

Inefficient Hiring in Entry-Level Labor Markets

Amanda Pallais*

Harvard University and NBER

February 25, 2013

Abstract

Hiring inexperienced workers generates information about their abilities. If this information is public, workers obtain its benefits. If workers cannot compensate firms for hiring them, firms will hire too few inexperienced workers. I determine the effects of hiring workers and revealing more information about their abilities through a field experiment in an online marketplace. I hired 952 randomly-selected workers, giving them either detailed or coarse public evaluations. Both hiring workers and providing more detailed evaluations substantially improved workers' subsequent employment outcomes. Under plausible assumptions, the experiment's market-level benefits exceeded its cost, suggesting that some experimental workers had been inefficiently unemployed.

*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, Lawrence Katz, 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 seminar participants at Booth, Boston University, Columbia, Cornell, Harvard, Michigan State, MIT, Princeton, Stanford, UC-Irvine, UC-San Diego, University of Chicago, University of Virginia, Wharton, and Yale 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. Amanda Pallais's email address is apallais@fas.harvard.edu.

1 Introduction

Young workers are more likely to be unemployed than older, more experienced workers.¹ A key question in designing policies to improve young workers' labor market outcomes is whether their poor outcomes result from human capital deficiencies or barriers to labor market entry. If it is the former, then these workers may need to engage in intensive education or training programs to succeed in the labor market. If it is the latter, then programs that simply give these workers a foot in the door may have long-lasting benefits.

This paper evaluates whether inexperienced workers would benefit, on average, from simply obtaining a job because it would give them a chance to demonstrate their abilities. Employers are uncertain about the abilities of inexperienced workers. Hiring these workers generates information about their abilities (*e.g.*, Farber and Gibbons, 1996; Altonji and Pierret, 2001). This information is valuable because it allows firms to hire higher-ability workers in the future. But, firms might not have sufficient incentive to generate this information because, if it is partially public, workers receive part of its value as higher earnings. Hiring workers is costly. Managers must spend time explaining the jobs to workers and monitoring workers' progress. Moreover, firms incur an opportunity cost of lost time if jobs are not completed correctly or timely. If workers cannot compensate firms for hiring them and producing the information, for example because a minimum wage or adverse selection (*e.g.*, Weiss, 1980) prevents wages from falling or because bonding is prohibitively difficult (*e.g.*, Dickens *et al.*, 1989), firms will hire inefficiently few entry-level workers.

Through a field experiment in an online marketplace, this paper assesses the impact of (1) giving jobs to relatively inexperienced workers and (2) giving the market more information about workers' job performance on their future employment outcomes and the market as a whole. The online marketplace, oDesk, consists of workers all over the world who complete approximately 200,000 hours of work per week remotely.² Importantly, when an oDesk employer terminates a job, it is required to publicly disclose a rating of the worker on a one-to-five scale and can, if it chooses, provide a short comment.

In this experiment, I invited low-wage data-entry specialists to apply for 10-hour data-entry jobs. When workers applied, they proposed hourly wage rates for the job. The 3,767 workers who applied proposing wages of \$3 per hour or less formed the experimental sample. Workers in the sample were randomized into three groups: two treatment groups (with

¹For example, in December 2012, 12.8% of workers 20 to 24 years old were unemployed, compared with only 6.7% of workers 25 to 54 years old. These statistics are from the Bureau of Labor Statistics's Table A-13 (<http://www.bls.gov/web/empsit/cpseea13.pdf>), last accessed January 10, 2013.

²All statistics about the marketplace describe oDesk in July 2010, immediately after the experiment. This statistic was generated by oDesk and is reported at <https://www.odesk.com/economy/activity> (last accessed February 13, 2012).

476 workers in each) and a control group (containing the remaining 2,815 workers). The size of the treatment groups (together approximately 25% of the sample) was determined by financial constraints. I did not hire control group workers. I hired workers in both the "coarse evaluation treatment" and "detailed evaluation treatment" groups. I provided workers in both treatment groups with a public one-to-five rating, calculated from their actual performance statistics and normed to match the distribution of ratings in the market. The ratings were calculated without reference to what treatment group the workers were in. The difference between the two treatment groups was the amount of information about workers' performance that was in the public comment I provided. The comments workers in the coarse evaluation treatment received were designed to be as uninformative as the comments typically left in the marketplace. However, in the detailed comment treatment, workers receiving a rating of four or higher received a detailed comment with objective information about their data entry speed, accuracy, following of directions, and timely task completion. Due to IRB restrictions, I was not allowed to provide detailed evaluations to workers with low ratings. Thus, for workers earning below four, the detailed evaluation treatment was identical to the coarse evaluation treatment. Because of the large fraction of workers earning high ratings on oDesk, only 17% of workers earned ratings below four.

