

Inefficient Hiring in Entry-Level Labor Markets

Amanda Pallais*
Harvard University

July 1, 2011

Abstract

When a firm hires a novice worker, it obtains both labor services and information about the worker's productivity. The firm must pay a sunk cost to hire the novice, but the information generated has option value: novices found to have high productivity can be rehired in subsequent periods. If competing firms also observe the novice's productivity, however, the option value of hiring accrues to the worker, not the employer. Firms will accordingly under-invest in discovering novice talent unless they can claim the benefit, for example, by workers buying their entry-level jobs or signing contracts transferring the option value to the firm.

I formalize this intuition in a model of the labor market in which positive hiring costs and publicly observable output lead to inefficiently low novice hiring. I test the model's relevance in an online labor market by hiring 952 workers at random from an applicant pool of 3,767 for a 10-hour data entry job. In this market, worker performance is publicly observable. Consistent with the model's prediction, novice workers hired at random obtain significantly more employment and have higher earnings than the control group, following the initial hiring spell. A second treatment confirms that this causal effect is likely explained by information revelation rather than skills acquisition. Providing the market with more detailed information about the performance of a subset of the randomly-hired workers raised earnings of high-productivity workers and decreased earnings of low-productivity workers. Due to its scale, the experiment significantly increased the supply of workers recognized as high-ability in the market. This outward supply shift raised subsequent total employment and decreased average wages in occupations affected by the experiment (relative to non-treated occupations), implying that it also increased the sum of worker and employer surplus. Under plausible assumptions, this additional total surplus exceeds the social cost of the experiment.

*I would especially like to thank my advisors, David Autor, Esther Duflo, and Daron Acemoglu. I would also like to thank Jason Abaluck, Leila Agha, Josh Angrist, Bob Gibbons, Michael Greenstone, Jon Gruber, Daniel Keniston, Danielle Li, Christopher Palmer, Iuliana Pascu, Jim Poterba, David Powell, Michael Powell, Simone Schaner, Joe Shapiro, Pian Shu, Jean Tirole, Juuso Toikka, Sarah Turner, Heidi Williams and participants at MIT's labor and applied micro lunches for their many helpful comments and suggestions. I am grateful to Anand Hattiangi, Dmitry Diskin, and the oDesk Corporation for help understanding and accessing the oDesk data. Financial support from the George and Obie Shultz Fund, David Autor, and Esther Duflo is gratefully acknowledged.

1 Introduction

Young workers are more likely to be unemployed than workers who have had time to accumulate labor market experience. In August, 2010, for example, 14.7% of workers 20 to 24 years old were unemployed, compared with 9.5% of the general population.¹ Competition for entry-level jobs is intense. Newman (1996) found that the average success rate of applications to fast food restaurants, a common job for entry-level workers, in Harlem, NY was only 7%.

There are many potential explanations for the high unemployment rates of young workers, including that they are simply less skilled than more experienced workers. This paper proposes the hypothesis that firms hire inefficiently few inexperienced workers because they do not receive the full benefit of discovering workers' talent. It tests this hypothesis in a field experiment involving over 3,700 workers and finds strong support for it.

Employers are uncertain about the abilities of workers who lack experience. Hiring these workers generates information about their abilities (e.g., Farber and Gibbons, 1996; Altonji and Pierret, 2001), producing an option to hire the high-ability workers in the future. Generating this option is costly for firms. Managers must spend time explaining the jobs to workers and monitoring their progress. Moreover, firms incur an opportunity cost of lost time if jobs are not completed correctly or timely. If information about worker quality is partially public, high-ability workers receive part of the option value as higher wages. In the extreme case of pure public learning, workers receive the entire option value. If workers cannot compensate firms for generating this option, for example because they cannot accept sub-minimum wages or credibly commit to low-wage long-term contracts, firms will hire inefficiently few inexperienced workers.²

Tervio (2009) proposed that this type of informational inefficiency could generate excessively high wages for CEO's and entertainers. This paper shows that it may affect entry-level labor markets. There is substantial uncertainty about the abilities of entry-level workers, particularly those with little education and few credentials. Firms cannot conceal whether they have fired, retained, or promoted a worker, an important signal of entry-level worker performance. Entry-level labor markets have developed several institutions that reduce this inefficiency. Internships and hiring subsidies for young workers reduce firms' costs of hiring inexperienced workers. Fixed-term contracts (in Europe) reduce firms' costs of hiring inexperienced workers by allowing firms to dismiss low-ability young workers more easily. In some European training contracts, unions and industry consortiums pay workers' initial salaries. These allow the parties who benefit from talent discovery to pay for it.

Discovering a worker's ability is similar to general skills training: both produce a form of human capital that raises workers' value to firms, but require up-front investments. A large literature examines firms' provision of general skills training. Similar to this paper's model, Becker (1964)

¹This is from the Bureau of Labor Statistics's Table A-13, September 2010.

²These are not the only reasons why workers may not be able to compensate firms for creating the option. This inefficiency occurs at above-minimum wages if a wage rigidity, such as potential adverse selection of workers, prevents wages from falling.

shows that, because workers receive the benefits of general skills training, it will be underprovided if firms cannot be compensated for providing it. More recent theoretical work identifies circumstances in which firms receive part of the return on their training investments. If firms have monopsony power in the labor market (Acemoglu and Pischke, 1999), obtain private information about worker quality (Acemoglu and Pischke, 1998), or can use training to screen workers (Autor, 2001), they will provide some general skills training. There is some empirical evidence that firms provide general skills training that is not fully offset by lower wages (e.g., Autor, 2001; Loewenstein and Spletzer, 1998). However, neither the theoretical nor the empirical literature has shown that firms recoup the full value of their training investments and, thus, will provide the optimal level of training. This paper provides evidence that another form of human capital, information about worker ability, is underprovided by firms.

If not provided by firms, general human capital can be provided by schools; but output and information about workers' abilities may be jointly produced. Worker attributes such as reliability, enthusiasm, and maturity are difficult to verify outside of an employment context. Thus, if firms do not generate this information, there are few alternative mechanisms for its production.

This paper first develops a stylized model that demonstrates this inefficiency and generates testable predictions. In the model, firms have to pay a non-wage cost to hire novice workers. Hiring these workers generates information about their productivities, generating an option to hire the high-ability workers in the subsequent period. However, the market observes this information so the option value accrues to high-ability workers as higher wages. Because workers cannot provide their labor services below a certain wage (e.g., because of adverse selection or an inability to take sub-minimum wages), firms underinvest in discovering novice talent. The model predicts that providing inexperienced workers with an opportunity to demonstrate their ability through employment would allow some of them to be recognized as high-ability. This would increase inexperienced workers' future employment, earnings, and reservation wages. Moreover, by increasing the supply of workers recognized as high ability, it would increase future market employment, decrease market wages, and increase total market surplus: the sum of worker surplus and firm profits. The model also predicts that providing the market with more information about the job performance of inexperienced workers would increase the employment, earnings, and reservation wages of workers who performed well and impair the analogous employment outcomes of workers who performed poorly. Overall, it would increase the option value of hiring inexperienced workers, and thereby, increase their future earnings.

The paper then tests these predictions through a field experiment in a large, online marketplace: oDesk. oDesk workers are located all over the world and complete approximately 200,000 hours of work per week remotely. Many oDesk workers work full-time in the marketplace, though many others apply for jobs but never get work. Employers are required to provide public evaluations of all of their workers. They must rate workers on a one-to-five scale and may, at their option, leave a short comment. Workers cannot remove the numerical rating.³

³While they can remove the comment, only four percent of them do.

For the experiment, I posted 10-hour data entry jobs to the marketplace and invited low-wage data-entry specialists to apply. Out of an applicant pool of 3,767 workers, I hired 952 workers at random. Like all employers in the marketplace, after the job was complete, I rated the performance of each worker on a one-to-five scale. I gave most of these workers an uninformative comment, but I gave a randomly-selected subset a detailed comment with objective information about their data entry speed, accuracy, following of directions, and timely completion of the task. Using the marketplace's administrative data, I then observe all subsequent employment outcomes of the experimental workers on oDesk, including the jobs they obtain, the hours they work, and their total earnings. Three weeks later, I assess their reservation wages by inviting a randomly-selected subset of these workers to apply to a new job offered by a different employer with a randomly-selected wage.

The data support the model's predictions. In the two months after the experiment, workers randomly selected to receive treatment jobs were more likely to be employed, requested higher wages, and had higher earnings than the control group. Although the experiment did not target workers who were particularly likely to benefit from the treatment, it had very large effects. Excluding the experimental jobs, within two months of the treatment, the average treatment group worker had earned in excess of \$20 more than the average control group worker. I spent less than \$17 to hire each treatment group worker. Workers who received positive detailed comments earned substantially more after the experiment than workers who received uninformative comments. Workers who received negative detailed comments earned even less. Overall, workers who received detailed evaluations earned in excess of \$23 more than those who received less informative comments.

The large size of the experiment relative to the marketplace and the wide variation in its impact on the marketplace's 74 job categories allow me to evaluate the experiment's effects on the marketplace as a whole. Using a difference-in-difference strategy, I find that, after the experiment, employment increased and wages decreased in job categories more heavily affected by the experiment. I use these estimates to estimate the experiment's effects on total market surplus and find that, under plausible assumptions, the benefits of the experiment to firms and workers outweighed its social cost. This result obtains despite the fact that this experiment was not targeted to increase surplus, output created during the experimental treatment is not counted in the total, and I only consider benefits within eight weeks of the experiment. This result suggests that inefficiently low hiring of novice workers in this market led to diminished output and employment.

The interpretation of these results depends on the mechanism that generates them. The paper considers whether three alternative mechanisms could generate these outcomes: (1) the experimental jobs provided workers with human capital, (2) hiring the workers in itself gave the market a positive signal about their abilities, and (3) job receipt induced workers to apply to more oDesk jobs, but did not change employers' beliefs about worker ability. These mechanisms cannot explain all of the experiment's results. None of them can explain the effects of giving workers a more detailed evaluation. Moreover, the treatment is unlikely to have substantially increased workers' human capital because, on average, it only lasted 7.6 hours and represented a small increase in

workers' total offline work experience. Being hired in itself did not improve workers' outcomes. Job receipt only affected employment outcomes after the market observed the workers' evaluations, not when it observed only that they had been hired. The evaluations did change the market's assessment of these workers' value. After the experiment, treatment group workers were more likely to receive any given job they applied to than were control group workers.

The experimental results inform the literature on the effect of reputation on contracting. The theoretical literature shows that when legal or feasibility constraints prevent writing the optimal contract, agents' reputations affect the set of feasible outcomes (e.g., Kreps and Wilson, 1982; Tirole, 1996). Empirical research also shows reputations are important determinants of contracts (e.g., Banerjee and Duflo, 2000; McMillan and Woodruff, 1999; Resnick, *et al.* 2006). However, developing a reputation is a cost of market entry. An important question is to what extent this cost restricts economic activity. This is a particularly large concern in the trade of goods and labor with foreign countries where contracts are difficult to enforce and domestic firms have little familiarity with trading partners. In the oDesk marketplace, reputations are important because workers have limited liability constraints and cannot pledge their future output. The paper shows that developing a reputation is a significant barrier to entering this labor market, which reduces total market employment and output.⁴

The rest of the paper is organized as follows. Section 2 describes the online marketplace, Section 3 presents the theoretical framework, and Section 4 lays out the experimental design. Section 5 analyzes the worker-level effects of the experiment and analyzes whether they could be generated by alternative mechanisms. Section 6 presents the experiment's market-level effects and Section 7 concludes.

2 The Marketplace

oDesk is an online marketplace in which employers hire independent contractors to perform tasks remotely. The marketplace is large: as of July, 2010, oDesk workers completed approximately 200,000 hours of work per week, the equivalent of 5,000 full-time employees. oDesk workers are located around the world. A plurality (37%) lives in the United States, while India (15%) and the Philippines (14%) are the next most common countries of residence. Employers are also located around the world, but approximately 80% are located in the United States.

Employers post job openings in a wide variety of fields, the majority being in web programming, website design, and data entry. Some jobs last only a few hours, while others constitute full-time employment. The average oDesk job lasts 69 hours.

When posting a job, employers choose one of two pay schemes: "hourly" and "fixed wage." Hourly jobs, the type of jobs created in this experiment, constitute 70% of jobs on oDesk. In this type of employment, workers earn a fixed hourly wage and oDesk tracks the number of hours

⁴ As the majority of oDesk jobs are offshored, these results speak directly to the feasibility of offshoring. Reputation concerns do not limit all trade, but they prevent offshoring from reaching the efficient level.

worked. oDesk guarantees that the employer will pay for the hours worked regardless of the quality of the final product, though the employer can stop the job at any time or limit the number of hours worked. In fixed-wage jobs, workers and employers agree to a price for the entire project, but at the end of the project, employers have complete discretion over how much they pay. In these jobs, hours worked are not recorded.

Once a job is posted, workers can apply to the job directly or employers can invite them to apply. Under either method, the worker proposes a price: an hourly wage (in an hourly job) or an amount for the entire project (in a fixed wage job). After reviewing their applicant pools, employers may hire as many or few applicants as they deem suitable.

Workers post public profiles, describing their skills and the types of jobs they are seeking. To verify these skills, workers can take over 200 multiple-choice skills tests on topics such as Microsoft Word or C++ proficiency and post the results to their profiles. Each worker posts her preferred hourly wage rate at the top of her profile. Workers can suggest a different wage to employers when applying for jobs, but employers see their posted wages as well. An example worker profile is displayed in Figure 1.⁵

As soon as a hired worker begins working in an hourly job, the job title, number of hours worked, and hourly wage rate are automatically posted to the worker's profile.⁶ When a job is complete, employers and workers must rate each other from one to five on six dimensions: availability, communication, cooperation, deadlines, quality, and skills. These scores are averaged to form each party's overall rating for the job. Both composite ratings are automatically posted to the worker's profile.⁷ A worker cannot remove the ratings without refunding the remuneration received. Employers' ratings are typically very positive: 64% of workers receive a rating of exactly five, while 83% average at least four.

Worker and Employers may also choose to provide short comments about the employment experience. These are also automatically posted to the worker's profile. Comments are generally one to two positive sentences providing little objective information. Unlike the numerical ratings, workers can remove employer comments without penalty, but only 4% of them do.

3 Model

This section provides a simple framework that formalizes the insight that inefficiently few inexperienced workers will be hired when firms do not receive the benefit from discovering talented novices. It also develops implications that are testable in the oDesk setting.

