the task.[16] After clicking this link, workers are randomized into one of eighteen experimental groups. Each group is offered a different menu of bonus compensation options for completing the ten-minute data-entry task.

In the first treatment group, participants are offered a choice between a flat bonus of $0.10 for completing the task or a piece rate of $0.03 per correctly typed sentence. In the second treatment group, participants are offered a choice between a flat $0.15 bonus or the same $0.03 piece rate. Additional treatment groups follow the same structure, increasing the flat bonus offer by multiples of $0.05, up to a maximum of $1.75. A control group is offered the $0.03 piece rate for each correctly typed sentence, with no alternative option. Experimental conditions are summarized in Table 1.

After receiving detailed instructions on the data-entry task, treated workers are presented with their group's bonus options in randomized order, as shown in Figure 5B. Once workers choose their compensation scheme, they are brought to a new page that states,"For performing this task, you will receive $1.00, plus your chosen bonus of [*bonus choice*]." A random 50 percent of workers who chose lower-valued bonus options receive a modified message that increases their base payment by enough to equalize their total compensation

[16]The task is hosted on the Qualtrics platform. Readers can view and perform a replication of the task here.

with the most generous offer ($1.00 + $1.75 = $2.75). For example, half of those who select the $0.25 bonus are told "you will receive $2.50, plus your chosen bonus of $0.25."

Once participants are notified of their bonus compensation and click "Begin Task," they are presented with a handwritten sentence and a text box. The worker types a sentence in the box and clicks the "Next" button, bringing them to a new page with a different sentence. This process continues for five minutes. Worker output is validated in real time, so workers can see a running tally of their "score" (the number of correctly typed sentences) and their bonus earnings in the lower-left corner of each page. Workers also see a countdown timer displaying the number of minutes and seconds remaining in the task.[17] When the timer reaches zero, the screen refreshes to an end-of-task page displaying a performance summary and a unique survey code to redeem their earnings. Workers are paid the $1.00 reward plus any bonus earnings within 24 hours of completing the task. Figure 6 provides a timeline of the experimental protocol.

Importantly, requesters on the Prolific platform have the ability to reject or approve a given worker's assignment. Rejected assignments do not earn rewards and lower workers' approval ratings. The reputational damage from rejected assignments is a salient concern among workers on Prolific and similar platforms. As in most labor markets, this threat of rejection threat creates an incentive for online workers to maintain a minimum standard of performance, even if they are paid a flat hourly wage.

4 Results

This section describes early results from the experiment, which is currently in the field.

The bar chart in Figure 7 shows the share of borrowers in each treatment group who accepted their hourly wage offer instead of the $0.03 piece rate. Unsurprisingly, the relative supply of hourly workers increases with the offered wage.

[17] Figure 5C provides a screenshot of the task. The display and submission methods for this task designed to prevent workers from cheating through automation software or "bots." While it is possible that some participants may have tried to make use of such software, performance statistics suggest any such attempts were unsuccessful at increasing output.

Tables 3 through 5 report reduced-form estimates of selection and treatment effects of hourly wages, respectively.

Figures 8 through 10 plot semiparametric estimates of supply and value curves under both hourly wage and piece-rate counterfactuals. Value curves are estimated using a second-degree local polynomial regression of residualized hourly output value against predicted hourly supply. Shaded regions represent 95% confidence intervals.

4.1 External Validity

The design of this experiment raise several potential concerns regarding external validity. First, my experimental setting and task may not generalize to other contexts. For example, the pattern of selection on data-entry skills likely differs from how delivery workers would sort on driving ability. But given the economy-wide division of labor into increasingly specialized roles, such limits to generalizability are nearly ubiquitous in applied research on worker incentives. Whether they come from rideshare drivers (Angrist et al., 2021; Cook et al., 2021), agricultural workers (Brune et al., 2022; Bandiera et al., 2010), cashiers (Mas and Moretti, 2009), or automotive glass repairers (Lazear, 2000), parameter estimates concerning worker productivity are usually difficult generalize beyond narrowly defined labor markets.[18] While my study is not exempt from these limitations, several elements of my experiment are designed to mitigate these concerns. First, workers in my experiment are recruited using a widely used and well-established freelancing platform with over 100,000 workers (Difallah et al., 2018). The ubiquity of such platforms (e.g., MTurk, Upwork, Fiverr) means that even the most conservative interpretation of my estimates holds non-trivial welfare implications. Second, workers are not aware that they are part of an experiment until after they perform the task, so estimates are biased by their potential desire to generate a particular result.[19] Third, my experimental typing task requires a dimension of effort and

[18]Indeed, Herbst and Mas (2015) finds that for one particular parameter—peer effects on worker output—estimates vary dramatically from one study to another, regardless of whether estimates are taken from the lab or the field.

[19]Rather than elicit consent prior to the task, I debriefed participants on the nature of the experiment after its conclusion. The experimental task resembles those commonly requested for non-research purposes.

skill commonly needed for gig work. "Traditional keyboarding" is a job requirement for 66 percent of American workers (Bureau of Labor Statistics, 2024), suggesting my estimates of selection and incentive effects could plausibly generalize to a variety of labor markets.

A second threat to external validity concerns the selection of workers into the experiment itself. Even if the eligible population for my study were representative of the broader population of gig workers, I still only observe outcomes for those who initially agreed to the task. If my job posting only attracted low-productivity workers, my estimates would exclude selection among high-types because they never received an offer. I mitigate this concern by advertising a generous up-front fee for accepting the task. By posting a guaranteed \$1.00 plus the \$0.03-per-entry piece rate offered to all treatment groups, I am likely to attract a broad swath of workers who meet my screening criteria.

A related concern is the \$0.03-per-entry piece rate. While I set this rate to roughly correspond to rates commonly seen both in my own observations and in previous literature (DellaVigna and Pope, 2022), it may not perfectly reflect the true market value of workers' labor product. If, for example, \$0.03 was an overestimate of the true market value of a typed data entry, estimates of reservation wages and marginal values would be inflated above what one would expect if workers' outside options more accurately represented the (uninsured) value of their labor. In Appendix A, I show that under fairly benign assumptions concerning worker utility, these curves are proportional to the per-unit price at which workers can sell their labor product. In particular, if workers' contract preferences exhibit constant relative risk aversion, I can normalized reservation wages and marginal values to a piece rate of \$1, allowing me to express welfare loss on a per-dollar basis.

Finally, external validity might suffer if the identifying variation in contracts is too localized. To credibly estimate welfare effects of information asymmetries, the observed variation in wage offers must cover the entire range of contracts that might exist under a benchmark counterfactual where firms are fully informed. This limitation is especially relevant in research involving information asymmetries, because the presence of those asymmetries can often limit the set of contracts observed in existing markets. As a result, methods us-

ing real-world wage contracts are likely to understate the consequences of adverse selection (Einav et al., 2010b). My design holds a distinct advantage over these "under-the-lamppost" methods, as my experiment allows me to create a market for hourly wage contracts that real-world employers might deem unprofitable due to market unraveling.