Using the marketplace's administrative data, I then observed the experimental workers' subsequent oDesk employment outcomes. Workers benefitted from obtaining an experimental job. After the experiment, workers in the coarse evaluation treatment were more likely to be employed, requested higher wages, and had higher earnings than control group workers. In the two months after the experiment, inexperienced workers' earnings approximately tripled as a result of obtaining a job. Providing workers with more detailed evaluations also increased their earnings and the wages they requested. This is consistent with the idea that more information about worker quality makes workers more valuable to firms. As theory suggests, the benefits of detailed evaluations were not universal: detailed performance evaluations helped those who performed well and hurt those who performed poorly.

The interpretation of these results depends on whether the treatments affected outcomes by revealing information about worker ability or through another mechanism. I consider whether five alternative mechanisms could explain the experiment's results: (1) the treatment jobs provided human capital, (2) the act of hiring workers led the market to positively update its belief about their abilities, (3) the fact that workers received detailed evaluations led the market to positively update its belief about their abilities, (4) obtaining an experimental job induced workers to apply to more oDesk jobs, but did not change employers' beliefs about workers' abilities, and (5) I gave workers more positive ratings than they deserved. None of these alternative explanations can explain all the experiment's results. For example,

the first two explanations cannot explain the results of the detailed evaluation treatment. Moreover, it does not appear that the act of hiring workers in itself led the market to positively update its beliefs about workers abilities: obtaining a job did not improve workers' employment outcomes when the market observed only that they had been hired (and not their evaluations), but outcomes improved immediately after the evaluations became public.

The fact that the treatments benefitted treatment group workers does not imply that they increased overall market welfare. That is, treatment group workers could have simply displaced other equivalent oDesk workers. While I do not have experimental variation that allows me to estimate the effect of the experiment on the market as a whole, I compare how employment and wages changed after the experiment across oDesk's 74 job categories, based on the intensity with which the categories were affected by the experiment. I find that, after the experiment, total employment increased in more-affected job categories relative to less-affected ones, while average wages decreased in the former relative to the latter. I use these results to estimate the experiment's effects on total market surplus. Under plausible assumptions, the benefits to market participants of the increased employment induced by the experiment outweighed the experiment's social cost (the time workers spent working and I spent managing them). While I do not directly observe outcomes outside of oDesk, this calculation does account for workers' and firms' opportunity costs of the increased oDesk employment. The result obtains despite the fact that the experimental jobs did not create usable output and I consider only benefits within six months of the experiment. It suggests that inefficiently low hiring of novice workers led to diminished employment and output in this market.

This paper directly relates to three strands of the literature. First is the literature on firm provision of general skills training. Public information about workers' abilities is similar to general human capital. While public information does not increase workers' output conditional on their working, it increases the aggregate output of a group of workers by allowing firms to hire only the highest-ability workers. Thus, discovering a worker's ability is similar to general skills training: both produce future productivity benefits, but require up-front investments. Becker (1964) 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 work shows that 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 shows that firms recoup the full

value of their training investments resulting in their providing the optimal level of training. This paper provides evidence that information about worker ability is underprovided by firms.

If not provided by firms, general skills training can be provided by schools; but output and information about workers' abilities are 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 may be few alternative mechanisms for its production.

The most closely related paper to this one is Tervio (2009), which proposes that the combination of hiring costs and publicly-observable performance could generate inefficiently high wages and low employment for CEOs and entertainers. This paper shows that a similar inefficiency may lead to inefficiently low employment in entry-level labor markets, making interventions that give workers a chance to demonstrate their abilities particularly effective. There is substantial uncertainty about the abilities of entry-level workers, particularly those with little education and few credentials. Firms often cannot conceal whether they have fired, retained, or promoted a worker, an important signal of entry-level worker performance. Workers' expected output is low, so minimum wages may be binding, making it difficult for workers to compensate firms for hiring them.

Entry-level labor markets have institutions and policies that, in theory, reduce this inefficiency: some reduce firms' cost of hiring inexperienced workers and some directly credential workers in return for compensation. For example, internships and hiring subsidies for young workers reduce firms' costs of hiring inexperienced workers. Fixed-term contracts (in Europe) reduce firms' hiring costs by allowing firms to dismiss low-ability young workers more easily. In many occupations, workers can pay to take tests demonstrating their competence at a given activity. Temporary help firms play a similar role. They screen workers for a variety of competencies (*e.g.*, Microsoft Word skills) and, in return for endorsing the worker, receive part of the worker's compensation.³ However, it is difficult for private firms to entirely remove the inefficiency. While tests can determine workers' skills and aptitude, they may not capture workers' dedication and enthusiasm. Moreover, while policies may reduce firms' costs of hiring workers, in most cases, workers are not legally allowed to pay firms for hiring them or agree to indentured servitude contracts.

Finally, this paper relates to the large literature evaluating whether programs that help young and disadvantaged workers enter the labor market can improve their long-term out-

³oDesk, the online marketplace that is the setting for the paper's empirical work, has similar institutions. Workers can take approximately 300 skills tests to demonstrate proficiency in subjects ranging from English to Microsoft Excel and C++. Moreover, Stanton and Thomas (2012) discuss agencies, which allow established workers to vouch for inexperienced workers in return for a percentage of their earnings.

comes. The findings from this literature are mixed.⁴ This paper has two primary advantages relative to this literature. First, while other programs typically combine many different elements, this experiment is able to isolate the effect of information about workers' abilities from on-the-job training, job placement services, or stigma from participating in a given program. Second, because the experiment was so large relative to the marketplace, this paper can address the concern that benefits for hired workers came entirely at the expense of other non-studied workers. I find that this was not the case: that the benefits to experimental workers far outweighed any cost to other oDesk workers.