⁵This worker was not included in the experiment because he did not join oDesk until after the experiment.

⁶In fixed-wage jobs, only the job title and agreed job price are posted.

⁷According to oDesk's website, oDesk posts the worker's evaluation of the firm beside the firm's evaluation of the worker in order to allow "both sides of the story" to be presented.

3.1 Model Set-Up

The marketplace comprises a mass 1 of firms and potential workers. Workers (indexed by i) live for two periods (a “novice” period and a “veteran” period) while firms live for one period. Each period, one generation of workers exits the market and a new generation enters. Each generation has mass $\frac{1}{2}$.

Workers vary in their ability (a_i), which is unknown to all market participants. Before the worker’s novice period, all market participants observe a normally distributed signal of the worker’s ability from the characteristics, x_i , listed on her resume. They use Bayesian updating to form their beliefs about worker ability:

$$a_i \sim N\left(\hat{a}_{i0}, \frac{1}{h_{x_i}}\right) \quad (1)$$

where \hat{a}_{i0} denotes a worker’s expected ability before her novice period. This expected ability is also distributed normally across workers.

Each firm (indexed by j) offers an identical task in which workers’ output is $y_i = a_i$. It must pay a firm-specific fixed cost c_j to hire a worker. This cost includes the time to explain the job to the worker and monitor her progress as well as any related overhead costs, such as for equipment or office space. It is continuously distributed across firms with cumulative distribution function F_c . A worker’s net marginal product in the task is

$$a_i - c_j. \quad (2)$$

The output of employed workers is imperfectly observable. All market participants observe a signal of output, \hat{y}_i :

$$\hat{y}_i = a_i + \varepsilon_{iy} \text{ where } \varepsilon_{iy} \sim N\left(0, \frac{1}{h_y}\right). \quad (3)$$

Firms update their beliefs about the abilities of novice workers using this signal and Bayesian updating. The term \hat{a}_i designates any worker’s expected ability, while \hat{a}_{i1} designates a worker’s expected abilities before her veteran period. The term F_0 is defined as the exogenous cumulative distribution function of novice workers’ expected abilities and F_1 as the endogenous cumulative distribution function of veteran workers’ expected abilities.

The market clears each period. Wages, w_{ij} , are restricted to be non-negative. For simplicity, workers and firms are risk neutral and do not discount the future. If firms do not fill their jobs, they receive 0, while if workers do not receive a job, they earn their outside option: $w_0 = 0$.

3.2 Market Equilibrium

Proposition 1 *The Perfect Bayesian Equilibrium of this game involves a threshold, \bar{c} , such that all firms with fixed costs $c_j \leq \bar{c}$ will hire a worker, while no firm with $c_j > \bar{c}$ will. All workers with expected ability $\hat{a}_i \geq \bar{c}$ and only these workers will be employed. These workers will earn wages $w_{ij} = \hat{a}_i - \bar{c}$.*

Risk-neutral firm j earns expected profit

$$\hat{\pi}_{ij} = \hat{a}_i - c_j - w_{ij} \quad (4)$$

from hiring worker i . Because novice expected ability is normally distributed, there exists a mass of novice workers whose expected ability is below every firm's fixed cost. No firm will incur an expected loss by hiring them at non-negative wages. Thus, not all workers will be hired and not all firms can hire a worker.

Firms vary only by their fixed costs. Therefore, there exists a fixed cost threshold, \bar{c} , such that all firms with $c_j \leq \bar{c}$ will hire a worker.⁸ All workers weakly prefer a market job to their outside option. To the market, they vary only by their expected abilities. Thus, there is an ability threshold, \underline{a} , such that all workers with $\hat{a}_i \geq \underline{a}$ will be hired.

Firms must all be indifferent to hiring any employed worker at her market wage, so all employed workers must be offered a wage

$$w_{ij} = \hat{a}_i - \bar{w} \quad (5)$$

for some constant \bar{w} . The marginal firm must earn zero profit from hiring this worker so $\bar{w} = \bar{c}$ and wages equal

$$w_{ij} = \hat{a}_i - \bar{c}. \quad (6)$$

The threshold worker must earn a zero wage; otherwise firms could earn a higher profit by offering slightly lower-ability workers lower wages. This implies $\underline{a} = \bar{c}$. These thresholds are determined by equating the number of hired workers with the number of filled jobs:

$$\underbrace{\frac{1}{2}(1 - F_0(\bar{c})) + \frac{1}{2}(1 - F_1(\bar{c}))}_{\text{Supply of Labor}} = \underbrace{F_c(\bar{c})}_{\text{Demand for Labor}} \quad (7)$$

where F_1 is the cumulative distribution function of veteran workers when all novice workers with $\hat{a}_{i0} \geq \bar{c}$ are hired. The left-hand side of this equation is the mass of workers hired: the mass of novice and veteran workers with $\hat{a}_i \geq \bar{c}$. The right-hand side is the mass of firms hiring: the mass of firms with $c_j \leq \bar{c}$.

3.3 The Social Planner's Solution

Proposition 2 *The desired equilibrium of a social planner who has the same information as the market and maximizes market surplus is as follows. There exists a threshold, $c^* > \bar{c}$, such that every firm with $c_j \leq c^*$ and only these firms hire workers. All veteran workers with expected ability $\hat{a}_{i1} \geq c^*$ are employed as are all novice workers with $(\hat{a}_{i0} - c^*) + \text{Pr}(\hat{a}_{i1} \geq c^*) \times E[\hat{a}_{i1} - c^* | \hat{a}_{i1} \geq c^*] \geq 0$.*

⁸Throughout the paper, there are measure zero sets of workers and firms that are indifferent between entering the market and taking their outside options (e.g., firms with $c_j = \bar{c}$). Throughout the paper, I use weak inequalities indicating that these agents enter the market. However, no other features of the equilibria would change if these measure zero sets of workers and firms did not enter.

More novice workers are employed than in the market equilibrium.

The social planner will have only the most efficient firms hire workers. Thus, there is a threshold, c^* , such that every firm with $c_j \leq c^*$ will hire a worker and no firm with $c_j > c^*$ will. Veteran workers are hired only when hiring them weakly increases expected worker surplus in the current period. That is, when

$$\hat{a}_{i1} \geq c^*. \quad (8)$$

In the social planner's solution, novice workers will be hired if the expected benefit from hiring them, calculated over both periods of life, is non-negative. All novices expected to generate positive surplus in the current period ($\hat{a}_{i0} \geq c^*$) will be hired. Workers expected to generate a loss in their novice period ($\hat{a}_{i0} < c^*$), but whose veteran-period option value is greater than this expected loss will also be hired. Novice workers will be hired if

$$\underbrace{(\hat{a}_{i0} - c^*)}_{\text{Novice-Period Benefit}} + \underbrace{\Pr(\hat{a}_{i1} \geq c^*) \times E[\hat{a}_{i1} - c^* | \hat{a}_{i1} \geq c^*]}_{\text{Veteran-Period Option Value}} \geq 0. \quad (9)$$

The first term in this equation is the expected current-period benefit of hiring a novice worker. The second term is the option value of being hired for workers with $\hat{a}_{i0} < c^*$. (This equation is automatically satisfied for workers with $\hat{a}_{i0} \geq c^*$.) If not employed in their novice periods, these workers will not be employed in their veteran periods. However, if employed in their novice periods, these workers will be hired with probability $\Pr(\hat{a}_{i1} \geq c^*)$ in their veteran periods and generate expected surplus $E[\hat{a}_{i1} - c^* | \hat{a}_{i1} \geq c^*]$ if hired. In contrast, in the market equilibrium, because firms do not value workers' veteran-period option value, novice workers are hired only if their novice-period benefit is non-negative.

The social planner's cutoff, c^* , is determined by the market clearing condition. Define N_{c^*} as the fraction of novices for whom Equation 9 holds and F_1^* as the cumulative distribution of veteran ability if novices for whom Equation 9 is satisfied are hired. Then, c^* is chosen so that

$$\underbrace{\frac{1}{2}N_{c^*} + \frac{1}{2}(1 - F_1^*(c^*))}_{\text{Supply of Labor}} = \underbrace{F_c(c^*)}_{\text{Demand for Labor}}. \quad (10)$$

For any threshold, c , more novices and veterans are hired in the social planner's solution than in the market equilibrium, so there would be an excess supply of labor in the social planner's solution if $c^* = \bar{c}$. The social planner would want to employ novices with $\hat{a}_{i0} < \bar{c}$, which would lead to a larger mass of veterans with $\hat{a}_{i1} \geq \bar{c}$. To equilibrate supply and demand, more firms must hire workers in the social planner's solution than the market equilibrium: $c^* > \bar{c}$. With more firms hiring workers and a higher veteran hiring threshold than in the market equilibrium, the social planner's solution must involve hiring more novice workers than the market equilibrium. Otherwise, there would be excess demand for labor.

The social planner's solution would be enacted as the competitive equilibrium if workers could

accept negative wages in their novice period. In this case, novices employed in the social planner's solution, but not by the constrained market, would receive wages of

$$w_{ij} = \hat{a}_{i0} - c^* < 0. \quad (11)$$

3.4 Adverse Selection

In the above model, wages for the marginal worker must equal zero or (equivalently) a binding minimum wage, for the inefficiency to be present. This is a function of the simplified set-up, not a true requirement for the model. This inefficiency will be present at higher wages if a wage rigidity prevents wages from dropping below a certain level. I consider a simple extension in which potential adverse selection of workers keeps wages artificially high.

While adverse selection is a more general problem, I illustrate its effects with a simple extension of the model with three groups of workers: A , B , and C , which comprise fractions γ_A , γ_B , and γ_C of the population, respectively. A worker's ability, a_i , and outside option, w_0 , depend on to which of three groups a worker belongs. Group A and B workers have the same (non-zero) distribution of ability and outside options $w_0 = w_A$ and w_B , respectively that satisfy

$$w_A > w_B > 0. \quad (12)$$

Group C workers all have ability $a_i = 0$ and an outside option $w_0 = 0$. These groups represent the fact that workers with higher productivity in a given market will, in general, have better outside options. However, outside options vary among workers with the same market productivity. For simplicity, in this section, I assume that employers face equal uncertainty about each novice worker: $h_{x_i} \equiv h_x$ for all i . The rest of the model is unchanged.

Before workers' novice period, firms observe characteristics x_i from each worker's resume. Firms do not know workers' group affiliations and these characteristics provide no information about affiliations. Workers and firms have the same information except that workers know their group affiliations because they know their outside options. Firms learn workers group affiliations when workers are employed.

Define \hat{a}_i as firms' expectation of a worker's ability if she were known to be in group A or B . Define $\hat{\eta}_i$ as firms' expectation of the ability of workers who will accept a given wage offer, conditional on \hat{a}_i and their beliefs about workers' strategies. For veteran group A or B workers who worked in their novice period, $\hat{a}_{i1} = \hat{\eta}_{i1}$.

I make two assumptions about the model's parameters. The first guarantees that group A 's outside option is sufficiently high that novice group A workers accept their outside options over some positive market wages. If group A workers never accept their outside option over positive wage offers, then there is no adverse selection problem. The second guarantees that labor demand is sufficiently low that at the ability levels at which group A workers select out of the market, firms will not hire only group B and C workers. Without this assumption, adverse selection would not

lead to reduced market employment. Before specifying these assumptions mathematically, I define two functions, $c(a)$ and $G(a)$.

The function $c(a)$ gives the fixed cost of the marginal hiring firm if all novices with $\hat{a}_{i0} \geq a$, all veteran group A workers with $\hat{a}_{i1} - c \geq w_A$, and all veteran group B workers with $\hat{a}_{i1} - c \geq w_B$ are hired. It is defined implicitly by

$$\underbrace{\frac{1}{2}[1 - F_0(a)]}_{\text{Novice Workers Hired}} + \underbrace{\frac{1}{2}\gamma_A[1 - F_{1a}(w_A + c(a))]}_{\text{Veteran Group A Workers Hired}} + \underbrace{\frac{1}{2}\gamma_B[1 - F_{1a}(w_B + c(a))]}_{\text{Veteran Group B Workers Hired}} = \underbrace{F_c(c(a))}_{\text{Hiring Firms}}. \quad (13)$$

where F_0 is the cumulative distribution function of \hat{a}_{i0} for group A workers and F_{1a} is the cumulative distribution function of \hat{a}_{i1} for group A veterans if all group A novices with $\hat{a}_{i0} \geq a$ are hired. The left-hand side of the equation is the mass of hired workers and the right-hand side is the mass of hiring firms.

The function $G(a)$ is the expected net surplus a group A novice worker with expected ability $\hat{a}_{i0} = a$ receives from working when the workers listed above are also employed. That is

$$G(a) = \underbrace{((\gamma_A + \gamma_B)a - c(a) - w_A)}_{\text{Net Wage in Novice Period}} + \underbrace{\Pr(\hat{a}_{i1} - c(a) \geq w_A) \times E[\hat{a}_{i1} - c(a) - w_A | \hat{a}_{i1} - c(a) \geq w_A]}_{\text{Veteran-Period Option Value}}.$$

The worker's novice-period wages equal her expected net marginal product at the marginal firm. Her expected ability is $(\gamma_A + \gamma_B)a$ because all three groups of novices with this expected ability work in the market. In her veteran period, her group affiliation is known and she will work in the market when her market wages exceed her outside option. Finally, I define \bar{a}_A as the expected ability level at which a novice group A worker is indifferent between working in the market and accepting her outside option:

$$G(\bar{a}_A) \equiv 0. \quad (14)$$

Then, the two assumptions are:

$$(\gamma_A + \gamma_B)\bar{a}_A - c(\bar{a}_A) > 0 \quad (15)$$

$$\left(\frac{\gamma_B}{\gamma_B + \gamma_C}\right)\bar{a}_A - c(\bar{a}_A) < 0. \quad (16)$$

The left-hand side of Equation 15 is the market wage received by the indifferent novice group A worker. I assume this wage is positive. The left-hand side of Equation 16 is the expected marginal product of group B and C workers with $\hat{a}_{i0} = \bar{a}_A$ at the marginal hiring firm. I assume it is negative. This ensures firms will not hire workers with $\hat{a}_{i0} < \bar{a}_A$: at these expected ability levels they would only hire group B and C workers who would not generate positive net output.