5 Conclusion

Short-term labor contracts offer workers the benefits of flexible schedules (Mas and Pallais, 2017) and liquid income (Garin et al., 2020; Koustas, 2018). However, these benefits come at a cost—a gig worker can choose their own hours, but they often face more uncertainty over what they'll earn during those hours. This trade off is likely not a coincidence—traditional employers can profitably sustain hourly wage contracts because their repeated interactions with workers mitigate information asymmetries. Short-term employment relationships, on the other hand, offer less opportunity to reveal workers' latent productivity, making the implicit insurance of hourly wages more prone to moral hazard and adverse selection on worker productivity.

This paper uses an experimental approach to investigate information asymmetries in short-term labor markets. The experiment offers participants a choice between a performance-based piece rate and a randomized hourly wage, allowing me to separately identify selection and treatment effects of wage contracts.

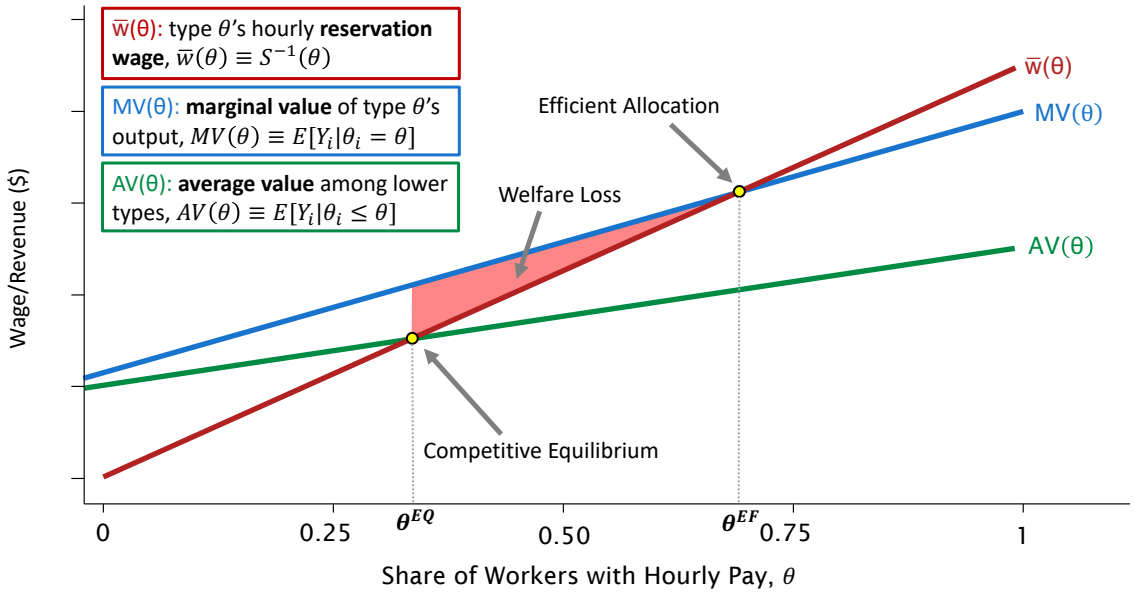
Findings from a preliminary pilot experiment reveal significant selection effects—a one-dollar increase in the offered hourly wage correlates with a \$0.32 (SE=0.156) increase in hourly output among those who prefer the piece rate. This result provides evidence of adverse selection in contract preference, wherein workers with higher latent productivity are more likely to forego guaranteed wages in favor of performanced-based compensation.

I place these experimental estimates into a theoretical framework that shows how the provision of hourly employment contracts is determined by two factors: a worker's reservation wage—the minimum compensation they will accept in exchange for an hour of labor—

and the average output of workers with comparatively lower reservation wages. These objects can be straightforwardly identified for workers in my experimental sample—first by comparing the shares of workers opting into hourly wages across offer treatments, then by comparing the average output among hourly workers in each group. I then show how to use these model estimates to quantify the welfare loss associated with inefficiently low provision of hourly positions.