The rest of the paper is organized as follows. Section 2 describes the online marketplace, lays out the experimental design, and assesses the randomization. Section 3 presents the theoretical framework and generates testable predictions for the effects of the experimental treatments. Section 4 analyzes the worker-level effects of the experiment and discusses whether they could have been generated by alternative mechanisms. Section 5 estimates the effect of the experiment on net market surplus. Section 6 concludes and discusses the application of these results to other settings as well as public policies that could potentially reduce the inefficiency.

2 Experimental Context and Design

2.1 The Marketplace

oDesk is an online marketplace in which employers hire independent contractors to perform tasks remotely. The marketplace is large: immediately following the experiment in 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. Right after the experiment, a plurality (37%) lived in the United States, while India (15%) and the Philippines (14%) were the next most common countries of residence.⁵ In contrast, approximately 80% of employers were located in the United States. The most common types of jobs on oDesk were web programming, website design, and data entry. In general, oDesk jobs were shorter than traditional, offline jobs. But, there was a lot of variation in the length of oDesk jobs: some jobs lasted for only a few hours, while others constituted full-time employment. The average job lasted 69 hours. Repeat interactions occurred, but were not the

⁴See, for example, Couch (1992), Bell and Orr (1994), Bloom *et al.* (1997), Autor and Houseman (2010), Redcross *et al.* (2009), Bloom *et al.* (2009), and Holister *et al.* (1984). Bloom (2010) and Stanley *et al.* (1998) provide summaries of the literature.

⁵All statistics in this section aside from the total number of hours per week worked on oDesk are from my calculations using the oDesk database.

norm. The average worker with any employment had 5.9 jobs with 4.7 unique employers.

Employers posted job openings in 74 job categories. These postings described the job and any necessary worker characteristics. When employers posted, they chose whether to offer hourly or fixed wage jobs. Hourly jobs, the type created in this experiment, constituted 70% of jobs on oDesk. In these jobs, oDesk tracked the number of hours worked and each worker earned an established hourly wage. oDesk guaranteed that the employer would pay for the hours worked, though the employer could stop the job at any time. In a fixed wage job, the worker and employer agreed to a price for the entire project, hours worked were not recorded, and the employer had complete discretion over how much it paid.

Workers posted public profiles, describing their skills and the types of jobs they were seeking. An example is displayed in Figure 1. (This worker was not in the sample because her proposed hourly wage was too high.) Workers could apply directly to jobs; alternatively, employers could search for workers and invite them to apply. When employers searched for workers, they could search for workers with different levels of experience, choosing to contact only workers who had worked a certain number of hours or had a certain feedback score. However, employers rarely invited workers in this sample to apply: less than 9% of applications sent by treatment group workers after the experiment were initiated by employers. Under either application method, the worker proposed 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 could hire as many or few applicants as they deemed suitable.

Each worker posted her preferred hourly wage rate at the top of her profile. When applying for a job, a worker could suggest a different wage to employers, but employers saw her posted wage as well. As soon as a worker began working in an hourly job, the job title, number of hours worked, and hourly wage were automatically posted to her profile. In fixed wage jobs, the job title and agreed job price were automatically posted. When an employer ended a job, it had to rate the worker from one to five on six dimensions: availability, communication, cooperation, deadlines, quality, and skills. These scores were averaged to form the worker's overall rating for the job. The worker rated the employer on the same six dimensions (before seeing her own rating); these scores were averaged to form the employer's overall rating. Because oDesk wanted the employee to be able to present her side of the story, both composite ratings were automatically posted to the worker's profile. A worker could not remove the ratings without refunding the remuneration received. Employers' ratings were typically very positive: before the experiment, 64% of low-wage data entry workers received a rating of exactly five, while 83% averaged at least four. Workers and employers could also choose to provide short comments about the employment experience, which were also automatically posted to the worker's profile. Comments were generally one or two positive

sentences providing little objective information. Unlike the numerical ratings, workers could remove employer comments without penalty, but overall only 4% of oDesk workers did.

In addition to the employer feedback mechanism, oDesk developed a number of ways to let workers demonstrate their abilities. Because workers' listed skills and experience could be hard to verify, oDesk developed its own skills tests: 40 minute, 40 multiple choice question tests on subjects such as written English, Microsoft Word, and C++. Workers' scores and performance relative to other oDesk workers could then be directly posted to their profiles. Workers could also display their "qualifications," certifications from other online platforms, and post a portfolio of their prior work.