Proposition 3 *If Equations 15 and 16 are satisfied, then the Perfect Bayesian Equilibrium of this game is as follows. There exists a threshold, \bar{c} , such that all firms with $c_j \leq \bar{c}$, and only these firms hire a worker. There exists a threshold, \bar{a}_A , such that all novice workers with $\hat{a}_{i0} \geq \bar{a}_A$ are hired*

at a wage of $(\gamma_A + \gamma_B)\bar{a}_A - \bar{c} > 0$ and workers with $\hat{a}_{i0} < \bar{a}_A$ remain unemployed. At this wage, group A novice workers with expected ability $\hat{a}_{i0} = \bar{a}_A$ are indifferent to taking a job, but group B and C workers with $\hat{a}_{i0} = \bar{a}_A$ strictly prefer a market job to their outside options. Group B and C workers with $\hat{a}_{i0} < \bar{a}_A$ would be willing to work for lower wages, but firms will not hire them. Total social surplus is maximized when some novice workers with $\hat{a}_{i0} < \bar{a}_A$ are also hired.

Proof. See Appendix A. ■

3.5 Testable Predictions

The model makes predictions about the effects of the experiment on both treated workers and the market. In the first treatment, novice workers are hired at random from the applicant pool. I then compare the employment outcomes of the treatment and control groups in their veteran periods. The model makes the following testable predictions about the employment outcomes of veteran treatment group workers relative to veteran control group workers.

1. Veteran treatment group workers will have higher employment rates.
2. Veteran treatment group workers will have higher earnings.
3. Veteran treatment group workers will have higher reservation wages.
4. The treatment will have larger effects on the employment, earnings, and reservation wages of workers about whom the market is more uncertain, conditional on ability. Unconditional on ability, the model makes no predictions about the relative sizes of these treatment effects.

Treatment group workers have a range of values of \hat{a}_{i0} . Figure 2 is a schematic representation of the distribution of novice workers' expected abilities. (For clarity, it assumes that the market faces the same uncertainty about the abilities of all novice workers, so that the option value of being hired and the social planner's novice hiring threshold depend only on workers' expected abilities.) Workers in the unshaded region would be hired by both the market and the social planner, workers in the light gray shaded region would be hired only by the social planner, and workers in the dark gray shaded region would be hired by neither the market nor the social planner. Treatment group workers come from all three sections of the graph.

Control group workers in the unshaded region (with $\hat{a}_{i0} \geq \bar{c}$) will be hired by other employers in their novice periods, so there is no difference between the veteran-period outcomes of treatment and control group workers in this region. Their employment, earnings, and reservation wages will be the same, on average. However, no control group worker in either of the gray shaded regions (with $\hat{a}_{i0} < \bar{c}$) will be hired in her novice period. Thus, their expected abilities will remain below the hiring threshold and they will have no earnings or employment in their veteran period. On the other hand, some treatment group workers with $\hat{a}_{i0} < \bar{c}$ will perform well enough in their treatment jobs that their veteran-period expected abilities will exceed the hiring threshold and they will

have positive veteran-period employment and earnings. Thus, as a group, veteran treatment group workers will have higher employment and earnings than veteran control group workers.

The reservation wages of veteran workers with expected ability below the hiring threshold ($\hat{a}_{i1} < \bar{c}$) is zero. These workers will not receive any positive wage offers, so they will be willing to accept any offer at least as high as their outside option. Veteran workers with expected ability above the hiring threshold have reservation wages equal to the positive wages they receive in the market ($\hat{a}_{i1} - \bar{c} > 0$). Receiving a treatment job affects the reservation wages of only workers in the gray shaded areas. By increasing some of these workers' expected abilities above the threshold, it increases their reservation wages above zero.

Novice workers in the gray shaded areas realize an increase in employment, earnings, and reservation wages of

$$\underbrace{\Pr(\hat{a}_{i1} \geq c^*)}_{\text{Increase in Employment}} \tag{17}$$

$$\underbrace{\Pr(\hat{a}_{i1} \geq c^*) \times E[\hat{a}_{i1} - c^* | \hat{a}_{i1} \geq c^*]}_{\text{Increase in Earnings and Reservation Wages}}. \tag{18}$$

Conditional on the novice worker's expected ability, \hat{a}_{i0} , both expressions are increasing in the variance of veteran expected ability, \hat{a}_{i1} . This simply says that the option value of being hired is increasing in the uncertainty about expected ability. However, these expressions vary directly with novice expected ability. If expected ability differs among workers about whom the market is more and less uncertain, the relative sizes of these expressions for workers about whom the market is more and less uncertain is indeterminate.

The second experimental treatment provides the market with more precise information about the performance of (randomly-selected) workers. This represents an increase in the precision of the signal \hat{y}_i (an increase in h_y) for randomly-selected workers. There are two testable predictions from this treatment that compare the veteran-period employment outcomes of treatment group workers (those who received detailed evaluations) to control group workers (hired workers who received coarse evaluations).

5. Treatment group workers who receive positive detailed evaluations will have higher employment, earnings, and reservation wages in their veteran period than control group workers. Treatment group workers who receive negative detailed evaluations will have lower employment, earnings, and reservation wages.
6. Overall, veteran treatment group workers will have higher earnings and reservation wages than veteran control group workers. The treatment's effect on workers' employment is indeterminate.

Revealing positive information about workers' job performance increases their veteran-period expected abilities. Thus, they will all have higher employment, earnings, and reservation wages. Revealing negative additional information leads to the opposite results.

As a group, veteran workers' total earnings equal the total amount they produce. A more precise signal increases the market's ability to distinguish between high-ability and low-ability workers. This increases the fraction of high-ability workers employed, increasing the number of workers who generate positive marginal products, and decreases the fraction of low-ability workers employed, decreasing the number of workers with negative marginal products. The result is an increase in the total amount produced by these workers and thus their earnings. Since workers' earnings equal their reservation wages, the treatment increases workers' reservation wages as well.⁹

If the market received only an imprecise signal of their novice-period job performance, some low-ability workers would be hired in their veteran periods. A more precise signal of job performance can decrease these workers' veteran-period expected abilities sufficiently that they are unemployed in their veteran periods. However, other workers would not be hired in their veteran periods if the market only received an imprecise signal of their performance. A more precise signal can increase these workers' veteran-period expected abilities sufficiently that they are employed in their veteran period. The relative size of these two effects (and the net effect of detailed comment receipt on employment) depends on the distribution of ability.

Finally, I consider the market-level effects of employing many novice workers. Employing these novices increases the mass of veteran workers with $\hat{a}_{i1} \geq \bar{c}$ (the supply of workers recognized to be high ability). This generates two predictions.

7. Total market employment will increase as a result of the experiment by less than the increase in treatment group workers' employment.
8. Average wages will decrease, conditional on ability.

Because the experiment increases the fraction of veteran workers whose expected abilities exceed the hiring threshold, \bar{c} , there is an excess supply of labor after the experiment. To clear the market, the hiring threshold increases to c' . As the number of firms hiring workers increases, total employment must increase. However, as the hiring threshold increases, non-experimental workers with $\bar{c} \leq \hat{a}_i < c'$ become unemployed. Thus, total market employment will increase less than the increase in employment of treatment group workers. Wages, conditional on ability, equal workers' marginal product at the marginal firm:

$$w_{ij} = \hat{a}_{i1} - \bar{c}. \tag{19}$$

When \bar{c} increases to c' , wages decrease.

3.6 Market Exit

This stylized model does not include the worker's choice to exit the market because including this choice adds little insight to the model. However, because market exit is an important, testable,

⁹Providing a detailed comment only affects workers' mean earnings because some workers are unemployed. If all workers were employed, the provision of a detailed comment would increase only the variance of earnings, not mean earnings.

outcome, I extend the model in this section to include this choice.

Each worker faces a distinct decision of whether to exit the market in each period of life. Before each period, a worker observes her period-specific outside option. With probability κ , her outside option is $w_1 > 0$; with probability $1 - \kappa$, her outside option is $w_0 = 0$. The worker first decides whether to accept the outside option. If she does, she exits the market. Otherwise, she remains in the market, costlessly signaling to firms she is available for jobs. After workers decide whether to exit the market, the market clears. Workers who remain in the market, but do not receive wage offers receive zero.

This section also introduces randomness in whether workers receive jobs. After the market clears, with probability ε , each worker who has not received a job receives a wage offer of $\delta > 0$.¹⁰

Proposition 4 *Veteran workers who have received a treatment job are less likely to exit the market than veteran control group workers. Veteran workers who receive a detailed evaluation may be either more or less likely to exit the market than workers who receive a coarse evaluation, depending on the relevant parameter values.*

Proof. See Appendix B. ■

4 Experimental Context and Design

4.1 Sample Selection

I recruited subjects for this experiment by posting hourly data-entry jobs to the marketplace and inviting workers to apply.¹¹ The job was expected to take approximately 10 hours and involved entering Census records from a PDF file into a Microsoft Excel spreadsheet. I invited applications from every oDesk worker who had a public profile, listed her specialty as data entry, posted an hourly wage of \$3 or less to her profile, and had applied for at least one job in the prior three months. Because hiring so many workers at one time would be both logistically difficult and a large shock to the market, I contacted workers in two waves, two weeks apart. Workers were randomly allocated to a wave. The workers who responded to the invitation and applied to the jobs, requesting a wage of \$3 or less, form the experimental sample.

Table 1 shows the sample selection. Slightly fewer than 10,000 workers fit the sample selection criteria, most of whom had never had an oDesk job. Thirty-nine percent of the workers applied to the jobs, all but 85 of whom requested a wage of \$3 or less. Workers with prior oDesk experience were substantially more likely to apply than inexperienced workers.

The final column of Appendix Table 1 provides sample summary statistics. The majority of the sample lived in the Philippines, while India, Bangladesh, and Pakistan were also well-represented.

¹⁰The randomness in hiring reflects that some firms hire workers without reading workers' resumes. It only affects the model by ensuring that workers with outside option $w_0 = 0$ do not exit the market and, thus, removing the multiplicity of equilibria.

¹¹I posted these jobs from the accounts of 23 different employers. Each employer posted 10 separate (but identical) jobs, so that no one employer or job applicant pool would appear too large. Workers in the sampling frame were randomly assigned to an employer and job.

Workers from the United States comprised 2.6% of the sample. The average worker had been in the marketplace for six months, sent 27 job applications, and passed 3.5 skills tests before the experiment. Workers with prior oDesk jobs had been in the marketplace twice as long and had sent seven times as many applications as inexperienced workers. The mean experienced worker had seven prior oDesk jobs, while the median experienced worker had four.

4.2 Experimental Protocol

There were three randomizations in this experiment. First, I randomized which workers I hired. Second, I randomized the amount of information I provided to the market about the job performance of hired workers. Third, three weeks after the initial hiring I invited randomly-selected experimental workers to apply to another data-entry job to assess their reservation wages.

In total, 952 workers, one quarter of the experimental sample, were selected to receive treatment jobs. The randomization was stratified on whether workers had a previous oDesk job: inexperienced workers were selected to receive treatment jobs with higher probability than were experienced workers. Hired workers were instructed to send an email to the employer (me) to receive a PDF file with the data to be entered, the data entry instructions, and a data entry template. They were given a maximum of 10 hours over one week to enter the data. They were told that, if they could not complete the task within 10 hours, they should send the file back unfinished.

When workers sent back the files, I recorded several measures of their job performance: their data entry speed, their error rate, the date they returned the data file, and three measures of whether they had followed the data entry instructions. I used this information to rate all the workers on a one-to-five scale. The distribution of scores from my job was designed to match the distribution of scores of low-wage data entry workers in the marketplace, adjusted for the fact that a worker in my sample was more likely to be inexperienced than a typical oDesk worker.¹² Approximately 18% of workers did not return the file or log any hours. Under oDesk’s protocol, these workers are not rated.

While all hired workers were rated using identical criteria, they did not all receive the same type of employer comment. Workers who earned a rating of least four were randomly selected to receive either an informative comment or an uninformative one. No workers with a rating below four received a detailed comment due to human subject concerns.¹³

The uninformative comment was chosen to be short and positive, like most of the comments in the marketplace. It read (where only the words in brackets vary by worker):

¹²In fact, the distributions of feedback scores received by experienced and inexperienced workers are not statistically distinct.

¹³The human subjects committee was concerned that giving workers negative evaluations would harm workers. They allowed me to give low numerical ratings (which were essential to the experiment). However, they permitted me to provide detailed information about the performance only of workers who did well overall on the task. There is no censoring of the comment for people receiving a rating of four or above, so the detailed comments do provide negative information about these workers (e.g., they were in the bottom 10% in speed or accuracy). All analyses assessing the effect of the detailed comment condition on workers having received a rating of at least four and, thus, being eligible to receive a detailed comment.

“It was a pleasure working with [x].”

The detailed comment provided objective information on the worker’s data entry speed, accuracy, ability to meet deadlines, and ability to follow instructions. Additionally, it repeated the uninformative comment, so the only difference between the two comment types was the additional objective information. The detailed comment read:

“[x] completed the project [y days before the deadline, by the deadline, z days after the deadline] and [followed our instructions perfectly, followed our instructions, followed most of our instructions, did not follow our instructions]. [x] was in the [top 10%, top third, middle third, bottom third, bottom 10%] of providers in speed and the [top 10%, top third, middle third, bottom third, bottom 10%] in accuracy. It was a pleasure working with [x].”

Because such a high fraction of oDesk workers receive a rating of five, many workers who received this rating were in the bottom third of speed, accuracy, or both. If employers, particularly those new to the oDesk market, did not realize that a high fraction of oDesk workers receive a rating of five, these detailed comments would appear very negative. Workers have the option to remove the detailed comment, but in the marketplace as a whole only 4% of comments are removed. Workers who received a rating between three and four all received uninformative comments, while those who received a score less than three received no comment at all.

I also randomized whether the comment revealed that the employer had extensive experience hiring workers on oDesk. This was to test whether any effect of the detailed comment could be attributed to the market believing that firms who left a detailed comment were larger and more competent. The uninformative comment that provided information about the employer read:

“It was a pleasure working with [x]. Our organization has hired hundreds of providers on oDesk.”

The analogous informative comment read:

“[x] completed the project [y days before the deadline, by the deadline, z days after the deadline] and [followed our instructions perfectly, followed our instructions, followed most of our instructions, did not follow our instructions]. [x] was in the [top 10%, top third, middle third, bottom third, bottom 10%] of providers in speed and the [top 10%, top third, middle third, bottom third, bottom 10%] in accuracy. This is based on our experience with hundreds of providers on oDesk. It was a pleasure working with [x].”