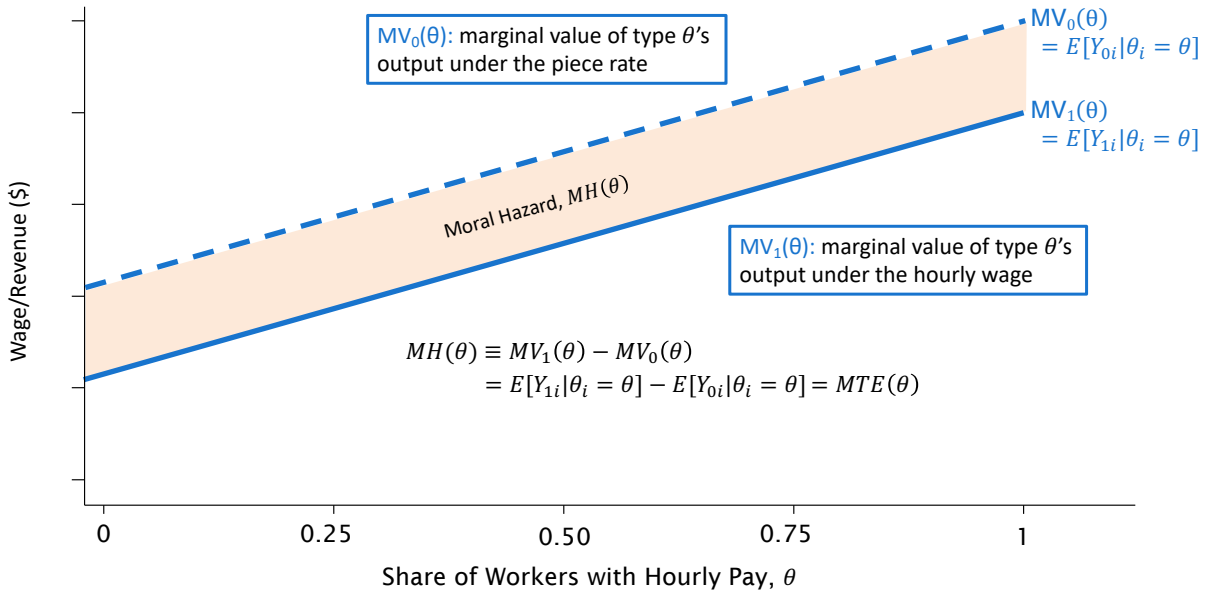
Figures and Tables

Figure 1: Model of Asymmetric Information in Wage Contracts



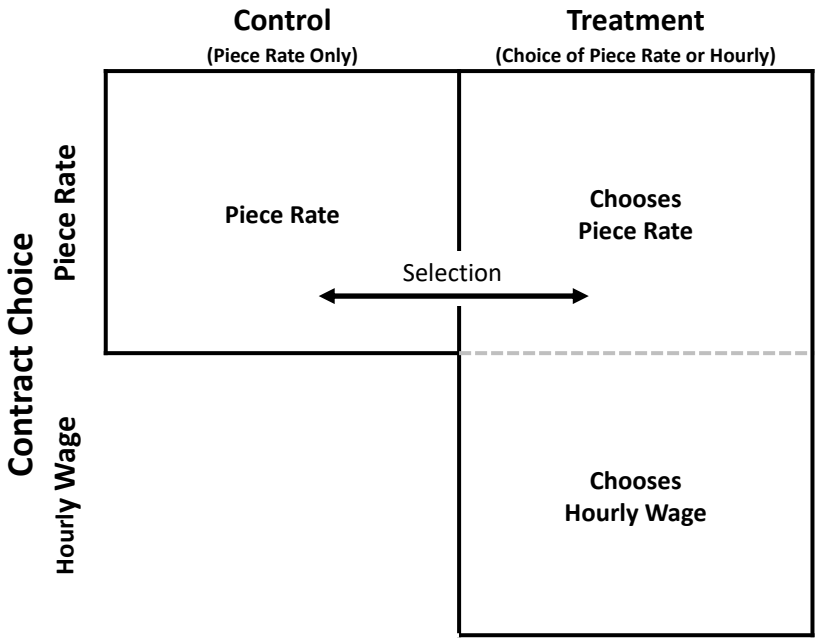
Note: This figure provides a graphical representation of market unraveling for hourly wages. On the horizontal axis, types θ are enumerated in ascending order based on their hourly reservation wage, $\overline{w}_i$. The blue line plots the marginal value curve, $MV(\theta)$, which is equal to expected worker output value conditional on their type, $MV(\theta) \equiv E[Y_i|\theta_i = \theta]$. The red line plots hourly reservation wage, $\overline{w}(\theta)$, which equals the inverse of labor supply, $\overline{w}(\theta) \equiv S^{-1}(\theta)$. The green line plots the average value curve, $AV(\theta)$, which corresponds to the average expected output among lower-type workers, $AV(\theta) \equiv E[Y_i|\theta_i \leq \theta]$.

Figure 2: Model of Asymmetric Information in Wage Contracts: Moral Hazard Effects



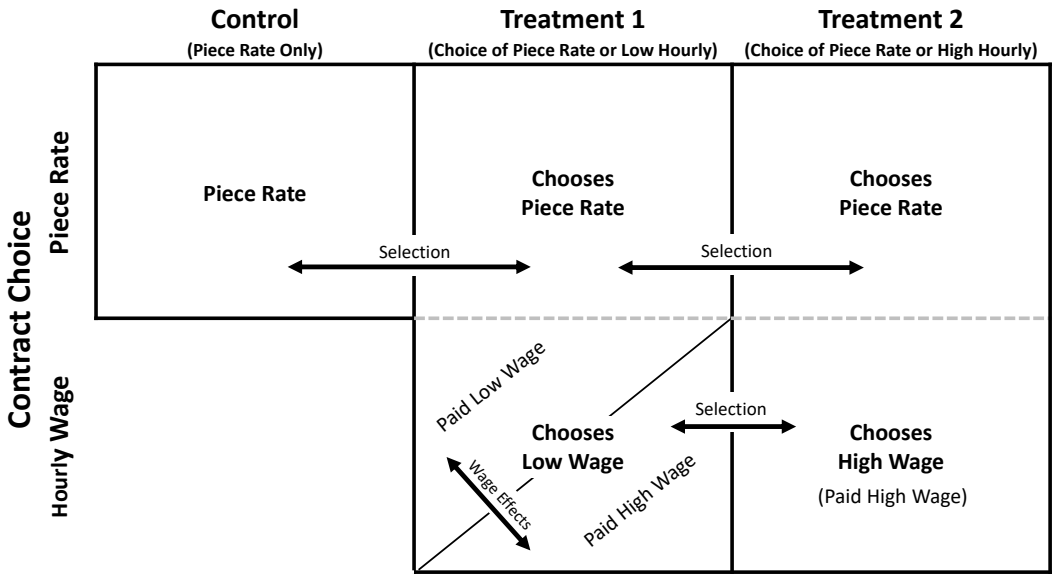
Note: This figure provides a graphical representation of moral hazard in my model. On the horizontal axis, types θ are enumerated in ascending order based on their hourly reservation wage, $\overline{w}_i$. The solid blue line plots $MV_1(\theta)$, which is equal to the expected output value among workers of type θ under the hourly wage, $MV_1(\theta) \equiv E[Y_{1i}|\theta_i = \theta]$. The shaded blue line plots $MV_0(\theta)$, which is equal to the expected output value among the same workers if they were instead paid a piece rate, $MV_0(\theta)\, E[Y_{0i}|\theta_i = \theta]$. The difference between the two marginal value curves identifies the moral hazard effect for a given type, $MH(\theta) \equiv MV_1(\theta) - MV_0(\theta)$, which is equivalent to the marginal treatment effect of the hourly contract among those whose resistance to treatment (quantile reservation wage, $\theta_i \equiv S(\overline{w}_i)$) is equal to the propensity score (share of hourly workers, $\theta = S(w)$) for their assigned instrument (wage offer, w_i).

Figure 3: Experimental Design: Single Treatment



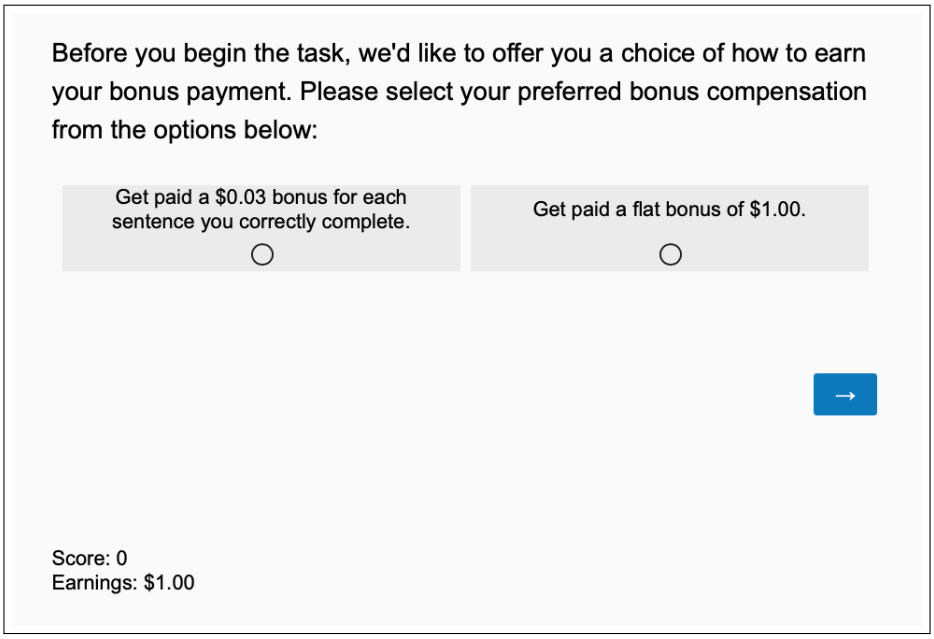
Note: This figure provides a graphical representation of a single-offer version of my experimental design. Columns denote experimental groups with different menus of wage options, and rows denote the realized wage contracts chosen by workers within each group. The control group, represented by the left column, is not offered an hourly wage option and is compensated through the piece-rate contract (upper box). The treatment group, represented by the right column, is separated into those who accept the piece-rate contract (upper box) and those who accept the hourly contract (lower box). The solid arrow denotes comparison groups to measure adverse selection—groups that were offered different menus of contracts but ultimately face the same repayment terms.