Workers could also join agencies, groups of workers typically coordinated by an established worker, in return for a fraction of their earnings. The profile of each agency-affiliated worker contained the agency's average feedback score as well as the worker's own feedback score. Stanton and Thomas (2012) show that agencies were a way for workers to signal their quality. They find that agency-affiliated workers were much more likely to obtain a first job and earned higher wages in their first jobs than non-affiliated workers. However, once the market observed feedback on the workers' own abilities, workers no longer benefitted from being in an agency. While agencies were more common among high-wage workers (only 7% of my treatment group workers were in an agency), their presence suggests the difficulties oDesk workers had in developing reputations.

2.2 Sample Selection

I recruited subjects for this experiment by posting hourly data-entry jobs to the marketplace and inviting workers to apply.⁶ The jobs were 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 3,767 workers who applied to the jobs and requested a wage of \$3 or less formed the experimental sample.

Appendix 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

⁶I 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 a job.

with prior oDesk experience were substantially more likely to apply than inexperienced workers (54% and 33% applied, respectively).

2.3 Experimental Protocol

Figure 2 displays the experimental design. Workers were first randomized into either the control group or one of two treatment groups: the detailed evaluation or the coarse evaluation treatment group. Randomization into any treatment group was stratified on prior oDesk experience, such that workers without oDesk experience had a higher chance of being in any treatment group (32%) than experienced workers (15%). Conditional on receiving any treatment, all workers had a 50% chance of receiving the detailed evaluation treatment. Inexperienced workers constituted approximately three quarters of each treatment group.

The coarse evaluation treatment was designed to be equivalent to being hired (and, thus, evaluated) by a typical employer in the marketplace. The detailed evaluation treatment was identical to the coarse evaluation treatment except that it provided the market with more information about some workers' job performance. Workers in both treatment groups were hired and given a maximum of 10 hours over one week to enter the data. They were told that if, after spending 10 hours on the task, they had not completed it, they should send the file back unfinished. I recorded objective measures of workers' 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 rated all hired workers on a one-to-five scale using a weighted average of workers' scores on these performance measures. The distribution of scores from my job was designed to match the distribution of scores low-wage data entry workers received in the marketplace, adjusted for the fact that a worker in my sample was more likely to be inexperienced than a typical oDesk worker.⁷ The scores were calculated in the same way for workers in both treatment groups. Approximately 18% of workers did not return the file or log any hours. Under oDesk's protocol, these workers were not rated. Thus, the treatments should be considered as an intent to hire.

The particular treatment group to which workers were assigned affected only the type of comment workers were eligible to receive. No workers in either treatment group received a comment if they earned a rating below three. The remaining workers in the coarse evaluation treatment received an uninformative comment. The remaining workers in the detailed evaluation treatment received a detailed comment if they scored at least a four and an uninformative comment if they scored between three and four. (The human subjects committee

⁷In fact, the distributions of feedback scores received by experienced and inexperienced workers were not statistically distinct.

permitted detailed evaluations only for workers scoring at least four.)⁸ Workers in the detailed evaluation treatment did not know that they would receive a detailed evaluation until it was posted.

The uninformative comment was chosen to be short and positive, like most of the comments in the marketplace. The detailed comment provided objective information on the worker’s data entry speed and accuracy, whether the worker met the deadline, and whether she followed the job’s instructions. Additionally, it repeated the uninformative comment, so the only difference between the two comment types was the objective information provided in the detailed evaluation.

The uninformative comment read as follows, where only the words in brackets varied by worker.

“It was a pleasure working with [x].”

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].”⁹

Because such a high fraction of oDesk workers generally earned a rating of five, many treatment group workers who received this rating were in the bottom third of speed, accuracy, or both. If employers, particularly those new to oDesk, did not realize that so many workers received fives, these detailed comments would have appeared very negative.

I did not hire workers in the control group. However, some of these workers were hired by outside employers while the treatment group worked on my job. This was rare for workers

⁸MIT’s human subjects committee was concerned that giving workers negative evaluations would harm workers. It allowed me to give low numerical ratings, which were essential to the experiment. However, it permitted me to provide detailed comments only to 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 aspects of these workers’ performance (*e.g.*, they were in the bottom 10% of workers I hired in speed or accuracy).

⁹In order to test whether any effect of the detailed comment was a result of it simply being longer than the coarse comment or signaling that the worker was hired by a larger or more competent firm, I also randomized whether the comment mentioned that the hiring firm was large. I added the (true) sentence "Our organization has hired hundreds of providers on oDesk" to randomly-selected coarse comments and the sentence "This is based on our experience with hundreds of providers on oDesk," to randomly-selected detailed comments. These sentences had no effect on workers’ subsequent employment outcomes.

without prior experience: only 4% of inexperienced control group workers worked during this period. Unsurprisingly, a higher fraction of experienced control group workers, 27%, obtained jobs from other employers during this period. Because inexperienced workers comprised over three quarters of the treatment group, this suggests that fewer than 9% of treatment group workers would have been hired during this period in the absence of the experiment.