Figure 3 gives a schematic representation of these treatments. Dashed connecting lines indicate random assignment while solid connectors indicate non-random assignment.

To elicit workers’ reservation wages, three weeks after the initial randomization, I randomly selected 630 workers from my sample to contact again from the account of a different employer. I

contacted both workers who received treatment jobs and those who had not. These workers were invited to apply to another data-entry job requiring skills similar to the initial job. Unlike the previous job, however, these jobs had fixed wage rates: some paid \$0.75 per hour, some paid \$1 per hour, and some paid \$2 per hour. The wage of the job each worker was invited to apply to was randomly selected. I recorded who accepted the invitation and offered a job to a randomly-selected 5% of applicants.

4.3 Data Collection

I hand-collected data on workers' job performance and whether they accepted the invitation to apply to the subsequent jobs. The remaining worker characteristic and outcome data used in this project are administrative data obtained from oDesk's server with oDesk's permission. oDesk's server automatically records information on workers' employment, job applications, and profiles. The primary worker-level outcomes of interest are measures of workers' employment, earnings, and reservation wages.

I consider three measures of employment: whether a worker obtained any job after the experiment, the number of jobs obtained, and the number of hours worked (in hourly jobs). In addition to assessing workers' reservation wages directly by testing their willingness to apply to subsequent jobs, I use the wages workers post to their profiles as a measure of their reservation wages. All workers must post a wage on their profiles, so this measure is free from selection concerns. In a fully competitive market, the wage workers advertise would be their reservation wage. However, even if workers post wages above their reservation wages, there is no reason to believe the treatment affects the relationship between workers' posted wages and their reservation wages. Finally, I evaluate the effect of the treatments on workers' earnings from all jobs.

5 Worker-Level Effects

5.1 Effect of Job Receipt on Employment Outcomes

In this section, I evaluate the effect of the first treatment – providing workers with employment in their “novice period” – on their earnings, employment, and posted wages in the two months after the experiment: their “veteran period.” The model predicts that employing these workers allows the market to identify high-ability workers, generating an option to hire the high-ability workers as veterans. Because these newly-identified high-ability workers will be rehired, the treatment increases workers' employment rates. Since workers receive the option value of being hired, the treatment also increases workers' earnings and reservation wages.

Appendix Table 1 shows that the randomization produced similar treatment and control samples on pre-experiment observable characteristics (conditional on the stratification). Because the randomization was stratified on whether workers had prior oDesk experience, all results are conditional on this variable.

Table 2 compares the employment, earnings, and posted wages of the treatment and control groups *excluding* the experimental jobs in the two months following the experiment. It shows that obtaining a treatment job substantially increased the future employment of both experienced and inexperienced workers. The treatment increased the probability that an inexperienced worker obtained any job in the subsequent two months from 12% to 30%. It also almost tripled the number of jobs the average inexperienced worker obtained from 0.28 to 0.81. Workers with prior oDesk experience generally worked much more following the experiment than did inexperienced workers. Obtaining a treatment job increased the probability an experienced worker was employed in the subsequent two months from 55% to 61% and increased the average number of jobs she obtained by 25% from 2.0 to 2.5.

Receiving a treatment job increased the average wage inexperienced workers posted on their profiles, a proxy for their reservation wages, from \$2.03 to \$2.28. This is an increase of approximately 12%. Among experienced workers, the treatment group’s average posted wage was \$0.12 larger than the control group’s, but this difference is not significant. The treatment almost tripled the earnings of inexperienced workers from \$10.06 to \$28.43. This increase in earnings of \$18.37 exceeds the \$16.03 I spent, on average, to hire these workers. Among experienced workers, the treatment group’s earnings were more than \$23 higher than the control group’s, which also exceeds the \$18.72 I spent, on average, to hire them, but this treatment effect is imprecisely measured.

The treatment’s effect on workers’ earnings corresponds, in the model, to the option value of hiring them. It is striking that it is so large given that the treatment was not targeted towards workers thought to have particularly high option values. Some treatment group workers had such low expected abilities and option values that they were in the dark gray region of Figure 2 and would not have been hired by the social planner. Others had expected abilities in the unshaded region and were already recognized as high ability before the experiment. The large average option value of hiring these workers suggests that some of them may have been inefficiently employed before the experiment. However, workers are only inefficiently unemployed if the *social* benefit of hiring them exceeds the *social* cost. The increase in the treatment group’s earnings overstates the benefits to workers if some of the jobs they obtained as a result of the experiment would have otherwise gone to other workers or if treatment group workers face a non-zero cost of effort. However, it does not include any increase in firm profits. I estimate the experiment’s effect on market surplus in Section 6.

The model predicts that the treatment’s effects should be larger for workers about whom the market is more uncertain, conditional on expected ability. The market should be more uncertain about the ability of workers lacking experience, suggesting that the treatment should have larger effects for these workers. While the treatment effects are much larger in percentage terms for inexperienced workers, in absolute terms, the effect for inexperienced workers is only statistically different (larger) for one of the five outcomes: the probability that workers obtained an additional job. This is most likely because workers with and without previous jobs do not have the same expected abilities. The model predicts the treatment will have the largest effects for workers with

expected abilities close to, but below, the hiring threshold. Workers with previous jobs had expected ability above the hiring threshold when they were first hired. After their jobs, if their expected abilities are no longer above the hiring threshold, they will be relatively close to it. On the other hand, workers without prior oDesk experience may have expected abilities relatively far away from the hiring threshold.¹⁴

Table 3 estimates the effect of job receipt on the employment outcomes of the entire sample. It displays estimates of the regression

$$y_i = \alpha + \beta job_i + \gamma_1 prevjob_i + X_i \gamma_2 + \varepsilon_i \quad (20)$$

where i indexes the worker, y_i is one of the five employment outcomes analyzed in Table 2, and job_i is an indicator for whether the worker received the treatment job. An indicator for whether the worker had a previous job, $prevjob_i$, is included because of the stratification, while ε_i is the mean-zero error term. The vector X_i contains worker characteristics measured before the experiment. The regressions in Panel A do not include any worker characteristics except for a dummy for whether the worker was in the second (randomly-selected) wave of the experiment. Regressions in Panel B add controls for the wage the worker requested when applying to the treatment job, controls for how long the worker had been on oDesk before the experiment, and country fixed effects.¹⁵ Panel C additionally includes dummies for whether workers took each of the six most popular oDesk skills tests among low-wage data-entry workers and whether they scored above average on each of these tests.¹⁶

The table shows that the treatment significantly increases the employment, earnings, and reservation wages of experienced and inexperienced workers taken together. The estimates are very similar across panels, which is consistent with random assignment and the balanced covariates of the treatment and control groups. While the treatment effects are large, they are smaller than the coefficients on having had a previous job, γ . This is as expected given that this coefficient includes both the treatment effect of having had several previous jobs and the effect of differential selection into employment.

The employment outcomes of treatment group workers vary greatly with the rating I gave them. After two months, workers who received ratings of five earned \$34 more than the control group, while workers who received ratings of one and two earned \$23 less. This is consistent with employers learning about worker quality from the ratings. Alternatively, it could be that workers who received scores of five are more talented than the average member of the control group and

¹⁴I have attempted to test this explanation by estimating the treatment effect conditional on expected ability. However, the worker characteristics in the data are not very predictive and, thus, the standard errors are too large to meaningfully test this explanation.

¹⁵These regressions include dummies for whether the worker offered a wage of \$1 to \$1.99, a wage of \$2 to \$2.99, and a wage of exactly \$3. The excluded category is offering a wage less than \$1. They also include dummies for whether the worker had been on oDesk from 60 to 120 days, 121 to 180 days, 181 to 240 days, or more than 240 days. The excluded category is having joined oDesk in the last 60 days.

¹⁶The six tests were each taken by over 1,000 of the experimental workers. They include two tests about how oDesk works, two English tests, a Windows XP test, and a Microsoft Office test.

would have earned more even without the rating, while workers who received ratings of one or two are less talented.

5.2 Effect of Detailed Evaluations on Employment Outcomes

This section analyzes the effect on workers’ future employment outcomes of providing the market with more detailed information about their job performance. If the content of the comment affects the market’s expectation of worker ability, revealing positive information about workers’ performance will improve their employment outcomes, while revealing negative information will worsen their employment outcomes. Because providing this information allows employers to identify and hire the high-ability workers, the model predicts that workers who received detailed comments should have higher average earnings and reservation wages after the experiment than workers who received coarse comments.

Appendix Table 2 shows that the sample chosen to receive detailed comments was similar on observables to the sample chosen to receive uninformative comments. It considers only workers who received a rating of four or five because only these workers were eligible to receive detailed comments.

I determine the effect of receiving a particular type of detailed comment by comparing the employment outcomes of workers with the same performance who received coarse and detailed comments. I estimate the following regression:

$$y_i = \alpha + \beta_1(x_i \times detailed_i) + \beta_2x_i + \beta_3detailed_i + \varepsilon_i. \quad (21)$$

The dependent variable, y_i , is one of the five employment outcomes examined in Table 2. The variable $detailed_i$ indicates that a worker received a detailed comment while x_i is a positive measure of worker performance. The regression is limited to workers who received a score of at least four because receipt of a detailed comment was randomized only among these workers. The model predicts that β_3 , the effect of the detailed comment for workers without the (positive) characteristic (e.g., workers who did not meet the deadline), is negative. It predicts that β_1 , the additional effect of the detailed comment for workers with these (positive) characteristics, is positive as is $\beta_1 + \beta_3$, the overall effect of the comment for workers with each characteristic.

Table 4 displays the results of estimating this equation where x_i is an indicator for meeting the deadline. It shows that revealing that workers did not meet the deadline decreased earnings by \$35 on average. Revealing that a worker met the deadline increased earnings by \$59 relative to revealing she did not and by \$24 relative to leaving only a coarse comment. These are large effects; they are even bigger than the effect of receiving a job. There are large effects consistent with the model on the other employment outcomes as well, though except for hours worked, they are imprecisely measured.

This is the only piece of information that appeared to matter to employers. Appendix Table 3 shows that revealing that workers followed instructions, were in the top third of workers in speed, or

were in the top third of workers in accuracy did not affect employment outcomes. This is most likely because meeting deadlines was the first piece of information revealed in the comment. Because of the length of the comment, the parts about speed and accuracy were not immediately visible on most workers’ profiles; one had to click on the continuation to see them.

To test whether receiving a detailed comment increased average earnings and reservation wages, I estimate the following regression:

$$y_i = \alpha + \beta_1 \text{detailed}_i + \beta_2 \text{rate45}_i + \beta_3 \text{job}_i + \gamma_1 \text{prevjob}_i + X_i \gamma_2 + \varepsilon_i \quad (22)$$

where the variable rate45_i indicates that workers received a rating of at least four. The worker characteristics, X_i , are the same controls included in Table 3 and β_1 is the coefficient of interest. Table 5 displays the results. The coefficient estimates are similar in all three panels as one would expect. The table shows that receiving a detailed comment led to a large increase in earnings (\$23.85) for the average worker within two months. It also led to an increase in workers’ reservation wages measured by the wages posted on their profiles.¹⁷ The measured effect of the treatment on the number of jobs obtained is large and positive, but far from significant.¹⁸

5.3 Reservation Wage Treatment

In addition to using the wages posted on workers’ profiles as a proxy for their reservation wage, I directly measured workers’ reservation wages by inviting a randomly-selected sample to apply for additional jobs three weeks later. These jobs had a stated hourly wage, either \$0.75, \$1, or \$2, and workers were randomly assigned to invitation pools with these different wages. To test the prediction that receiving a treatment job increased workers’ reservation wages, I estimate the following regression:

$$\text{apply}_i = \alpha + \beta_1 \text{job}_i + \beta_2 \text{prevjob}_i + \gamma_w + \gamma_w \times \text{job}_i + \gamma_w \times \text{prevjob}_i + \varepsilon_i. \quad (23)$$

Here, apply_i is a dummy for whether the worker applied for the job.¹⁹ The variables γ_w are fixed effects for the job’s wage and $\gamma_w \times \text{job}_i$ are interactions of this wage with an indicator for receiving a treatment job. Because I stratified job receipt on having a previous job, I must also interact these wage dummies with whether workers had a previous job.

Table 6 shows the results of this regression. Receiving a treatment job made workers 14 percentage points more likely to apply to these jobs. This is consistent with the model extension that predicts that receiving a treatment job will induce workers to remain in the market. Workers with

¹⁷Both of these coefficients are significant only at the 10% level. The increase in earnings after one month (\$13.46) is significant at the 5% level.

¹⁸It is significant at the 10% level after one month.

¹⁹The invitation to apply to each job explained that the job paid a stated wage, either \$0.75, \$1, or \$2 per hour, and that workers who requested a higher wage would not be considered. The invitation described how to comply with this requirement. Still, there were some workers who applied to the jobs with wages higher than specified. I consider these workers as choosing not to apply to the job. However, the results are very similar if these workers are counted as applying.

previous jobs were also more likely to apply to these jobs. Unsurprisingly, workers were more likely to apply to higher wage jobs.

If receiving a treatment job increased workers' reservation wages, treatment group workers would be relatively more likely to apply to higher wage jobs than control group workers. The coefficients show this pattern. Treatment group workers were (an insignificant) 6 percentage points more likely to apply to the \$0.75 job than control group workers and (a significant) 20 percentage points more likely to apply to the \$2 job. However, the differences between these coefficients are not significant.²⁰

5.4 Mechanism for the Effect

The previous sections show that the experimental results are consistent with the model's predictions. However, the interpretation of these results depends on the mechanism that generates them. This section analyzes whether the results could be generated by three alternative mechanisms. The first is that completing the treatment job provides workers with human capital and this additional human capital improves workers' employment outcomes. The second is that the act of hiring a worker or leaving a detailed comment may by itself cause the market to positively update its belief about worker quality, improving employment outcomes. The third is that the treatment affects outcomes by changing worker behavior, not employer beliefs. The first two explanations seem unlikely. I find that the third may account for some, but not all, of the treatment effect.

It is relatively unlikely that working in the treatment job substantially increased workers' human capital. Workers worked for a maximum of 10 hours; the average hire worked for only 7.6 hours. Given the workers' offline experience, this is a very small increment to their total work experience. I did not provide training or guidance as it was impractical given the number of workers hired at one time. Moreover, as all hired workers completed the same task, human capital accumulation cannot explain why the detailed comment increased average earnings.