Figure 4: Experimental Design: Multiple Treatments

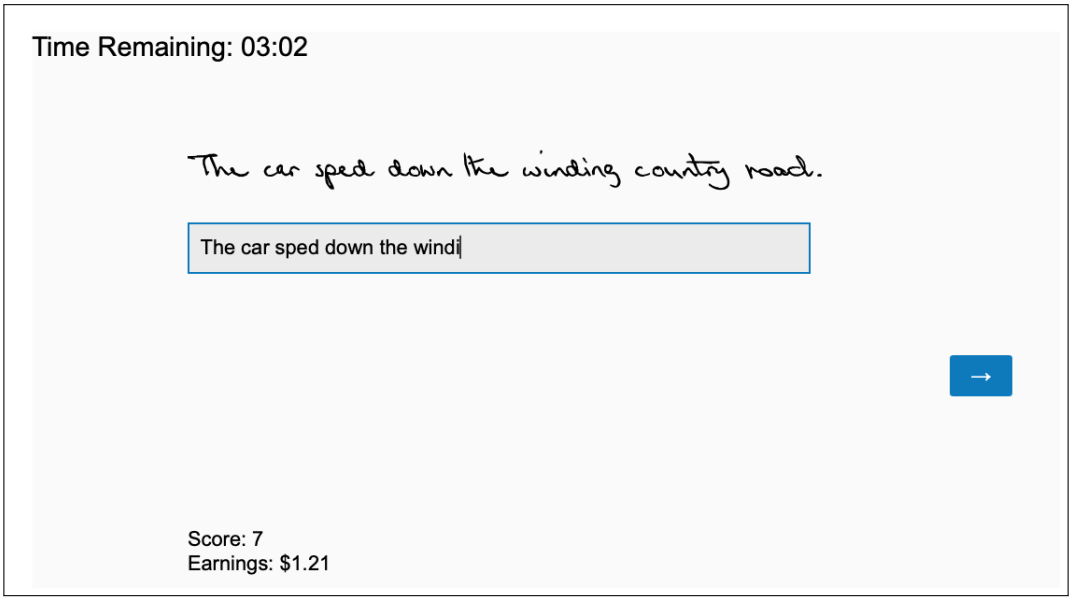


Note: This figure provides a graphical illustration of my two-stage experimental design with two treatment offers. Columns denote initial hourly wage offers, and rows denote the type of contract that workers choose. The diagonal split in the bottom box of Treatment 1 (low-wage offer) represents the second stage of randomization, in which some workers accepting the low hourly wage are promised the higher wage before they begin the task. Horizontal arrows denotes comparison groups to measure adverse selection—groups that were offered different menus of contracts but ultimately face the same repayment terms. The diagonal arrow denotes comparison groups to measure wage effects. The treatment effect of a hourly wages relative to piece rates (moral hazard) is identified by instrumenting for a given (paid) hourly wage with initial wage offers.

Figure 5: Example Job Posting



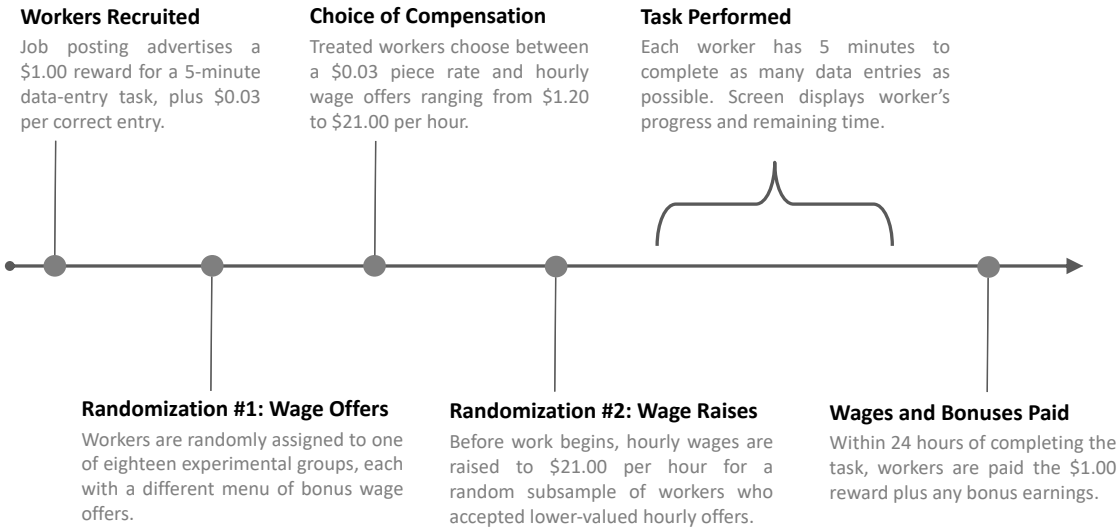
(A) Example Wage Offer



(B) Typing Task

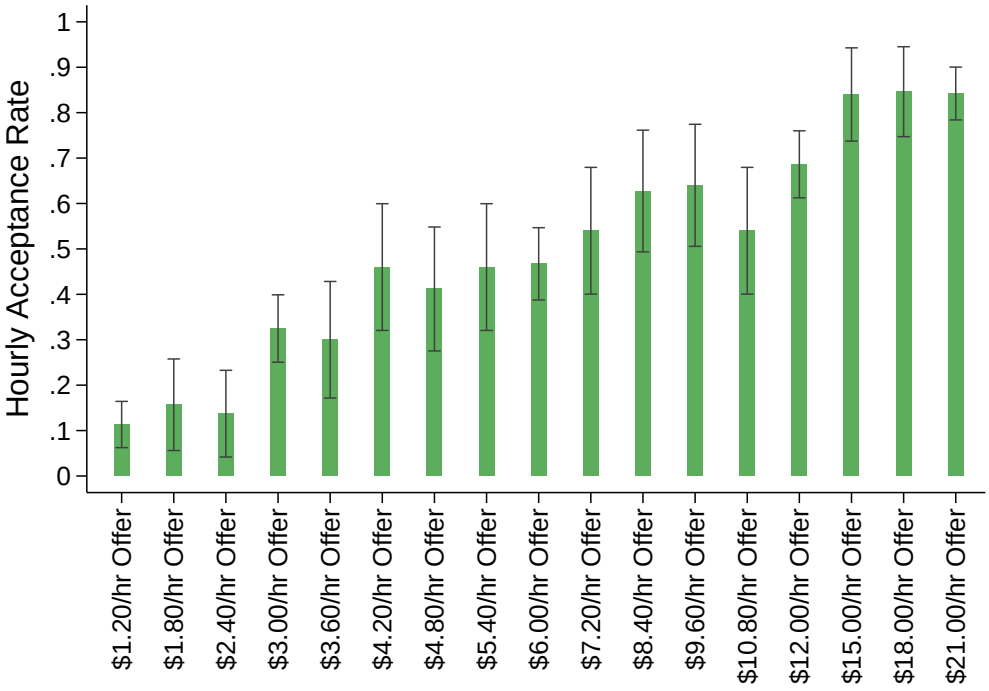
Note: This figure provides screenshots of the experimental intervention. Panel A shows an example wage offer participants see before they begin the task. Panel B shows the sentence-typing task while it is being performed.