2.4 Data Collection

I directly collected data on workers' job performance. 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' profiles, job applications, and employment. The primary worker-level outcomes 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). I also use the wages workers posted to their profiles as a measure of their reservation wages. All workers had to post a wage to their profiles, so this measure is free from selection concerns. I observe the wage workers posted before the experiment and the timing of all subsequent changes to this posted wage, so I can determine the wage posted at any point in time. In a fully competitive market, workers would post their reservation wages. While workers do accept wages below their posted wages, there is no reason to believe the treatment affected the relationship between workers' posted wages and their reservation wages. Finally, I calculate workers' earnings from all oDesk jobs.

Three weeks after the initial randomization, I invited 630 workers to apply to another data-entry job with a fixed wage rate of either \$0.75, \$1, or \$2 per hour. These workers were randomly selected without reference to their prior experience or whether they had been placed into the control group or a treatment group. The invitation was sent from a new employer and workers were randomized into either the \$0.75, \$1, or \$2 job. I recorded which workers applied and offered a job to a randomly-selected 5% of applicants. I use data on whether workers applied to this job to calculate their opportunity cost of working in the welfare calculations.

2.5 Randomization Assessment

Tables 1 and 2 assess the randomization and present descriptive statistics about the sample. Table 1 shows that the majority (63%) of workers were from the Philippines, while relatively few (under 3%) were from the United States. On average, workers without prior experience

had been on oDesk for just over 4 months, passed 2.7 oDesk skills tests, and sent about 23 applications. Workers with previous jobs had been on oDesk for about twice as long and sent over seven times as many applications. They had an average of seven previous jobs (the median worker had four).

Table 1 compares the pre-experiment characteristics of workers in a treatment group with those of workers in the control group. It compares workers separately by prior experience because the randomization stratified on this variable. The treatment and control samples look similar based on covariates. Out of the 24 comparisons examined, one is statistically different at the 5% level and one is significantly different at the 10% level. In neither case, is there a significant difference between workers in the treatment and control groups when I pool the sample of workers with and without previous experience and control for workers' prior experience.

Table 2 compares workers randomized into the detailed and coarse evaluation treatment groups. Conditional on being in a treatment group, all workers had a 50% chance of being in either treatment group. The table also separately compares the covariates of workers who received ratings of four or five in the two treatment groups since these are the only workers for whom the two treatments differed. In both cases, the randomization produced similar samples.

3 Model

This section provides a simple framework that formalizes the insight that firms will hire inefficiently few inexperienced workers when they do not receive the benefit from discovering talented novices. It then defines two shocks to the market that are the model equivalents of the coarse and detailed evaluation treatments and generates predictions about the effects of these shocks.

3.1 Model Set-Up

The marketplace comprises a mass 1 of firms and potential workers. Workers (indexed by i) live for two periods (period 0, the “novice” period and period 1, the “veteran” period). Each period, one generation of workers with mass $\frac{1}{2}$ exits the market and a new generation enters. Firms (indexed by j) live for one period.

Workers vary in their ability (a_i), which is normally distributed in the population.¹⁰ Each

¹⁰The key aspect of this assumption is that it ensures that some workers have expected abilities below firms' hiring costs, and thus, they will not be hired. This assumption seems reasonable on oDesk: five percent of the workers I hired charged time for the job, but never turned in the entered data and 30% of those who

firm offers one identical task in which workers' output is $y_i = a_i$. Before a worker's novice period, the market observes her expected ability: $\hat{a}_{i0} = a_i + \varepsilon_{ia}$ where $\varepsilon_{ia} \sim N(0, \sigma_{ai}^2)$ and is independent of a_i . If the worker does not work in her novice period, the market does not update its beliefs about her ability and her expected ability before her veteran period (\hat{a}_{i1}) equals \hat{a}_{i0} . If a worker is employed in her novice period, the market observes a signal of her output, \hat{y}_{iM} :

$$\hat{y}_{iM} = a_i + \varepsilon_{iM} \text{ where } \varepsilon_{iM} \sim N(0, \sigma_M^2).$$

Firms use this signal and Bayesian updating to update their beliefs. For simplicity, workers have the same information about their abilities as the market.

Each firm 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 as well as any related overhead costs, such as for equipment or office space. It is continuously distributed across firms on $[0, \infty)$, generating a downward-sloping labor demand curve. Worker i 's net marginal product at firm j is $a_i - c_j$.

Each period, firms make wage offers to workers, who can accept at most one offer. Each agent is either in an employment relationship with wages w_{ij} or takes her outside option. Wages are restricted to be non-negative. Both firms' and workers' outside options are zero. Workers and firms are risk neutral and discount the future at rate $\beta < 1$. If agents are indifferent between an employment relationship and their outside option, they enter the employment relationship. The timing of events within each period is as follows.

1. A new generation of firms and novice workers enters the market.
2. Firms observe each worker's novice-period expected ability (\hat{a}_{i0}) and novice-period output signal (\hat{y}_{iM}), if it exists. They calculate each worker's expected ability.
3. Firms make wage offers to workers.
4. Workers accept or reject their wage offers.
5. Hired workers work, producing output y_i and receiving their wages.
6. Veteran workers and all firms exit the market. Novice workers become veterans.

did turn in the data entered over a third of the cells incorrectly. Hiring these workers for an actual job would provide low benefits compared to the time cost of administering the job and the opportunity cost of waiting for the worker to complete the task. In a more general context, it seems reasonable that some workers would have expected marginal products below the cost of hiring them (including any minimum wage). Low-skill workers who steal or do not show up for work impose large costs on firms compared to the benefits they produce if they perform well.