Hiring a worker in itself would lead the market to positively update its beliefs about the worker's ability if employers received different signals about worker quality. This explanation is less likely here than in a traditional labor market because, on oDesk, all employers see exactly the same resume and there are no face-to-face interactions. However, employers can interview workers via online chat or telephone. Moreover, employers may place different values on the same information. For example, some employers may know Filipino universities well, while others may not.

An empirical test of this explanation relies on the fact that the market observes immediately that a worker has begun working, but cannot see the worker's rating until nine to 11 days later. Panel A of Table 7 evaluates the treatment's effect in the week workers were working, while Panel B evaluates the treatment's effects the following week. The table shows that there is no effect of the treatment before workers were rated, but there is a large effect afterwards.²¹

²⁰I don't display the effect of the detailed comment on workers' reservation wages as measured by this treatment because the standard errors of this regression are so large as to not be informative. These results are available upon request.

²¹This alternative explanation also cannot account for the effect of the detailed treatment, since the market observes

A final alternative is that the treatments do not change workers’ expected abilities; rather, they change workers’ willingness to apply for jobs. For example, while completing the treatment job, workers might realize that oDesk jobs are more desirable than they believed and apply to more jobs as a result. It is difficult to rule out the possibility that this explanation causes part of the effect, given that the model extension in Section 3.6 also predicts that workers who received a treatment job will apply to more oDesk jobs. In fact, workers who received a treatment job did apply to more jobs than control group workers. However, this alternative hypothesis cannot explain the fact that workers who received treatment jobs were more likely to obtain the jobs they applied to than control group workers or why they had higher reservation wages. It also cannot explain the effects of receiving a detailed comment which did not change application patterns.

Panel A of Table 8 analyzes the effect of receiving a treatment job on applications. The first two columns display the results of estimating Equation 20 where the dependent variables are (1) an indicator for sending at least one application within two months and (2) the number of applications a worker sent. The average treatment group worker was 24 percentage points more likely to send at least one application and sent 24 more applications on average as a result of receiving the treatment job. Both of these effects are much greater for workers who did not have a previous job.

The third column of the panel examines the effect of the treatment on the competitiveness of the jobs to which workers applied. The data contains several objective features of jobs that correlate with the jobs’ competitiveness. Employers posting a job must indicate whether the job is an hourly job (as opposed to a fixed wage job), its job category, and whether the employer has a preference for workers with a given English ability, number of oDesk hours, level of oDesk feedback, and a maximum or a minimum wage. I also observe the number of applicants to each job.

Using these job characteristics, I predict how difficult each job is to obtain. I consider all jobs that these workers applied to in the month before the experiment. I then regress a dummy for whether the worker’s application was successful (multiplied by 100 for ease of interpreting the coefficients) on these job characteristics and worker fixed effects. The results of this regression are displayed in Appendix Table 4 and, as expected, hourly jobs, jobs with employer preferences, and jobs with more applicants are more difficult to obtain.

I then use these estimates to predict a worker’s relative probability of obtaining every job she applied to in the two months after the experiment. This relative predicted probability is the “ease index” of the job. The last column of Panel A displays the estimate of the following regression:

$$\widehat{ease}_j = \alpha + \beta job_i + \gamma_1 prevjob_i + X_i \gamma_2 + \varepsilon_{ij} \quad (24)$$

where the unit of observation is an application that worker i sent to job j . The dependent variable, \widehat{ease}_j is the index and the only control variable, X_i , included is the wave of the experiment.

that both workers who received detailed and coarse comments were hired. However, it is possible that the market inferred that the hiring firm was competent and credible from the fact that the hiring firm left a long, detailed comment. I find directly revealing, in the comment, that a worker was hired by firm that had hired hundreds of workers on oDesk (and was therefore presumably competent and credible) had no effect on workers’ outcomes.

Standard errors are clustered at the worker level. This regression shows that the average job treatment group workers applied to was substantially more difficult to obtain. This effect is sizeable: 13% of the variable’s mean.

Panel B of Table 8 evaluates the effect of receiving a detailed comment on the same three outcomes. The first two columns display the results of estimating Equation 22 where the dependent variables are, respectively, whether the worker sent any application in the ensuing two months and the number of applications she sent. The only control variable included is the wave of the experiment. The measured effects are very small, go in opposite directions, and are far from being significant. The third column displays the results of estimating the regression

$$\widehat{ease}_j = \alpha + \beta_1 detailed_i + \beta_2 rate45_i + \beta_3 job_i + \gamma_1 prevjob_i + X_i \gamma_3 + \varepsilon_{ij} \quad (25)$$

where observations are applications, not workers, and the only control variable included is the wave of the experiment. It shows that receiving a detailed comment did not affect the competitiveness of the jobs to which workers applied.

Table 9 evaluates the treatment’s effect on workers’ application success. The model extension that includes applications predicts that workers who receive the treatment group job will be more likely to obtain the jobs they apply to than control group workers, controlling for job characteristics (but not necessarily unconditionally). Panel A displays the result of the regression

$$success_{ij} = \alpha + \beta_1 job_i + \gamma_1 prevjob_i + X_{ij} \gamma_3 + \varepsilon_{ij} \quad (26)$$

where the unit of observation is worker i ’s application to job j and $success_{ij}$ is an indicator for whether this application was successful multiplied by 100 (for ease of viewing the table). The first column of the panel shows that when no worker or job controls, aside from the wave, are included, there is actually a sizable (insignificant) negative correlation between receiving a treatment job and sending a successful application.

The regression in Column 2 controls for the job’s predicted ease index because the treatment affects the competitiveness of jobs that workers apply to. It also includes employer fixed effects because this predicted ease index does not control for all aspects of a job. Once the difficulty of the job is (at least partially) controlled for, receiving a treatment job is estimated to have a sizable positive (but insignificant) effect on application success. Because receiving a treatment job affects both which workers apply to jobs and the wage they ask for, Column 3 adds controls for worker characteristics and the wage requested.²² Once these worker characteristics are included, receiving a treatment job is estimated to significantly increase the probability that a worker’s application is successful by 12%.²³

²²The worker characteristics included are dummies for the number of tests the worker has passed, the number of qualifications she has, and the number of applications she sent before the experiment, whether she took and scored above average on the most popular oDesk skills test, and the total number of jobs she had before the experiment.

²³I have controlled for selection on observables. However, if there is also selection on unobservables that is positively correlated with the selection on observables, this estimate will underestimate the true treatment effect.

Panel B estimates the effect of receiving a detailed comment on the probability that a worker obtained a given job she applied to by estimating the equation

$$success_{ij} = \alpha + \beta_1 detailed_i + \beta_2 rate45_i + \beta_3 job_i + \gamma_1 prevjob_i + X_{ij}\gamma_2 + \varepsilon_{ij}. \quad (27)$$

The controls in each column are the same as in Panel A. The estimates change across columns, but not very much, consistent with the findings that receiving a detailed comment did not significantly change application patterns. In all three columns, the effect of receiving a detailed comment on application success is sizable, but imprecisely measured. This is similar to the large, but imprecise effect of receiving a job on the number of jobs a worker obtained.

6 Market-Level Effects

6.1 Effect of the Experiment on Market Employment and Wages

The model predicts that employing many novice workers affects the market by increasing the supply of workers recognized to be high ability. It thereby increases future market employment and decreases future market wages, conditional on ability. The model predicts that market employment should increase by less than the increase in treatment group workers' employment because treatment group workers take jobs that would have been held by other workers in the absence of this experiment. These predictions can be tested utilizing the variation in the effect of the experiment on different job categories. Employers must choose one of 74 job categories when they post a job.

The experimental workers comprise a large fraction of the oDesk marketplace. In the month prior to the experiment, they sent 11% of all the applications on oDesk and obtained 8% of the jobs. They sent 34% of all applications to data entry jobs, 33% of all applications to web research jobs, and 26% of all applications to administrative support jobs. On the other hand, they sent less than 1% of applications to jobs in 20 other job categories, primarily ones that require specific computer skills such as web programming or game development.

First, I create a variable measuring the effective number of treatment jobs created in each job category. Since workers do not choose one of these job categories, I allocate the treatment jobs to these categories using the share of their applications that treatment group workers sent to jobs in these categories in the month prior to the experiment. That is

$$jobscreated_j = \frac{\text{treatment group applications to category } j \text{ jobs}}{\text{total treatment group applications}} \times 952 \quad (28)$$

where 952 is the total number of treatment group jobs created.

I estimate the effect of having an additional experienced, treatment group worker in a given job category on overall employment in that category by estimating the equation

$$y_{jt} = \alpha + \beta_1(jobscreated_j \times after_t) + \delta_t + \delta_j \times t + \varepsilon_{jt} \quad (29)$$

where the unit of observation is job category j in week t . The dependent variable is the number of jobs created in week t in job category j and δ_t are week fixed effects. The coefficient of interest, β_1 , measures the number of additional jobs created in a job category with an additional treatment group worker in a week after the experiment. Job category-specific time trends $\delta_j \times t$ are included to control for the fact that even in the absence of the experiment, these categories might have grown at different rates. The regression includes all sixteen weeks in 2010 until eight weeks after the experiment, excluding the weeks of the experiment itself. Observations are weighted by their size in a pre-period.²⁴ Standard errors are clustered by job category.

The first column of Table 10 shows the results of this regression. It shows that each additional treatment group worker in a job category led to an additional 0.051 jobs created per week. Treatment group workers obtained an extra 0.062 jobs per week as a result of the treatment.²⁵ The market-level increase in employment (0.051 jobs per week) is smaller than the worker-level increase (0.062 jobs per week) as the model predicts, but not substantially or significantly so. This suggests that treatment group workers crowded out relatively little employment of other workers.

I estimate the effect of the experiment on market wages using the same identification strategy. The second and third columns of Table 10 display the results of the regression

$$wage_{ijt} = \alpha + \beta_1(jobscreated_j \times after_t) + X_i\beta_3 + \delta_t + \delta_j \times t + \varepsilon_{ijt}. \quad (30)$$

Observations in this regression are hourly jobs and the dependent variable is the hourly wage in the job. Because the model's predictions are conditional on worker ability, it also includes controls for worker ability.²⁶ Column 2 shows that having an additional treatment group worker in a given job category decreases average hourly wages by a significant \$0.0015 on average. Including worker characteristics does not affect these estimates.

6.2 Market-Level Welfare Analysis

The experiment's effect on market surplus can be estimated using its effects on market employment and wages. Figure 4 shows the effect of the experiment on the oDesk market. The inelastic supply curve (S) represents the number of workers recognized to be high ability. The experiment shifts this supply curve outward, increasing employment and decreasing wages. Total market surplus increases by the dark gray shaded area. The dark gray shaded triangle is the increase in firm profits from the increased employment and the dark gray shaded rectangle is the increase in worker surplus from the increased employment. The light gray shaded rectangle is a transfer from workers to firms.

Assuming linearity of the demand curve, the area of the triangle is one half the change in

²⁴Because one job category had no new jobs in the pre-period, it is excluded from the regression.

²⁵This is calculated by dividing the treatment effect in Table 3 by the number of weeks.

²⁶These are indicators for the number of tests the worker took, the number of qualifications the worker listed, and the number of portfolio items the worker had before the experiment.

(future) employment generated by the experiment times one half the change in wages. This is:

$$\frac{1}{2} \times \underbrace{\left(952 \text{ jobs} \times 8 \text{ weeks} \times 1.31 \frac{\text{hours}}{\text{job} \times \text{week}} \right)}_{\text{Increase in Employment}} \times \underbrace{\left(952 \text{ jobs} \times \frac{\$0.0015}{\text{job} \times \text{hour}} \right)}_{\text{Decrease in Wages}} \approx \$6,900. \quad (31)$$

This calculation measures the increase in employment using a change in hours to correspond to the hourly wages. Estimating Equation 29 where the dependent variable is total hours worked in a job category-week shows that each treatment job created led to 1.31 additional hours worked in hourly jobs in each of the eight weeks after the experiment. Multiplying this figure by the number of weeks in the calculation (8) and the number of treatment jobs created (952) gives the overall increase in market employment resulting from the experiment. Similarly, multiplying the market-level change in wages per experimental job (\$0.0015) by the number of experimental jobs (952) yields the overall decrease in market wages. This calculation shows that the increase in firm profits as a result of this experiment, not including the transfer from workers, is \$6,900.

The area of the dark gray shaded rectangle, the increase in worker surplus due to increased market employment, is the increase in market employment multiplied by the prevailing wage level:

$$\underbrace{\left(952 \text{ jobs} \times 8 \text{ weeks} \times 1.31 \frac{\text{hours}}{\text{job} \times \text{week}} \right)}_{\text{Increase in Employment}} \times \underbrace{\$2.09}_{\text{Wage Level}} \approx \$20,800. \quad (32)$$

Here, \$2.09 is the average hourly wage experimental workers received after the experiment. However, the above calculation gives the increase in worker surplus only if workers face zero cost of effort. Panel A of Table 11 estimates the increase in total market surplus under different assumptions about the cost of worker effort.

These calculations may underestimate the increase in surplus on oDesk because they only consider gains for eight weeks after the experiment. However, they do not include any losses in other markets due to the increased employment on oDesk. If firms decrease their offline hiring or workers decrease their offline work when oDesk employment increases, the experiment's effect on oDesk social surplus exceeds its effect on total social surplus.

The social cost of this experiment includes the fixed cost of hiring workers and the cost of worker effort. The money paid to the workers in the treatment jobs is not a social cost; it is a transfer. If this transfer is funded by the government, there is a deadweight loss of taxation.²⁷ In a traditional job, these costs would be offset by the value of the output produced. However, in this experiment, there was no direct value to the output produced. Because workers were not expected to produce usable output, the fixed cost of employing them was relatively low.

Panel B of Table 11 calculates the social cost of this experiment under different assumptions about the size of these costs. The most difficult piece to estimate is the cost of worker effort. Given

²⁷Part of the funding for the experiment came from government grants, so there is some deadweight loss of taxation in this context. More generally, the deadweight loss of taxation is included to show the cost of this program if it were not run as an experiment, but was subsidized by the government.

that 49% of the treatment group was willing to apply to a job offering a wage of \$1 in the reservation wage treatment, I use this as the “best-guess” estimate of workers’ cost of effort. However, I also estimate the social cost of the experiment using both higher and lower costs of effort.

This intervention was not designed optimally, nor intended, to increase net market surplus. Nonetheless, unless workers’ cost of effort is estimated to be their full wage, the benefits of this intervention exceed its costs. This suggests that novice employment on oDesk was inefficiently low before the experiment.