Figure 6: Experiment Timeline



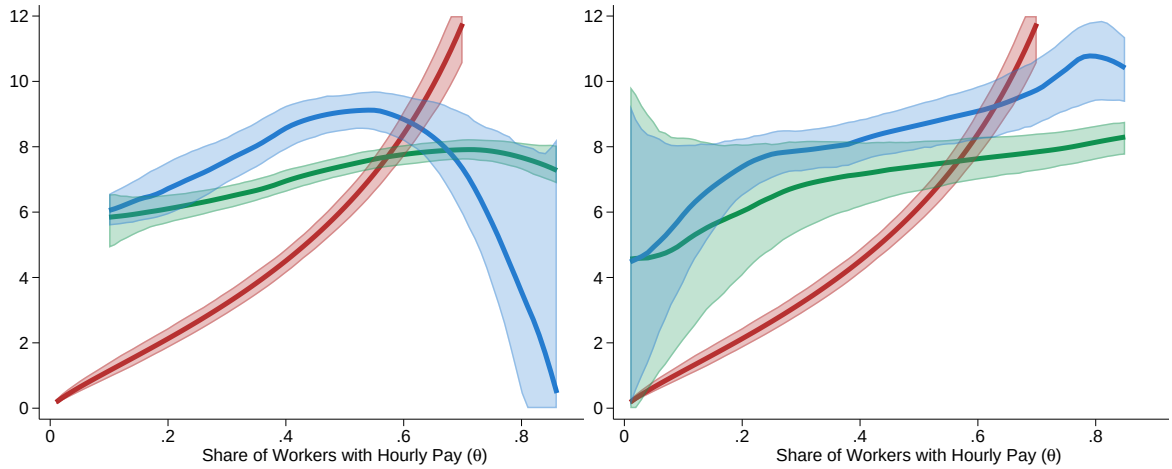
Note: This figure provides a timeline for a single wave of the experiment.

Figure 7: Hourly Wage Take-Up



Note: This figure reports hourly-wage acceptance rates by treatment group. The y-axis measures the share of borrowers in each group who declined the $0.03 piece rate in favor of the hourly wage offer displayed on the x-axis. Note that both piece rate and hourly wage options are offered as supplements to a base wage of $12.00 per hour. Dotted bands indicate 90% confidence intervals.

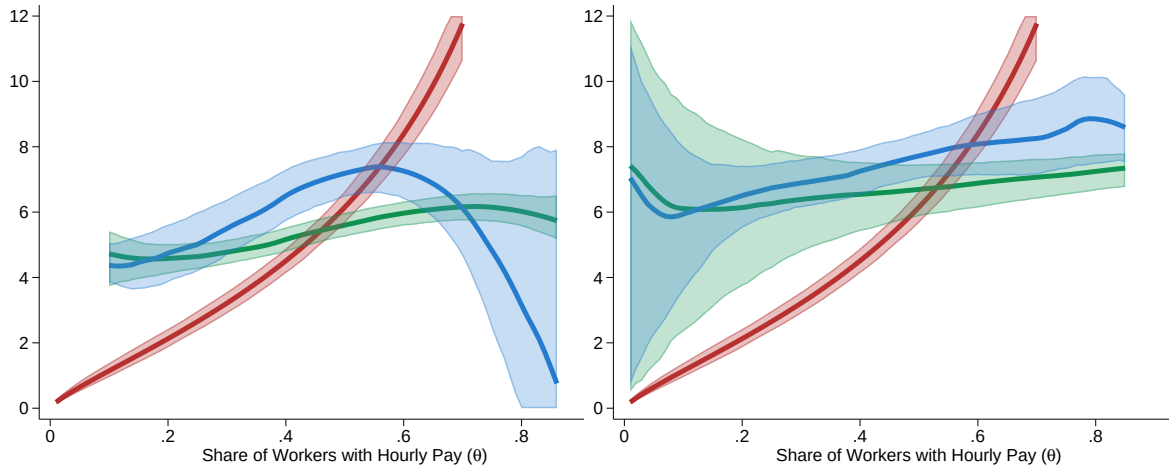
Figure 8: Estimates of Marginal and Average Value Curves



(A) Potential Value Under Hourly Wage (B) Potential Value Under Piece Rate

Note: This figure plots estimates supply and value curves, where output values reflect the number of typed sentences multiplied by the piece rate. In the left panel, the blue and green lines plot semiparametric estimates of the marginal value, $MV_1(\theta)$, and average value $AV_1(\theta)$, under hourly wages, as defined in Figure 1. In the left panel, blue and green lines plot these same curves ($MV_0(\theta)$ and $AV_0(\theta)$) under a piece-rate counterfactual. In both panels, the red line plots estimated hourly supply curve from a logit regression of hourly take-up against experimental wage offers. Value curves are estimated using a second-degree local polynomial regression of residualized hourly output value against predicted hourly supply. Shaded regions represent 95% confidence intervals.

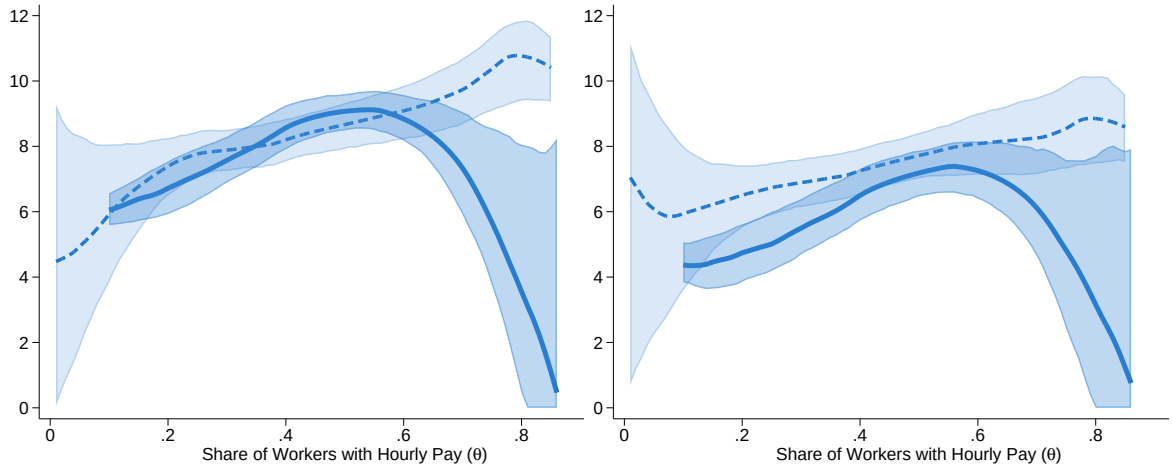
Figure 9: Estimates of Marginal and Average Value Curves: Quality Adjusted



(A) Potential Value Under Hourly Wage (B) Potential Value Under Piece Rate

Note: This figure plots estimates supply and value curves, where output values reflect the number of *correctly* typed sentences multiplied by the piece rate. In the left panel, the blue and green lines plot semiparametric estimates of the marginal value, $MV_1(\theta)$, and average value $AV_1(\theta)$, under hourly wages, as defined in Figure 1. In the left panel, blue and green lines plot these same curves ($MV_0(\theta)$ and $AV_0(\theta)$) under a piece-rate counterfactual. In both panels, the red line plots estimated hourly supply curve from a logit regression of hourly take-up against experimental wage offers. Value curves are estimated using a second-degree local polynomial regression of residualized hourly output value against predicted hourly supply. Shaded regions represent 95% confidence intervals.