This model is stylized, but the assumptions are empirically-motivated. The assumptions that firms live only one period and that wages are non-negative rule out bonding and long-term contracts. In the oDesk setting, these assumptions seem reasonable. More than 50% of oDesk firms offered only one job ever. oDesk did not allow contracts with negative wages and it would have been difficult for third-world oDesk workers to transfer money to U.S. firms outside of the marketplace. More generally, Dickens *et al.* (1989) suggests bonding is difficult in the labor market. Workers may be liquidity constrained and bonding may negatively affect workers' attitudes and morale and harm firms' public images. Bonding may generate moral hazard problems where firms have incentives to report that workers performed poorly. Moreover, courts will not enforce contract provisions that call for workers to pay large penalties for poor performance and most indentured servitude contracts are unenforceable.

There was a minimum wage of 0 on oDesk. However, this assumption is not necessary. Any wage rigidity that prevents wages from falling will lead to a similar inefficiency where firms hire too few novice workers. For example, Weiss (1980) presents a model where adverse selection prevents wages from falling. While some unemployed workers would be willing to work for lower wages, because outside options are correlated with unobserved ability, high-ability workers would select out of the job if wages fell. Thus, firms will not decrease wages.

The assumption that firms live only one period precludes firms from having monopsony power over veteran workers as they do in models of asymmetric information (*e.g.*, Waldman, 1984 and Greenwald, 1986). While there is evidence of asymmetric information (*e.g.*, Gibbons and Katz, 1991 and Kahn, 2012), there is also evidence that public learning about workers' abilities is important for inexperienced workers and early career earnings dynamics (*e.g.*, Farber and Gibbons, 1996, Altonji and Pierret, 2001, and Schönberg, 2007). It is certainly possible that the oDesk employers that remain in the market have private information about their previous hires. To the extent that private information or mobility costs allow firms to hire previously-employed workers at wages below the workers' expected marginal products, this would increase the benefit of hiring novice workers and reduce the inefficiency. Asymmetric information would also decrease the benefits of being hired for an experimental job because the market would negatively update its beliefs about workers' abilities from the fact that I did not rehire them.

3.2 Market Equilibrium and Social Planner's Solution

Proposition 1 *The perfect Bayesian equilibrium of this game exists and has unique actions along the equilibrium path. There is 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 $w_{ij} = \hat{a}_i - \bar{c}$.

Proof. See Appendix A. ■

Because there is a mass of workers with expected abilities below every firm's fixed cost, not all workers will be employed. In particular, low-ability novices expected to generate losses when working at the marginal firm will not be hired. However, they would like to work and would compensate firms for hiring them because being hired in their novice period would increase their expected veteran-period earnings. If a low-ability novice does not work, her expected ability will not change and she will not be hired in her veteran period. If she works, with some probability, she will perform well enough to earn strictly positive veteran-period wages. However, novices cannot compensate firms for hiring them: they cannot accept negative wages, post bonds, or commit to accept low wages in the future. This generates the inefficiency. It also implies that this equilibrium would not change if workers knew their own abilities. High-ability workers would be willing to compensate firms more for hiring them than would lower-ability workers, but they would not be able to do so.

In specifying the social planner's solution, it is helpful to define \hat{a}_{i1H} : worker i 's veteran-period expected ability if she is hired in her novice period.

Proposition 2 *The solution of a utilitarian social planner who has the same information as the market and maximizes expected market surplus (the sum of expected worker and firm 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 $(\hat{a}_{i0} - c^*) + \beta \Pr(\hat{a}_{i1H} \geq c^*) \times E[\hat{a}_{i1H} - c^* | \hat{a}_{i1H} \geq c^*] \geq 0$. A larger mass of novice workers is employed in the social planner's solution than in the market equilibrium and $c^* > \bar{c}$.*

Proof. See Appendix A. ■

The social planner's solution is the equilibrium that would be enacted if novices could accept negative wages. This solution employs some novices who would be unemployed in the market equilibrium because they generate an expected loss in novice-period work. However, the expected veteran-period benefit these workers receive from novice-period work exceeds this expected loss. In general, hiring novices provides benefits in the subsequent period because it produces information about their abilities that allows them to be more efficiently allocated to either market work or unemployment. While firms have to pay the hiring cost to produce the information, workers obtain its benefits (they earn their expected marginal products which are higher because workers have been more efficiently allocated to sectors).