7 Conclusion

The market for entry-level workers is characterized by high unemployment and strong competition for jobs. This paper proposes that the high unemployment partially results from the fact that firms do not receive the full benefit of discovering talent. Hiring an inexperienced worker requires an investment from firms, but reveals information about worker ability, generating an option to hire a known, good worker in the future. To the extent that this information is public, workers receive the option value. If workers cannot compensate firms for their investments, too few inexperienced workers will be employed.

This paper formalizes this intuition into a model of the labor market and tests this model through a large field experiment. Consistent with the model’s predictions, giving workers the opportunity to demonstrate their abilities through short jobs increased their future employment, earnings, and reservation wages. By increasing the supply of workers recognized to be high-ability, it increased market employment, decreased market wages, and increased total market surplus. Despite the fact that this intervention was untargeted, under plausible assumptions, it increased market surplus by more than its social cost. This suggests that, before the experiment, firms hired inefficiently few novices and discovered inefficiently little novice talent.

Directly giving the market more information about workers’ job performance also affected workers’ employment outcomes in the way the model predicts. Providing the market with additional positive information increased workers’ future earnings while providing it with additional negative information decreased their earnings. Overall, giving the market more information about workers’ job performance increased their future earnings. The model predicts that, by allowing the market to more accurately identify high-ability workers, this additional information increases the option value of being hired and workers’ future earnings.

Although these results come from a particular online marketplace, there are two settings to which these insights might generalize: traditional (offline) entry-level labor markets and international trade of goods and labor. Compared to traditional entry-level employers, oDesk employers have more uncertainty about the abilities of job applicants. They do not get to meet applicants, may not be familiar with workers’ credentials from foreign schooling systems or employers, and have very limited ability to verify these credentials. They may also be hiring workers for more technical tasks where ability varies more among applicants. This suggests that the inefficiency may

be larger in the oDesk marketplace than in a traditional labor market. On the other hand, the fixed costs of hiring an oDesk worker are particularly low. oDesk workers can be hired with the click of a mouse. The tasks they complete are typically well-defined, easy to explain, and require no training. If workers do not perform well, they can be fired instantaneously with no penalty. This suggests that the inefficiency may be larger in a traditional labor market.

Under the supposition that these insights apply to traditional entry-level labor markets, theory suggests several policies that would reduce the inefficiency. Subsidizing entry-level hiring compensates employers for discovering inexperienced workers' talent, inducing them to hire more novices. The market has already developed institutions that compensate employers for hiring novices. Internships allow firms to pay novices low wages while, in Europe, unions and industry consortiums directly pay some workers' initial salaries. Alternatively, workers could be hired for short-term employment by a separate, credentialing employer who would relay their job performance to the market. For this employer to improve market efficiency, its assessments of job performance would need to be credible. Workers would be willing to pay the employer directly to take this employment "test," but they may be credit-constrained, in which case public funds could subsidize the employer. A traditional test would not be sufficient to correct this inefficiency given the difficulty of determining workers' motivation and reliability through means other than employment.

As most oDesk jobs are offshored from U.S. employers to foreign workers, the experimental results in this paper may shed light on whether developing a reputation is a significant barrier to offshoring, and on a grander scale, trade between foreign and domestic firms. Unlike in other forms of offshoring and international trade, the only significant barrier to transacting on oDesk is the difficulty of building a reputation. Firms offshoring offline may face significant costs of identifying available labor. Similarly, foreign and domestic firms wanting to trade must invest in identifying and communicating with each other as well as, potentially, new plants and capital. On the other hand, oDesk workers and firms can join the marketplace and search for each other costlessly and quickly. This experiment shows that the cost of building a reputation, alone, is sufficient to reduce the volume of trade, but, when reputations are established, trade volume increases. The extent to which the results of this experiment can be applied to more general trade contexts is an important question for future research.

References

- ACEMOGLU, D., AND J.-S. PISCHKE (1998): "Why Do Firms Train? Theory and Evidence," *Quarterly Journal of Economics*, 113(1), 79–118.
- ACEMOGLU, D., AND J.-S. PISCHKE (1999): "Beyond Becker: Training in Imperfect Labour Markets," *The Economic Journal*, 109(453), pp. F112–F142.
- ALTONJI, J., AND C. PIERRET (2001): "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics*, 116(1), 313–350.

- AUTOR, D. H. (2001): “Why Do Temporary Help Firms Provide Free General Skills Training?,” *Quarterly Journal of Economics*, 116(4), 1409–1448.
- BANERJEE, A. V., AND E. DUFLO (2000): “Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry,” *Quarterly Journal of Economics*, 115(3), 989–1017.
- BECKER, G. (1964): “Human Capital,” *Columbia University for the National Bureau of Economic Research*.
- FARBER, H., AND R. GIBBONS (1996): “Learning and Wage Dynamics,” *The Quarterly Journal of Economics*, pp. 1007–1047.
- KREPS, D. M., AND R. WILSON (1982): “Reputation and Imperfect Information,” *Journal of Economic Theory*, 27(2), 253 – 279.
- LOEWENSTEIN, M. A., AND J. R. SPLETZER (1998): “Dividing the Costs and Returns to General Training,” *Journal of Labor Economics*, 16(1), pp. 142–171.
- MCMILLAN, J., AND C. WOODRUFF (1999): “Interfirm Relationships and Informal Credit in Vietnam,” *Quarterly Journal of Economics*, 114(4), 1285–1320.
- NEWMAN, K. (1996): “Job Availability: Achilles Heel of Welfare Reform,” *National Forum*, 76(3), 20–24.
- RESNICK, P., R. ZECKHAUSER, J. SWANSON, AND K. LOCKWOOD (2006): “The Value of Reputation on eBay: A Controlled Experiment,” *Experimental Economics*, 9, 79–101.
- TERVIO, M. (2009): “Superstars and Mediocrities: Market Failure in the Discovery of Talent,” *Review of Economic Studies*, 76(2), 829–850.
- TIROLE, J. (1996): “A Theory of Collective Reputations (with Applications to the Persistence of Corruption and to Firm Quality),” *The Review of Economic Studies*, 63(1), pp. 1–22.

8 Appendix A

This section proves Proposition 3. Only the most efficient firms will hire a worker, so there is a threshold, \bar{c} , such that all firms with $c_j \leq \bar{c}$ hire a worker and all firms with $c_j > \bar{c}$ do not. Firms must be indifferent to hiring any employed worker and the marginal firm must earn zero profit. Thus, the wages of employed workers equal

$$w_{ij} = \hat{\eta}_i - \bar{c}. \quad (33)$$

If all novice workers with $\hat{a}_{i0} \geq \bar{a}_A$ and only these novices are hired, then

$$\hat{\eta}_{i0} = (\gamma_A + \gamma_B)\hat{a}_{i0}. \quad (34)$$

Veteran workers who have worked have known types. Thus, experienced veteran group C workers will not receive wage offers, but experienced veteran group A and B workers will work if their market wages exceed their outside options:

$$w_{ij} = \hat{a}_{i1} - \bar{c} > w_0. \quad (35)$$

If only novices with $\hat{a}_{i0} \geq \bar{a}_A$ and veteran group A and B workers with $\hat{a}_{i1} - \bar{c} > w_0$ work in the market, then, by definition, $\bar{c} = c(\bar{a}_A)$.

Novice group C workers with $\hat{a}_{i0} \geq \bar{a}_A$ strictly prefer any positive wage offer to their outside options. Novice group A and B workers will work in the market when

$$\underbrace{((\gamma_A + \gamma_B)\hat{a}_{i0} - \bar{c} - w_0)}_{\text{Net Wage in Novice Period}} + \underbrace{\Pr(\hat{a}_{i1} - \bar{c} \geq w_0) \times E[\hat{a}_{i1} - \bar{c} - w_0 | \hat{a}_{i1} - \bar{c} \geq w_0]}_{\text{Veteran-Period Option Value}} \geq 0. \quad (36)$$

By the definition of \bar{a}_A and the fact that both terms of this expression are increasing in \hat{a}_{i0} , group A workers will work exactly when $\hat{a}_{i0} \geq \bar{a}_A$. Since $w_B < w_A$, Group B workers with $\hat{a}_{i0} \geq \bar{a}_A$ strictly prefer working in the market to their outside option.

By the assumption in Equation 16 of the text, only these workers: novices with $\hat{a}_{i0} \geq \bar{a}_A$ and veteran group A and B workers with $\hat{a}_{i1} - \bar{c} > w_0$ work in the market. Firms will not hire novice workers with $\hat{a}_{i0} < \bar{a}_A$. Group A workers with $\hat{a}_{i0} < \bar{a}_A$ do not work in the market, so firms could only hire group B and C workers at these ability levels. But, these workers' net marginal products are negative. Firms will also not hire veterans who have not worked in their novice periods. These workers all have

$$\hat{a}_{i1} = \hat{a}_{i0} < \bar{a}_A. \quad (37)$$

Novice workers benefit more from market work than do veterans because novices receive the option value from talent discovery. Group A workers with these expected abilities prefer their outside option to market work as novices, so they also prefer their outside options to market work as veterans. Firms will not hire only group B and C workers with these expected ability levels.

This equilibrium does not maximize total surplus. Consider a social planner who has the same information as employers; that is, he cannot distinguish group A , B , and C novices. He must, therefore, choose which novices to hire solely on the basis of their expected abilities. He will set an ability threshold $\hat{a}_{i0} = a^*$ such that all novices with $\hat{a}_{i0} \geq a^*$ and only these will be employed. Denote the fixed cost of the marginal firm who hires a worker in the social planner's solution by c^* . For notational simplicity, denote the marginal worker's expected net product at this marginal firm by w^* :

$$w^* = (\gamma_A + \gamma_B)a^* - c^*. \quad (38)$$

The total social surplus from hiring a worker with $\hat{a}_{i0} = a^*$ is a weighted average of the surplus from hiring group A , B , and C workers:

$$\begin{aligned} & \gamma_A[w^* - w_A + \Pr(\hat{a}_{i1} - c^* \geq w_A) \times E[\hat{a}_{i1} - c^* - w_A | a^* - c^* \geq w_A]] \\ & + \gamma_B[w^* - w_B + \Pr(\hat{a}_{i1} - c^* \geq w_B) \times E[\hat{a}_{i1} - c^* - w_B | a^* - c^* \geq w_B]] \\ & + \gamma_C[w^*] = 0. \end{aligned}$$

If the social planner used the same hiring threshold as the market, $a^* = \bar{a}_A$, then the same number of firms would hire a worker in the social planner's solution as the market equilibrium $c^* = \bar{c}$. Under these parameters, group A workers are indifferent to taking this wage offer, so the value of the first line is zero. However, group B and C workers prefer this wage offer to their outside option, so the values of the second and third lines are positive. Thus, the left-hand side of this equation is positive. For this equation to hold, the social planner's hiring threshold must be below the market's: $a^* < \bar{a}_A$. The market hires too few novice workers.

9 Appendix B

This section proves Proposition 4. No worker with $w_0 = 0$ will exit the market. If she exits the market, she receives 0, while if she remains in the market, her expected gain is at least $\varepsilon\delta$. Veteran workers with outside option w_1 will exit the market if

$$\hat{a}_{i1} < w_1 + \bar{c}. \quad (39)$$

They will earn

$$w_{ij} = \max(\hat{a}_{i1} - \bar{c}, 0). \quad (40)$$

if they remain in the market and will exit if their outside option is greater than this wage.

The first treatment, job receipt, only affects novice workers with expected abilities

$$\hat{a}_{i0} < \bar{c} \quad (41)$$

because control group novices with higher expected abilities are employed by other firms. All

control group workers with $\hat{a}_{i0} < \bar{c}$ have

$$\hat{a}_{i1} = \hat{a}_{i0} < \bar{c} < w_1 + \bar{c}. \quad (42)$$

They will all exit the market if offered outside option w_1 . On the other hand, some treatment group workers with $\hat{a}_{i0} < \bar{c}$ will perform well enough in the treatment job that their veteran-period expected ability exceeds the market-exit threshold, $w_1 + \bar{c}$. They will not exit the market even if offered outside option w_1 .

The second experimental treatment, providing workers with a more detailed evaluation, has an indeterminate effect on the fraction of workers whose expected ability exceeds the market-exit threshold. Some workers' veteran expected abilities would exceed the market-exit threshold if they received only a coarse comment. Receipt of a detailed evaluation decreases some of these workers' expected abilities below the market-exit threshold. Other workers' veteran expected abilities would be below the market-exit threshold if they only received a coarse comment. Receipt of a detailed evaluation increases some of these workers' expected abilities above the market-exit threshold. The distribution of workers' expected abilities determines the relative sizes of these two groups and the net effect of the treatment on market exit.

Figure 1. Example oDesk Profile



Angelo Penamayor - "PROFESSIONAL DATA ENTRY SPECIALIST WITH CUSTOMER SERVICE EXPERIENCE - Freelance Data Entry Professional, Philippines"

Permalink : <http://www.odesk.com/users/~~>

\$2.22/hr

Contact

Overview

Résumé

Work History & Feedback (4)

Tests (9)

Portfolio (0)

Over the last 3 years, I have also worked as a customer service representative for eBay USA and JP Morgan Chase Bank. I have provided services and information through phones, correspondence, email and chat. My core competency lies in providing high standard customer service along with my excellent communication and technical skills. Since I am also sending emails and correspondence to my customers, I have mastered my skills in all kinds of Microsoft Office applications. Now, I am seeking...

[more](#)

Recent Work History & Feedback

[See All Work History & Feedback \(4 items, with Feedback\)](#)

Employer ID	From/To	Job Title	Paid	Feedback
192528	08/2010 - Present	Simple data collection apply below 50c	\$20 (46 hrs @ \$0.44/hr)	<i>Job in progress</i>
216906	07/2010 - 08/2010	Profile Creator	\$10 (fixed-price)	★★★★★ 5.0 Great work. I'll definitely hire Angelo again. :) Contractor-to-Employer Feedback: ★★★★★ 5.0
136432	07/2010 - 08/2010	Data entry assistants needed no experience needed but great prospects	\$41 (73 hrs @ \$0.56/hr)	★★★★★ 5.0 Contractor-to-Employer Feedback: ★★★★★ 5.0
135179	08/2010 - 08/2010	Data Entry Specialist	\$10 (fixed-price)	★★★★★ 5.0 Contractor-to-Employer Feedback: ★★★★★ 5.0

Last 6 mos.