Figure 10: Estimates of Moral Hazard Effects



(A) Marginal Values: Raw Output (B) Marginal Values: Quality-Adjusted Output

Note: This figure plots marginal value curves under hourly and piece-rate counterfactuals, as defined in Figure 2. Solid lines denote $MV_1(\theta)$–the marginal worker's potential output value under an hourly wage. Dotted lines lines denote $MV_0(\theta)$–the marginal worker's potential output value under the piece. The right panel measures output value using the total number of sentences multiplied by the piece rate. The second panel uses *correctly* typed sentences multiplied by the piece rate. Value curves are estimated using a second-degree local polynomial regression of residualized hourly output value against predicted hourly supply. Shaded regions represent 95% confidence intervals.

Table 1: Experimental Group Assignment

Hourly Wage Offer	Piece-Rate Offer	Number of Participants
No Hourly Offer	\$0.03 per sentence	152
\$1.20/hr Offer	\$0.03 per sentence	150
\$1.80/hr Offer	\$0.03 per sentence	51
\$2.40/hr Offer	\$0.03 per sentence	51
\$3.00/hr Offer	\$0.03 per sentence	154
\$3.60/hr Offer	\$0.03 per sentence	50
\$4.20/hr Offer	\$0.03 per sentence	50
\$4.80/hr Offer	\$0.03 per sentence	51
\$5.40/hr Offer	\$0.03 per sentence	50
\$6.00/hr Offer	\$0.03 per sentence	152
\$7.20/hr Offer	\$0.03 per sentence	50
\$8.40/hr Offer	\$0.03 per sentence	51
\$9.60/hr Offer	\$0.03 per sentence	50
\$10.80/hr Offer	\$0.03 per sentence	50
\$12.00/hr Offer	\$0.03 per sentence	153
\$15.00/hr Offer	\$0.03 per sentence	50
\$18.00/hr Offer	\$0.03 per sentence	52
\$21.00/hr Offer	\$0.03 per sentence	152

Note: This table summarizes the treatment conditions and sample sizes for each experimental group in the pilot. *Piece-rate offer* denotes the performance-based bonus offer, which is awarded on a per-sentence basis and common across all experimental groups. *Hourly wage offer* denotes the time-based rate of compensation offered to participants for the 5-minute task, prorated to one hour. Note that both piece rate and hourly wage options are offered as supplements to a base wage of \$12.00 per hour.

Table 2: OLS Estimates of Selection by Wage Offer

	(1) Log Output Value	(2) Log Output Value	(3) Log Output Value
Log Hourly Wage Offer	0.0290*	0.0302*	0.0337**
	(0.0163)	(0.0161)	(0.0147)
Accepted	-0.379***	-0.377***	-0.359***
	(0.0801)	(0.0795)	(0.0758)
Accepted × Log Hourly Wage Offer	0.0553*	0.0524	0.0397
	(0.0329)	(0.0331)	(0.0300)
Accepted × Log Effective Hourly Wage	-0.000981	0.000153	0.00888
	(0.0305)	(0.0308)	(0.0288)
Experience Controls	No	Yes	Yes
Demographic Controls	No	No	Yes
R-squared	0.097	0.100	0.223
N	1358	1358	1358

Note: This table reports estimated coefficients from OLS regressions of output value (sentences ×\$0.03) against hourly wage offers interacted with a dummy for whether an individual accepted the hourly offer ("Accept") over the piece rate. The coefficient on "Log Hourly Wage Offer" captures the change in log output value among piece-rate workers for each unit increase in their log hourly wage offer. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 3: OLS Estimates of Selection by Wage Offer

	(1) Log Output Value	(2) Log Output Value	(3) Log Output Value
Log Hourly Wage Offer	0.0290*	0.0302*	0.0337**
	(0.0163)	(0.0161)	(0.0147)
Accepted	-0.379***	-0.377***	-0.359***
	(0.0801)	(0.0795)	(0.0758)
Accepted × Log Hourly Wage Offer	0.0553*	0.0524	0.0397
	(0.0329)	(0.0331)	(0.0300)
Accepted × Log Effective Hourly Wage	-0.000981	0.000153	0.00888
	(0.0305)	(0.0308)	(0.0288)
Experience Controls	No	Yes	Yes
Demographic Controls	No	No	Yes
R-squared	0.097	0.100	0.223
N	1358	1358	1358

Note: This table reports estimated coefficients from OLS regressions of output value (sentences ×\$0.03) against hourly wage offers interacted with a dummy for whether an individual accepted the hourly offer ("Accept") over the piece rate. The coefficient on "Log Hourly Wage Offer" captures the change in log output value among piece-rate workers for each unit increase in their log hourly wage offer. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 4: 2SLS Estimates of Treatment Effects of Hourly Wages Relative to Piece Rates

	(1) Log Output Value	(2) Log Output Value	(3) Log Output Value
Accepted Hourly Offer	-0.0823**	-0.0830**	-0.0683**
	(0.0382)	(0.0380)	(0.0347)
Experience Controls	No	Yes	Yes
Demographic Controls	No	No	Yes
R-squared	0.043	0.048	0.175
N	1509	1509	1509

Note: This table reports estimated coefficients from two-stage least-squares regressions of log output value—the number of typed sentences multiplied by the piece rate—against hourly wages, where I instrument for hourly wage with treatment-group assignment. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

Table 5: 2SLS Estimates of Treatment Effects of Hourly Wages Relative to Piece Rates on Quality-Adjusted Output Value

	(1) Log Output Value (Quality Adj.)	(2) Log Output Value (Quality Adj.)	(3) Log Output Value (Quality Adj.)
Accepted Hourly Offer	-0.234***	-0.236***	-0.220***
	(0.0617)	(0.0614)	(0.0582)
Experience Controls	No	Yes	Yes
Demographic Controls	No	No	Yes
R-squared	0.060	0.064	0.127
N	1473	1473	1473

Note: This table reports estimated coefficients from two-stage least-squares regressions of log output value against hourly wages, where output values reflect the number of *correctly* typed sentences multiplied by the piece rate. I instrument for hourly wage with treatment-group assignment. *** = significant at 1 percent level, ** = significant at 5 percent level, * = significant at 10 percent level.

References

Abraham, Katharine, John Haltiwanger, Kristin Sandusky, and James Spletzer (2017) "Measuring the gig economy: Current knowledge and open issues," *Measuring and Accounting for Innovation in the 21st Century*.