3.3 Model Predictions

Motivated by my experiment, I consider the comparative statics of two shocks to the market equilibrium. The first models the coarse evaluation treatment. In this shock, an outside employer (Employer C) enters the market for one period. This employer has the same information as the market at the beginning of the period, but hires novices without regard to their expected abilities. At the end of the period, it relays an output signal for each hired worker to the market. This output signal has the same distribution as the output signal generated in market jobs. (It represents the one-to-five rating left by employers.) The second shock models the detailed evaluation treatment. As in the prior shock, an outside employer (Employer D) enters the market for one period with the same information as the market, hires novice workers without regard to their expected abilities, and relays an output signal for each hired worker to the market. However, Employer D relays a more precise output signal than Employer C. It relays \hat{y}_{iD} where $\hat{y}_{iD} = a_i + \varepsilon_{iD}$, $\varepsilon_{iD} \sim N(0, \sigma_D^2)$ and is independent of all other variables and $\sigma_D^2 < \sigma_M^2$. (Here, \hat{y}_{iD} represents the one-to-five rating plus the detailed comment.) In both shocks, if a worker is not hired by the outside employer, she remains in the market and is either hired by another firm or takes her outside option.

An important assumption is that neither the fact that a worker was hired by an outside employer or the precision of the output signal conveys information about the worker's ability. The market updates its beliefs based only on the output signal itself. This assumption seems plausible. In the experiment, the market should not have updated its beliefs based on either of these factors since they were randomly determined (conditional on observables). Moreover, since all oDesk employers see the same worker profile and there is no face-to-face interaction, there is less scope for an employer to have private information before hiring on oDesk. The detailed evaluations were formulaic, often negative, and commonplace, characteristic of a particular employer, not the sign of a particularly talented worker. Nevertheless, I revisit this assumption when testing for alternative explanations in Section 4.2.

Proposition 3 *Relative to being in the market equilibrium, being hired by Employer C during the novice period weakly increases a worker's expected veteran-period probability of employment, earnings, and reservation wages. It strictly increases these outcomes for workers with $\hat{a}_{i0} < \bar{c}$. Relative to being hired by Employer C, being hired by Employer D during the novice period strictly increases a worker's expected veteran-period earnings and reservation wages, regardless of her novice-period expected ability, \hat{a}_{i0} . It increases her probability of veteran-period employment when $\hat{a}_{i0} < \bar{c}$ and decreases her probability of veteran-period employment when $\hat{a}_{i0} > \bar{c}$.*

Proof. See Appendix A. ■

The key intuition is that being hired by Employer C affects the expected veteran-period outcomes only of workers who would not have been hired in the market equilibrium (workers with $\hat{a}_{i0} < \bar{c}$). If they remain in the market equilibrium and are not hired in their novice periods, their expected abilities will not change and they will not be hired in their veteran periods. Their earnings and reservation wages will equal zero. But, if they are hired by Employer C, there is some probability that they will receive a sufficiently positive output signal to be employed with positive earnings and reservation wages in their veteran periods.

On the other hand, being hired by Employer D (relative to being hired by Employer C) affects workers of all expected abilities. The more precise performance signal causes the market to update its beliefs to a greater extent based on the output signal. This leads to lower probabilities of employment for workers who would have been employed without the output signal (those with $\hat{a}_{i0} > \bar{c}$) and to higher probabilities of employment for workers who would have been unemployed without the signal (those with $\hat{a}_{i0} < \bar{c}$). However, more updating increases expected earnings and reservation wages for workers of all expected abilities. Consider a worker with expected ability $\hat{a}_{i0} > \bar{c}$. If the market did not update its beliefs about her ability, she would be employed with certainty in her veteran period. Her earnings would equal her expected marginal product at the marginal firm, $\hat{a}_{i0} - \bar{c}$. This is the average of positive marginal products for states of the world when her true ability is above the hiring threshold and negative marginal products for states of the world when her true ability is below the threshold. If, instead, the market learned her true ability in her novice period, she would be unemployed in states of the world where her actual marginal product was negative. However, her expected earnings would no longer be driven down by the fact that she could generate a negative marginal product. She would still be rewarded for her positive marginal product in states of the world where her ability was above the hiring threshold, but states of the world where she was unemployed would contribute a zero marginal product. Thus, her expected earnings would be higher when the market knew her true ability even though her expected employment rate was lower.

In this model, a more precise output signal allows workers to be more effectively sorted into either market work or unemployment. This is the mechanism through which the signal's precision affects expected earnings. Without unemployment, the signal's precision would only affect the variance of earnings, not mean earnings.

Sections 4, 5, and 6 of Appendix A present and prove three additional propositions. The first, Proposition 5, says that the effect of being hired by Employer C, relative to remaining in the market equilibrium, is increasing in the initial uncertainty over the workers' ability (σ_{ai}^2). This is intuitive. When it is more uncertain about a worker's ability, the market puts

more weight on the output signal, so being hired by Employer C has a larger effect on the worker's expected veteran-period outcomes. However, counterintuitively, the proposition also says that the effect of being hired by Employer D relative to being hired Employer C is not necessarily increasing in σ_{ai}^2 . Employer D's output signal does cause the market to update its beliefs about workers with uncertain abilities more, but so does Employer C's output signal. Thus, relative to being hired by Employer C, being hired by Employer D does not necessarily have larger effects for workers with more uncertain abilities.