All-time

Feedback: ★★★★★ (5.00) 3 feedbacks

Hours: 119

Contracts: 4

[See all Work History & Feedback](#)

Location: Caloocan, Philippines (UTC+08:00)

English Skills: (self-assessed) ██████████ 5.0

Member Since: June 1, 2010

Last Worked: September 22, 2010

oDesk Ready: [?](#) Yes

Associated with:

[Topnotch Group > Ja G](#)

Hours: 0

★★★★★ [1 feedbacks](#)



Related links:

- Trends for [Data Entry Professionals](#)
- Trends for [Excel Consultants](#)

Figure 2. Expected Novice Ability

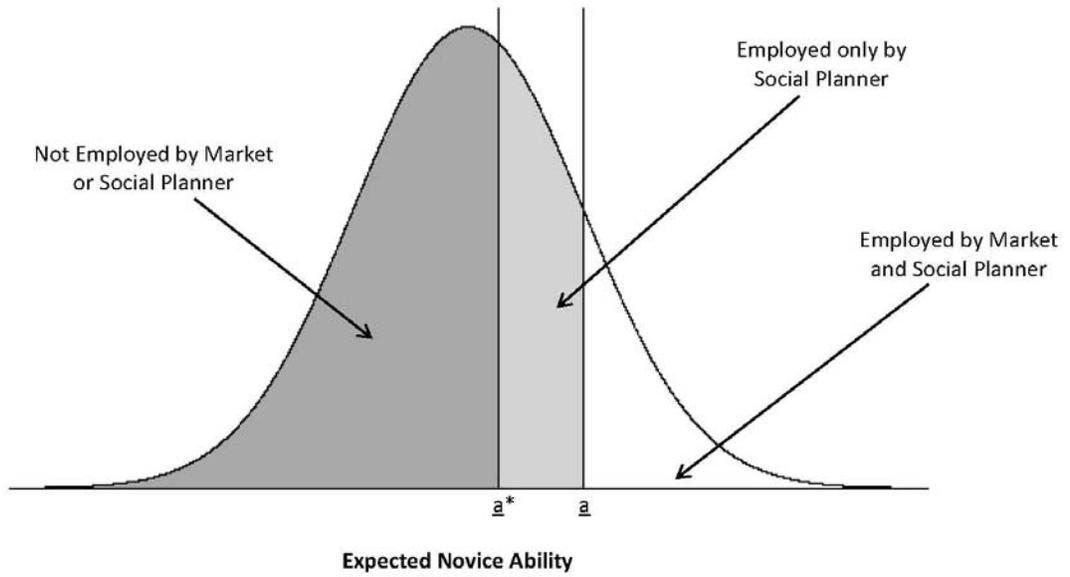


Figure 3. Experimental Design

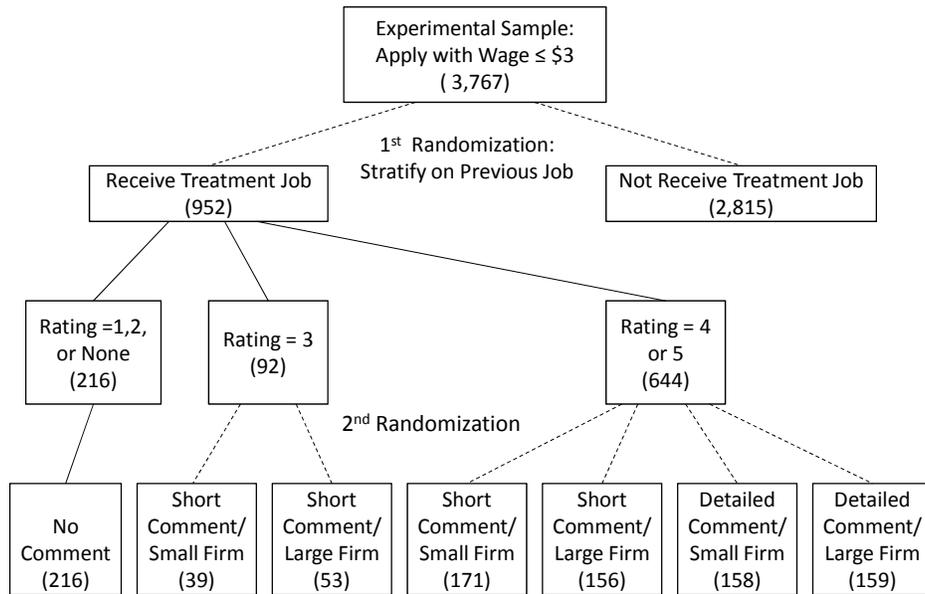


Figure 4. Experiment's Market-Level Effects

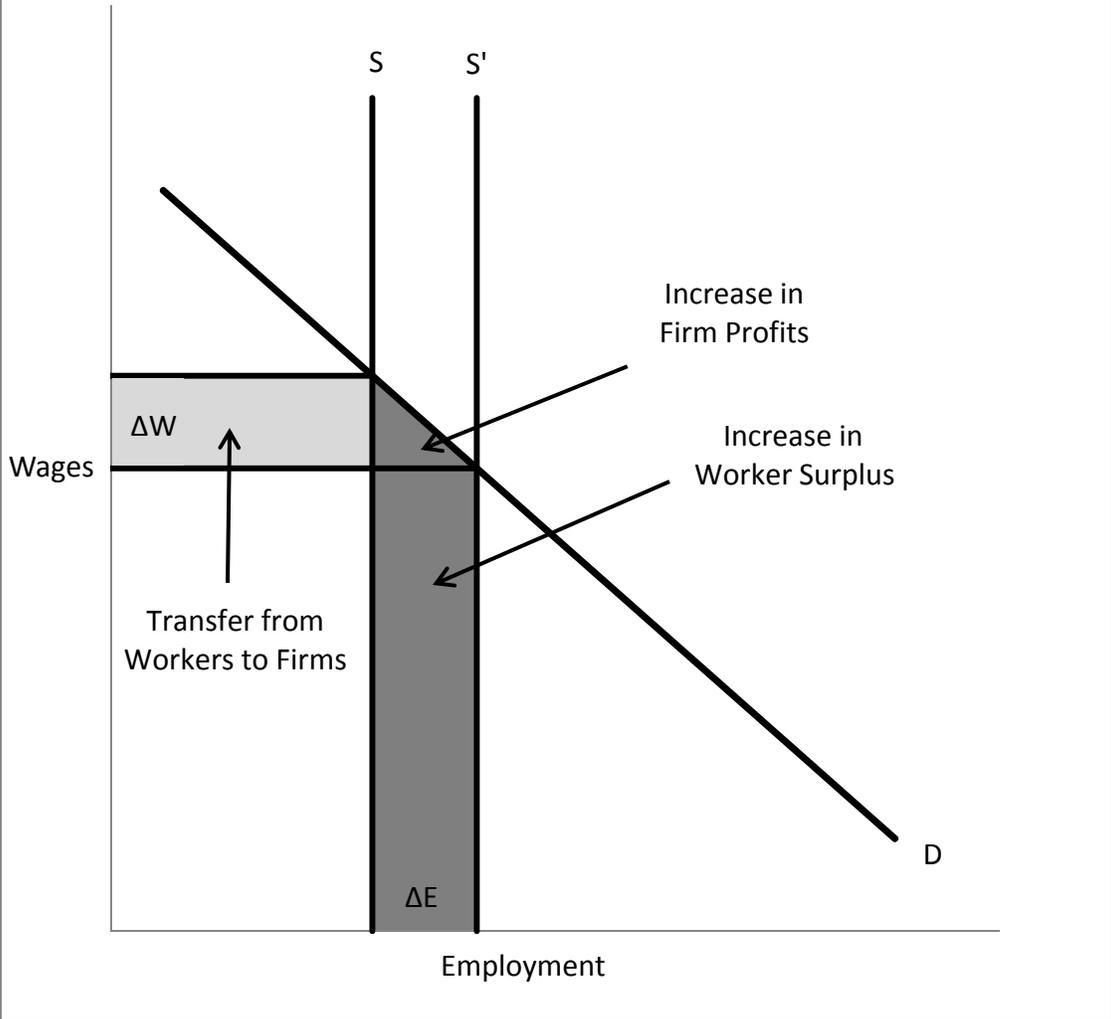


Table 1. Sample Selection

	No Previous Job	Any Previous Job	Total	Percentage
Contacted	7,136	2,826	9,962	100%
Applied	2,324	1,528	3,852	39%
Applied with Wage \leq \$3	2,298	1,469	3,767	38%
Received Treatment Job	736	216	952	25% of experimental sample

Notes: The first row enumerates the workers invited to apply to the job while the second counts those that accepted that invitation. The third row is the experimental sample: workers who applied to the job requesting a wage less than or equal to \$3. The final row counts the workers who were randomly selected to receive a treatment job. Unless otherwise indicated, percentages refer to the percentage of contacted workers.

Table 2. The Effect of the Treatment Job on Employment Outcomes
During the Two Months After the Experiment

	A. Workers with No Previous Job			B. Workers with Previous Jobs		
	Treatment (1)	Control (2)	Difference (3)	Treatment (4)	Control (5)	Difference (6)
Total Jobs	0.807	0.284	0.523**	2.463	1.958	0.505**
Standard Error of Difference			(0.073)			(0.244)
p-value: Equality of Differences						[0.947]
Any Job	0.299	0.117	0.182**	0.611	0.545	0.066*
Standard Error of Difference			(0.017)			(0.037)
p-value: Equality of Differences						[0.004]
Hours Worked	12.40	5.36	7.05**	60.99	47.80	13.19*
Standard Error of Difference			(1.59)			(7.12)
p-value: Equality of Differences						[0.45]
Posted Wage	2.28	2.03	0.25**	2.50	2.38	0.12
Standard Error of Difference			(0.05)			(0.08)
p-value: Equality of Differences						[0.25]
Earnings	28.43	10.06	18.37**	144.02	120.60	23.41
Standard Error of Difference			(3.57)			(18.14)
p-value: Equality of Differences						[0.80]
Observations	736	1,562	2,298	216	1,253	1,469

Notes: Cells in Columns 1, 2, 4, and 5 display the mean value of the employment outcome indicated by the row heading for workers indicated by the column heading within two months of the experiment. All experimental jobs and earnings are excluded. Columns 3 and 6 provide the difference in mean outcomes for treatment and control workers without and with previous jobs, respectively. The standard error of this difference is in parentheses. One asterisk indicates the difference is significant at the 10% level and two asterisks indicate the difference is significant at the 5% level. Column 6 presents, in brackets, the p-value from a test that the treatment effects for workers with and without previous jobs are equal.

Table 3. The Effect of the Treatment Job on Employment Outcomes
During the Two Months After the Experiment

	Total Jobs (1)	Any Job (2)	Hours Worked (3)	Posted Wage (4)	Earnings (5)
<u>A. No Controls</u>					
Treatment Job	0.498** (0.092)	0.150** (0.017)	9.62** (2.59)	0.21** (0.06)	20.36** (6.29)
Previous Job	1.674** (0.094)	0.406** (0.015)	43.50** (2.68)	0.33** (0.04)	111.42** (6.69)
<u>B. Controls for Wage Offered, Date Joined oDesk, and Country</u>					
Treatment Job	0.477** (0.093)	0.148** (0.017)	9.15** (2.58)	0.21** (0.05)	19.22** (6.26)
Previous Job	1.778** (0.103)	0.416** (0.016)	46.87** (2.96)	0.30** (0.04)	117.96** (7.48)
<u>C. Additional Controls for Test Scores</u>					
Treatment Job	0.483** (0.092)	0.147** (0.017)	9.14** (2.59)	0.21** (0.05)	19.47** (6.27)
Previous Job	1.658** (0.102)	0.388** (0.017)	44.89** (3.02)	0.28** (0.04)	110.71** (7.45)
Observations	3,767	3,767	3,767	3,767	3,767

Notes: This table displays the results of estimating Equation 20. The regressions in Panel A contain no controls except for an indicator for being in the second experimental wave. The regressions in Panel B additionally include controls for the wage requested when applying for the treatment job, the date the worker joined oDesk and country fixed effects. (See footnote 15 for more information.) The regressions in Panel C add indicators for whether workers took each of the six most popular skills tests among low-wage data-entry workers and scored above average on these tests. All experimental jobs are excluded. Huber-White standard errors are in parentheses. Two asterisks indicate the coefficient is significant at the 5% level.

Table 4. Effect of Receiving a Detailed Comment: Meeting Deadlines
During the Two Months After the Experiment

	Total Jobs (1)	Any Job (2)	Hours Worked (3)	Posted Wage (4)	Earnings (5)
Met Deadline × Detailed Comment	0.442 (0.446)	0.138 (0.334)	31.57** (15.19)	0.52 (0.61)	59.27** (21.08)
Met Deadline	0.791** (0.311)	-0.013 (0.190)	-0.47 (14.53)	0.07 (0.37)	19.90 (17.61)
Detailed Comment	-0.238 (0.388)	-0.095 (0.331)	-27.40** (13.98)	-0.28 (0.60)	-34.88** (15.51)
Sum of Detailed Coefficients	0.204	0.043	4.17	0.24	24.39
p-value of F-test	0.350	0.277	0.48	0.07	0.09
Observations	644	644	644	644	644

Notes: This table displays the results of estimating Equation 21 where x_i is an indicator for meeting the deadline. It includes only workers who received ratings of at least four and were thus eligible to receive a detailed comment. All experimental jobs are excluded. Huber-White standard errors are in parentheses. One asterisk indicates the coefficient is significant at the 10% level and two asterisks indicate the coefficient is significant at the 5% level.