Akerlof, George A (1970) "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 84 (3), 488–500.

——— (1982) "Labor contracts as partial gift exchange," *The quarterly journal of economics*, 97 (4), 543–569.

Angrist, Joshua D, Sydnee Caldwell, and Jonathan V Hall (2021) "Uber Versus Taxi: A Driver's Eye View," *American Economic Journal: Applied Economics*, 13 (3), 272–308.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul (2010) "Social incentives in the workplace," *The review of economic studies*, 77 (2), 417–458.

Björklund, Anders and Robert Moffitt (1987) "The estimation of wage gains and welfare gains in self-selection models," *The Review of Economics and Statistics*, 42–49.

Black, Dan A, Joonhwi Joo, Robert LaLonde, Jeffrey A Smith, and Evan J Taylor (2022) "Simple tests for selection: Learning more from instrumental variables," *Labour Economics*, 79, 102237.

Brune, Lasse, Eric Chyn, and Jason Kerwin (2022) "Peers and motivation at work: evidence from a firm experiment in Malawi," *Journal of Human Resources*, 57 (4), 1147–1177.

Bureau of Labor Statistics (2024) "Occupational Requirements in the United States," Website, `https://www.bls.gov/web/ors.supp.toc.htm` Accessed: 2024-03-28.

Carneiro, Pedro, James J Heckman, and Edward J Vytlacil (2011) "Estimating marginal returns to education," *American Economic Review*, 101 (6), 2754–2781.

Chiappori, Pierre-André and Bernard Salanie (2000) "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108 (1), 56–78.

Collins, Brett, Andrew Garin, Emilie Jackson, Dmitri Koustas, and Mark Payne (2019) "Is gig work replacing traditional employment? Evidence from two decades of tax returns," *Unpublished paper, IRS SOI Joint Statistical Research Program*.

Cook, Cody, Rebecca Diamond, Jonathan V Hall, John A List, and Paul Oyer (2021) "The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers," *The Review of Economic Studies*, 88 (5), 2210–2238.

DellaVigna, Stefano and Devin Pope (2018) "What motivates effort? Evidence and expert forecasts," *The Review of Economic Studies*, 85 (2), 1029–1069.

——— (2022) "Stability of Experimental results: Forecasts and evidence," *American Economic Journal: Microeconomics*, 14 (3), 889–925.

Difallah, Djellel, Elena Filatova, and Panos Ipeirotis (2018) "Demographics and dynamics of mechanical turk workers," in *Proceedings of the eleventh ACM international conference on web search and data mining*, 135–143.

Douglas, Benjamin D, Patrick J Ewell, and Markus Brauer (2023) "Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA," *Plos one*, 18 (3), e0279720.

Einav, Liran, Amy Finkelstein, and Mark R Cullen (2010a) "Estimating Welfare in Insurance Markets Using Variation in Prices," *The Quarterly Journal of Economics*, 125 (3), 877–921.

Einav, Liran, Amy Finkelstein, and Jonathan Levin (2010b) "Beyond testing: Empirical models of insurance markets," *Annu. Rev. Econ.*, 2 (1), 311–336.

Emanuel, Natalia and Emma Harrington (2024) "Working remotely? Selection, treatment, and the market for remote work," *American Economic Journal: Applied Economics (forthcoming)*.

Garin, Andrew, Emilie Jackson, Dmitri K Koustas, and Carl McPherson (2020) "Is new platform work different from other freelancing?" in *AEA Papers and Proceedings*, 110, 157–161, American Economic Association.

Garin, Andrew, Emilie Jackson, Dmitri K Koustas, and Alicia Miller (2023) "The Evolution of Platform Gig Work, 2012-2021,"Technical report, National Bureau of Economic Research.

Gibbons, Robert and Lawrence F Katz (1991) "Layoffs and lemons," *Journal of labor Economics*, 9 (4), 351–380.

Greenwald, Bruce C (1986) "Adverse selection in the labour market," *The Review of Economic Studies*, 53 (3), 325–347.

Heckman, James J and Edward Vytlacil (2005) "Structural equations, treatment effects, and econometric policy evaluation," *Econometrica*, 73 (3), 669–738.

Heckman, James J and Edward J Vytlacil (1999) "Local instrumental variables and latent variable models for identifying and bounding treatment effects," *Proceedings of the national Academy of Sciences*, 96 (8), 4730–4734.

——— (2007) "Econometric evaluation of social programs, part II: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments," *Handbook of econometrics*, 6, 4875–5143.

Hendren, Nathaniel (2017) "Knowledge of Future Job Loss and Implications for Unemployment Insurance," *American Economic Review*, 107 (7), 1778–1823.

Herbst, Daniel and Nathaniel Hendren (2021) "Opportunity Unraveled: Private Information and the Missing Markets for Financing Human Capital,"Technical report, National Bureau of Economic Research.

Herbst, Daniel and Alexandre Mas (2015) "Peer effects on worker output in the laboratory generalize to the field," *Science*, 350 (6260), 545–549.

Huber, Martin (2013) "A simple test for the ignorability of non-compliance in experiments," *Economics letters*, 120 (3), 389–391.

Jackson, Emilie, Adam Looney, and Shanthi P Ramnath (2017) "The rise of alternative work arrangements: Evidence and implications for tax filing and benefit coverage."

Jovanovic, Boyan (1982) "Favorable selection with asymmetric information," *The Quarterly Journal of Economics*, 97 (3), 535–539.

Kantarevic, Jasmin and Boris Kralj (2016) "Physician payment contracts in the presence of moral hazard and adverse selection: the theory and its application in Ontario," *Health economics*, 25 (10), 1326–1340.

Karlan, Dean and Jonathan Zinman (2009) "Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment," *Econometrica*, 77 (6), 1993–2008.

Katz, Lawrence F and Alan B Krueger (2019) "Understanding trends in alternative work arrangements in the United States," *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 5 (5), 132–146.

Koustas, Dmitri (2018) "Consumption insurance and multiple jobs: Evidence from rideshare drivers," *Unpublished working paper*.

Kowalski, Amanda E (2023a) "How to examine external validity within an experiment," *Journal of Economics & Management Strategy*, 32 (3), 491–509.

——— (2023b) "Reconciling seemingly contradictory results from the Oregon health insurance experiment and the Massachusetts health reform," *Review of Economics and Statistics*, 105 (3), 646–664.

Krueger, Alan B and Lawrence H Summers (1988) "Efficiency wages and the inter-industry wage structure," *Econometrica: Journal of the Econometric Society*, 259–293.

Kugler, Adriana D and Gilles Saint-Paul (2004) "How do firing costs affect worker flows in a world with adverse selection?" *Journal of Labor Economics*, 22 (3), 553–584.

Lazear, Edward P (1986) "Salaries and piece rates," *Journal of business*, 405–431.

——— (2000) "Performance pay and productivity," *American Economic Review*, 90 (5), 1346–1361.

Levine, David I (1991) "Just-cause employment policies in the presence of worker adverse selection," *Journal of Labor Economics*, 9 (3), 294–305.

Malcomson, James M (1981) "Unemployment and the efficiency wage hypothesis," *The Economic Journal*, 91 (364), 848–866.