Table 5. The Effect of Receiving a Detailed Comment on Employment Outcomes
During the Two Months After the Experiment

	Total Jobs (1)	Any Job (2)	Total Hours (3)	Posted Wage (4)	Earnings (5)
<u>A. No Controls</u>					
Detailed Comment	0.209 (0.208)	0.041 (0.038)	3.73 (5.58)	0.24* (0.13)	23.85* (13.43)
Rating = 4 or 5	0.612** (0.176)	0.159** (0.036)	11.88 (4.58)	0.26** (0.08)	16.46 (10.87)
<u>B. Controls for Wage Offered, Date Joined oDesk, and Country</u>					
Detailed Comment	0.217 (0.212)	0.041 (0.039)	4.32 (5.55)	0.22* (0.13)	25.15* (13.31)
Rating = 4 or 5	0.606** (0.181)	0.161** (0.036)	10.52** (4.59)	0.23** (0.06)	11.42 (11.03)
<u>C. Additional Controls for Test Scores</u>					
Detailed Comment	0.182 (0.210)	0.035 (0.038)	3.53 (5.55)	0.22* (0.13)	22.92* (13.25)
Rating = 4 or 5	0.612** (0.181)	0.161** (0.036)	10.89** (4.61)	0.22** (0.06)	12.16 (10.96)
Observations	3,767	3,767	3,767	3,767	3,767

Notes: This table displays the results of estimating Equation 22. The regressions in Panel A contain no controls except for an indicator for being in the second experimental wave. The regressions in Panel B additionally include controls for the wage requested when applying for the treatment job, the date the worker joined oDesk and country fixed effects. (See footnote 15 for more information.) The regressions in Panel C add indicators for whether workers took each of the six most popular skills tests among low-wage data-entry workers and scored above average on these tests. All experimental jobs are excluded. Huber-White standard errors are in parentheses. One asterisk indicates the coefficient is significant at the 10% level and two asterisks indicate the coefficient is significant at the 5% level.

Table 6. Effect of the Treatment Job on Reservation Wages

Dependent Variable: Indicator for Applying

	(1)	(2)
Treatment Job	0.141** (0.041)	
Treatment Job × \$0.75 Wage		0.059 (0.071)
Treatment Job × \$1 Wage		0.162** (0.066)
Treatment Job × \$2 Wage		0.195** (0.079)
\$1 Wage Job	0.141** (0.044)	0.099 (0.064)
\$2 Wage Job	0.292** (0.049)	0.212** (0.073)
Previous Job	0.075* (0.041)	
Interactions of Prev. Job with Wage	No	Yes
p-value: Equality of Treatment Coefficients		0.386
Mean of Dependent Variable	0.395	0.395
Observations	630	630

Notes: This table displays the results of estimating Equation 23. Only workers who were invited to apply to this job are included. Huber-White standard errors are in parentheses. One asterisk indicates the coefficient is significant at the 10% level and two asterisks indicate the coefficient is significant at the 5% level.

Table 7. Effect of the Treatment Job on Employment Before and After Rating

	A. Week Between Hiring and Rating		B. Week After Rating	
	Total Jobs (1)	Any Job (2)	Total Jobs (3)	Any Job (4)
Treatment Job	-0.010 (0.014)	0.000 (0.010)	0.086** (0.021)	0.057** (0.012)
Previous Job	0.252** (0.018)	0.170** (0.011)	0.247** (0.020)	0.167** (0.011)
Second Wave	0.008 (0.016)	0.001 (0.010)	-0.023 (0.017)	-0.014 (0.010)
Constant	0.037 (0.010)	0.033 (0.007)	0.046 (0.012)	0.034 (0.007)
Observations	3,767	3,767	3,767	3,767

Notes: Both panels present the results of estimating Equation 20. Panel A only includes jobs obtained in the week immediately following the worker's hire. Panel B only includes jobs obtained in the week immediately following workers' evaluations. Huber-White standard errors are in parentheses. Two asterisks indicate the coefficient is significant at the 5% level.

Table 8. Effect of the Treatments on Application Patterns
During the Two Months After the Experiment

	A. Effect of Treatment Job			B. Effect of a Detailed Comment		
	Sent Any Application	Applications Sent	Ease Index	Sent Any Application	Applications Sent	Ease Index
	(1)	(2)	(3)	(4)	(5)	(6)
Detailed Comment				-0.008 (0.019)	0.53 (5.33)	0.052 (0.132)
Rating = 4 or 5				0.266** (0.028)	29.17** (4.45)	-0.558** (0.171)
Treatment Job	0.238** (0.015)	23.88** (2.18)	-0.421** (0.085)	0.060** (0.027)	3.85 (2.38)	0.015 (0.145)
Previous Job	0.337** (0.013)	30.38** (1.89)	0.599** (0.080)	0.332** (0.013)	29.91** (1.89)	0.582** (0.081)
Mean of Dep. Var.	0.745	30.28	2.973	0.729	25.28	2.973
Observations	3,767	3,767	114,082	3,767	3,767	114,082

Notes: The first two columns of each panel present the results of estimating Equations 20 (Panel A) and 22 (Panel B) where the unit of observation is the worker and the dependent variables are an indicator for whether the worker sent any application and the number of applications sent in the two months following the experiment. Huber-White standard errors are in parentheses. Columns 3 and 6 present the results of estimating Equations 24 and 25 where the unit of observation is an application. In these regressions, standard errors are clustered by worker. All experimental jobs are excluded. Two asterisks indicate the coefficient is significant at the 5% level.

Table 9. Effect of the Treatments on Application Success
 Dependent Variable: Indicator that an Application is Successful $\times 100$

	A. Effect of the Treatment Job			B. Effect of a Detailed Comment		
	(1)	(2)	(3)	(4)	(5)	(6)
Detailed Comment				0.381 (0.331)	0.349 (0.269)	0.257 (0.257)
Rating = 4 or 5				-0.062 (0.499)	0.524 (0.331)	0.582* (0.302)
Treatment Job	-0.379 (0.244)	0.234 (0.198)	0.416** (0.188)	-0.482 (0.450)	-0.332 (0.305)	-0.158 (0.274)
Previous Job	2.035** (0.238)	0.752** (0.192)	0.731** (0.210)	2.045** (0.239)	0.777** (0.191)	0.727** (0.210)
Ease Index	No	Yes	Yes	No	Yes	Yes
Employer Fixed Effects	No	Yes	Yes	No	Yes	Yes
Worker Characteristics	No	No	Yes	No	No	Yes
Mean of Dep. Var.	3.59	3.59	3.59	3.59	3.59	3.59
Observations	114,082	114,082	114,082	114,082	114,082	114,082

Notes: Panels A and B display the results of estimating Equations 26 and 27, respectively. The regressions in Columns 1 and 4 include no controls. The regressions in Columns 2 and 5 control for the ease index and employer fixed effects. The regressions in Columns 3 and 6 add worker characteristics measured before the experiment: dummies for the number of tests the worker passed, the number of qualifications she had, whether she took and scored above average on the most popular oDesk skills test, the total number of jobs she had before the experiment, and the wage she requested when applying for the treatment job. All experimental jobs are excluded. Standard errors are clustered by worker. One asterisk indicates the coefficient is significant at the 10% level and two asterisks indicate the coefficient is significant at the 5% level.

Table 10. Market-Level Effects of the Experiment

	A. Jobs Created	B. Hourly Wages	
	(1)	(2)	(3)
Jobs Created x After	0.051** (0.023)	-0.145** (0.046)	-0.139** (0.049)
Worker-Level Effect	0.062		
Week Fixed Effects	Yes	Yes	Yes
Category-Specific Time Trends	Yes	Yes	Yes
Worker Characteristics	No	No	Yes
Observations	1,825	125,456	125,456

Notes: This table displays the results of estimating Equations 29 (Column 1) and 30 (Columns 2 and 3). In Column 1, observations are weighted by the number of jobs created in the category in a pre-period. These regressions contain outcomes from 16 weeks before the experiment to 8 weeks afterwards, excluding the weeks of the experiment. All standard errors are clustered by job category. Two asterisks indicate the coefficient is significant at the 5% level. See the text for the creation of the "Worker-Level Effect."

Table 11. Estimated Effect of the Experiment on Market-Level Efficiency

	High Benefit, Low Cost	Medium Benefit, Medium Cost	Low Benefit, High Cost
<u>A. Increased Market Surplus</u>			
Increased Firm Profit	\$6,860	\$6,860	\$6,860
Increased Worker Surplus Not Including Cost of Effort	\$20,772	\$20,772	\$20,772
Cost of Effort for Additional Hours Worked	-\$4,970 (\$0.50/hour)	-\$9,939 (\$1 per hour)	-\$20,772 (Full wage)
Total Surplus	\$22,662	\$17,693	\$6,860
<u>B. Social Cost</u>			
Fixed Cost of Employing (\$10 per hour spent)	\$476 3 min/worker	\$793 5 min/worker	\$1,587 10 min/worker
Deadweight Loss of Taxation (Total Cost: \$15,842)	\$0 (0%)	\$3,168 (20%)	\$3,961 (35%)
Worker Effort	\$3,622 (\$0.50/hour)	\$7,245 (\$1 per hour)	\$15,842 (Full wage)
Total Social Cost	\$4,098	\$11,206	\$21,390
<u>C. Overall Welfare Change</u>			
Total Surplus - Total Social Cost	\$18,564	\$6,487	-\$14,530

Notes: The "Total Surplus" row sums the three prior rows in Panel A. The "Total Social Cost" row sums the three prior rows in Panel B. The "Total Surplus - Total Social Cost" is the difference between these two summations.

Appendix Table 1. Verification of Randomization: Treatment Job

	No Previous Job		At Least 1 Previous Job		All Workers
	Treatment Job (1)	No Job (2)	Treatment Job (4)	No Job (5)	(6)
Posted Wage	1.97	1.94	2.01	2.04	1.98
Days Since Joining oDesk	137	126	251	257	179
Number of Applications Sent	25*	22*	160	167	27
Self-Assessed English Score	4.5	4.5	4.7	4.7	4.56
Number of Tests Passed	2.7	2.7	4.5	4.7	3.48
Philippines	63%	61%	63%	64%	63%
India	10%	11%	10%	12%	11%
Bangladesh	10%	10%	15%**	10%**	10%
Pakistan	6.3%	7.0%	5.1%	4.6%	5.9%
United States	2.9%	2.6%	2.3%	2.5%	2.6%
Number of Previous Jobs			7.3	6.9	6.9
Average Feedback Score			4.4	4.4	4.4
Observations	736	1562	216	1253	3,767

Notes: Each cell presents the mean value of the characteristic indicated by the row heading for workers indicated by the column heading before the experiment. One asterisk indicates the difference between treatment and control workers is significant at the 10% level and two asterisks indicate the difference is significant at the 5% level.

Appendix Table 2. Verification of Randomization: Detailed Comment

	Detailed Comment (1)	Uninformative Comment (2)
Posted Wage	1.99	2.03
Days Since Joining oDesk	163	164
Number of Applications Sent	55	53
Self-Assessed English Score	4.5	4.5
Number of Tests Passed	3.2	3.4
Philippines	67%	67%
India	9%	8%
Bangladesh	11%	9%
Pakistan	4.7%	4.9%
United States	1.6%	2.8%
Fraction with Previous Job	25%	25%
Workers with Previous Jobs Only		
Number of Previous Jobs	6.2	6.7
Average Feedback Score	4.6*	4.3*
Observations	317	327

Notes: Each cell presents the mean value of the characteristic indicated by the row heading for workers indicated by the column heading before the experiment. One asterisk indicates the difference between treatment and control workers is significant at the 10% level. Only hired workers who received a rating of four or higher are included.

Appendix Table 3. Effect of Detailed Comment: Instructions, Speed, and Accuracy
During the Two Months After the Experiment

	Total Jobs (1)	Any Job (2)	Hours Worked (3)	Posted Wage (4)	Earnings (5)
<u>A. Instructions</u>					
Follow All Instructions × Detailed Comment	0.579 (0.485)	0.097 (0.092)	-5.61 (15.78)	0.53** (0.26)	10.67 (37.72)
Follow All Instructions	0.022 (0.316)	0.024 (0.065)	4.43 (9.77)	-0.31 (0.18)	-6.95 (21.95)
Detailed Comment	-0.234 (0.415)	-0.032 (0.080)	8.20 (14.49)	-0.17 (0.20)	15.82 (34.57)
Sum of Detailed Coefficients	0.345	0.065	2.60	0.36	26.49
p-value of F-test	0.169	0.151	0.678	0.028	0.080
<u>B. Speed</u>					
Top Third × Detailed Comment	0.594 (0.471)	0.035 (0.081)	-15.14 (11.60)	0.11 (0.33)	-2.95 (29.08)
Top Third	0.108 (0.276)	0.017 (0.056)	12.91 (7.95)	0.14 (0.11)	17.28 (16.16)
Detailed Comment	0.004 (0.256)	0.030 (0.049)	9.74 (7.70)	0.21* (0.12)	25.82 (17.96)
Sum of Detailed Coefficients	0.598	0.065	-5.40	0.32	22.87
p-value of F-test	0.130	0.316	0.534	0.310	0.318
<u>C. Accuracy</u>					
Top Third × Detailed Comment	-0.075 (0.426)	0.063 (0.081)	-7.25 (11.98)	-0.16 (0.23)	-5.69 (29.60)
Top Third	-0.238 (0.283)	-0.072 (0.055)	0.09 (7.56)	-0.03 (0.11)	-0.93 (15.80)
Detailed Comment	0.227 (0.287)	0.016 (0.050)	6.45 (7.60)	0.30 (0.19)	26.06 (17.71)
Sum of Detailed Coefficients	0.152	0.079	-0.80	0.14	20.37
p-value of F-test	0.628	0.212	0.931	0.296	0.391
Observations	644	644	644	644	644

Notes: This table displays the results of estimating Equation 21 where x_i is an indicator for following all instructions (Panel A), an indicator for being in the top third of workers in speed (Panel B), or an indicator for being in the top third of workers in accuracy (Panel C). It includes only workers who received ratings of at least four and were thus eligible to receive a detailed comment. All experimental jobs are excluded. Huber-White standard errors are in parentheses. One asterisk indicates the coefficient is significant at the 10% level and two asterisks indicate the coefficient is significant at the 5% level.

Appendix Table 4. Estimates of Ease Index
 Dependent Variable: Indicator that an Application is Successful $\times 100$

Hourly Job	-0.755** (0.154)
Number of Applicants	-0.011** (0.000)
Data Entry	-1.693** (0.133)
Pref. for English Ability	-2.391** (0.134)
Pref. for Number of oDesk Hours	-0.588** (0.148)
Pref. for Level of oDesk Feedback	-1.457** (0.148)
Pref. for Maximum Wage below \$5	-0.791** (0.130)
Pref. for Minimum Wage above \$3	-0.007 (0.169)
Mean of Dependent Variable	3.57
Observations	156,184

Notes: This table presents the results of an ordinary least squares regression of an indicator for whether an application was successful (multiplied by 100) on the job characteristics listed in the left-most column and worker fixed effects. The unit of observation is an application. All applications sent by experimental workers in the month before the experiment are included. Standard errors are clustered by worker. Two asterisks indicate the coefficient is significant at the 5% level.