Mas, Alexandre and Enrico Moretti (2009) "Peers at work," *American Economic Review*, 99 (1), 112–145.

Mas, Alexandre and Amanda Pallais (2017) "Valuing alternative work arrangements," *American Economic Review*, 107 (12), 3722–3759.

Moen, Espen R and Åsa Rosen (2005) "Performance Pay and Adverse Selection," *Scandinavian Journal of Economics*, 107 (2), 279–298.

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

Pallais, Amanda and Emily Glassberg Sands (2016) "Why the referential treatment? Evidence from field experiments on referrals," *Journal of Political Economy*, 124 (6), 1793–1828.

Robinson, Peter M (1988) "Root-N-consistent semiparametric regression," *Econometrica: Journal of the Econometric Society*, 931–954.

Salop, Steven C (1979) "A model of the natural rate of unemployment," *The American Economic Review*, 69 (1), 117–125.

Shapiro, Carl and Joseph E Stiglitz (1984) "Equilibrium unemployment as a worker discipline device," *The American economic review*, 74 (3), 433–444.

Shearer, Bruce (1996) "Piece-rates, Principal-agent Models, and Productivity Profiles: Parametric and Semi-parametric Evidence from Payroll Records," *Journal of Human Resources*, 275–303.

——— (2004) "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment," *The Review of Economic Studies*, 71 (2), 513–534.

Spence, Michael (1973) "Job Market Signaling," *The Quarterly Journal of Economics*, 87 (3), 355–374.

Wald, Abraham (1940) "The fitting of straight lines if both variables are subject to error," *The annals of mathematical statistics*, 11 (3), 284–300.

Weiss, Andrew (1980) "Job Queues and Layoffs in Labor Markets with Flexible Wages," *Journal of Political economy*, 88 (3), 526–538.

——— (2014) *Efficiency wages: Models of unemployment, layoffs, and wage dispersion*, 1192: Princeton University Press.

Yellen, Janet L (1984) "Efficiency Wage Models of Unemployment," *The American Economic Review*, 74 (2), 200–205.

Appendix A Welfare Under Alternative Piece Rates

Let $\overline{w}(\theta;p)$ denote type θ's hourly reservation wage from Equation (2) when their outside option is selling their labor product, q, at a per-unit price, p. Given some distribution of potential output, $F_\theta(q)$, $\overline{w}(\theta;p)$ equals the certainty equivalent of type θ's earnings under the piece rate p, $\overline{w}(\theta;p) = u^{-1}\left(E\left[u(pq)|\theta\right]\right)$. Assuming preferences exhibit constant relative risk aversion,

$$\begin{aligned}
\overline{w}(\theta;p) &= u^{-1}\left(E\left[u(pq)|\theta\right]\right) && (25)\\
&= \left((1-\rho)E\left[\frac{(pq)^{1-\rho}}{1-\rho}|\theta\right]\right)^{\frac{1}{1-\rho}} && (26)\\
&= pu^{-1}\left(E\left[u(q)|\theta\right]\right) && (27)\\
&= p\overline{w}(\theta;1), && (28)
\end{aligned}$$

where ρ is the coefficient of relative risk aversion.

Now let $MV(\theta;p)$ denote the marginal value of type θ's labor product from Equation (3) when its sold at a per-unit price of p:

$$\begin{aligned}
MV\left(\theta;p\right) &\equiv E\left[pq_i|\theta_i=\theta\right] && (29)\\
&= pE\left[q_i|\theta_i=\theta\right] = pMV(\theta;1). && (30)
\end{aligned}$$

Equations (28) and (30) allow me to rewrite welfare loss from Equation (7) for a given piece-rate, p, as

$$\begin{aligned}
DWL(p) &= \int_{\theta_{EQ}}^{\theta_{EF}}\left(MV(\theta;p)-\overline{w}(\theta;p)\right)d\theta && (31)\\
&= \int_{\theta_{EQ}}^{\theta_{EF}}\left(pMV(\theta;1)-p\overline{w}(\theta;1)\right)d\theta && (32)\\
&= pDWL(1). && (33)
\end{aligned}$$

Equation (33) shows how welfare loss from the under provision of hourly wage contracts is

proportional to the per-unit value of workers' labor product. Under CRRA utility, I can therefore divide DWL by p to express welfare loss *per dollar earned* under the piece rate.

Note that these counterfactual welfare calculations assume worker production does not respond to different piece rates. This assumption might be violated if a higher piece rate (p) induces greater effort, resulting in higher output (q). To the extent the returns from this higher output exceeds the worker's disutility of effort, this incentive effect would attenuate counterfactual welfare estimates towards those calculated under the experimental piece rate.

Appendix B Identification of Wage Effects

Consider an individual i who receives a job offer, J_i, at one of two randomized wages: a high offer ($J_i = H$) or a low offer ($J_i = L$). Let D_{Ji} denote the individual's choice to accept a given offer, so that $D_{Hi} = 1$ if i would accept the high offer and $D_{Li} = 1$ if i would accept the low offer. Furthermore, let Y_{Hi} and Y_{Li} denote the potential output levels produced by i if they were paid hourly wages of H and L, respectively. Note that if realized wages reflected accepted offers, comparing output between those who accept H and those who accept L would yield the following:

$$\begin{aligned} &E\left[Y_i|J_i = H, D_{Hi} = 1\right] - E\left[Y_i|J_i = L, D_{Li} = 1\right] \\ &= \underbrace{E\left[Y_{Hi} - Y_{Li}|D_{Li} = 1\right]}_{\text{Wage Effect}} + \underbrace{E\left[Y_{Hi}|D_{Hi} = 1\right] - E\left[Y_{Hi}|D_{Li} = 1\right]}_{\text{Selection}}. \end{aligned} \tag{34}$$

This difference is the sum of both the wage effect and selection of H relative to L, which cannot be separated without observing $E\left[Y_{Hi}|D_{Li} = 1\right]$.

Now let W_i be an indicator whether individual i receives a surprise wage increase of $\Delta = H - L$ after accepting their contract. W_i is randomly assigned among those who received low offers ($J_i = L$) and accepted them ($D_{Li} = 1$) but is zero for everyone else. With this randomized wage raise, I can estimate wage effects by comparing output between low- and high-wage workers in the low-offer group:

$$\begin{aligned} \text{Wage Effect} &= E\left[Y_i|J_i = L, D_{Li} = 1, W_i = 1\right] - E\left[Y_i|J_i = L, D_{Li} = 1, W_i = 0\right] \\ &= E\left[Y_{Hi} - Y_{Li}|D_{Li} = 1\right] \end{aligned} \tag{35}$$

And I can estimate selection by comparing output between low- and high-offer groups with

high realized wages:

$$\begin{aligned}\text{Selection} &= E\left[Y_i|J_i = H, D_{Hi} = 1\right] - E\left[Y_i|J_i = L, D_{Li} = 1, W_i = 1\right] \\ &= E\left[Y_{Hi}|D_{Hi} = 1\right] - E\left[Y_{Hi}|D_{Li} = 1\right]. \end{aligned} \tag{36}$$