Proposition 6 says that while, *ex ante*, being hired by Employer D (relative to Employer C) increases all workers' expected earnings, *ex post*, being hired by Employer D can decrease workers' actual earnings. Workers who receive sufficiently poor output signals from Employer D have lower earnings and reservation wages than they would have if they had been hired by Employer C. Similarly, workers who receive sufficiently positive output signals from Employer D have higher earnings and reservation wages.

The stylized model presented above does not include the worker's choice to exit the market because including this choice adds little insight. However, because I observe workers exiting the market, I extend the model in Section 6 of Appendix A to give workers non-zero outside options and a choice to exit the market before each period. Proposition 7 says that being hired by Employer C in the novice period (relative to remaining in the market equilibrium) weakly increases workers' probabilities of remaining in the market in their veteran periods. This is because being hired by Employer C weakly improves all workers' subsequent market employment outcomes. On the other hand, being hired by Employer D (relative to being hired by Employer C) can increase or decrease the probability that a worker remains in the market. The market updates its beliefs more after a worker is hired by Employer D. This leads to lower probabilities of remaining in the market for workers who would have done so without the performance signal and higher probabilities of remaining in the market for those who would not have.

The final proposition considers the effect of Employer C on market employment, wages, and surplus.

Proposition 4 *If Employer C hires a non-zero fraction of novice workers in the subsequent period, the hiring threshold, \bar{c} , will increase to c' , total market employment will increase, wages will decrease conditional on expected ability from $\hat{a}_i - \bar{c}$ to $\hat{a}_i - c'$, and market surplus will increase.*

Proof. See Appendix A. ■

Employer C's hiring increases the mass of veterans with expected ability $\hat{a}_{i1} \geq \bar{c}$ in the subsequent period. This produces an excess supply of labor, which increases employment

and decreases wages, conditional on expected ability. Market surplus increases both because employment increases (more employment relationships generate surplus) and because workers hired by Employer C replace lower expected ability workers (the same mass of jobs produces more surplus).

Because being hired by Employer D decreases the expected veteran-period employment probabilities of workers with $\hat{a}_{i0} > \bar{c}$, Employer D's hiring a non-zero fraction of novices could decrease the mass of veterans with $\hat{a}_{i1} \geq \bar{c}$. This would decrease employment and increase wages, conditional on expected ability. This seems unlikely to happen in the empirical context as far less than 50% of control group workers were hired during the treatment period, suggesting that most workers had $\hat{a}_{i0} < \bar{c}$. Thus, I expect that the detailed and coarse evaluation treatments should have the same qualitative effect on market wages, employment, and surplus. However, even if Employer D decreased market employment, it would still increase market surplus by allowing employers to hire higher-ability workers.

4 Worker-Level Effects

4.1 Treatment Effects

I first assess the effects of the coarse evaluation treatment on workers' subsequent employment outcomes. Proposition 3 predicts that obtaining a job with a coarse evaluation will increase workers' employment rates, earnings, and reservation wages relative to being in the control group. Proposition 5 predicts that, conditional on workers' expected ability, the coarse evaluation treatment will have larger effects on workers about whom the market is more uncertain.

Table 3 compares the employment outcomes of the three experimental groups in the two months following the experiment. Workers are categorized by the treatment they were assigned to at the beginning of the experiment, even though 18% of workers in the treatment groups did not accept treatment jobs and workers earning low ratings in the detailed evaluation treatment did not receive detailed comments. Posted wages are measured at the end of the two-month period and the experimental jobs themselves are excluded from any outcomes. The results are presented separately for workers with and without prior oDesk experience as the randomization stratified on this variable. I use prior oDesk experience as a proxy for certainty about a worker's ability, assuming that the market is more uncertain about inexperienced workers' abilities. This is not a perfect proxy because experienced workers may have higher expected abilities than inexperienced workers.

The coarse evaluation treatment's effects reflect the model's predictions: it had positive

effects on employment outcomes for workers without prior oDesk experience, but no effect for experienced workers. Inexperienced control group workers performed poorly in the labor market: only 12% obtained any job in the next two months for an average earnings (unconditional on working) of approximately \$10. The coarse evaluation treatment significantly improved all five employment outcome measures for inexperienced workers. It almost tripled the fraction of these workers with any employment from 12% to 30% and the average worker's earnings from \$10 to \$27. It also increased the wage these workers posted on their profiles by approximately 10%.

Experienced control group workers performed much better than inexperienced control group workers: over half (55%) worked on oDesk in the two months after the experiment for an average earnings of \$121. However, the coarse evaluation treatment did not significantly improve any of the five employment outcomes for experienced workers.

Table 3 also allows me to assess the effects of the detailed evaluation treatment. Proposition 3 predicts that workers in the detailed evaluation treatment will have higher earnings and reservation wages than those in the coarse evaluation treatment, while the treatment's effect on employment is ambiguous. The table shows that, relative to the coarse evaluation treatment, the detailed evaluation treatment increased experienced workers' average earnings from \$101 to \$187 and their average posted wages by approximately 15%. The earnings gains did not come at the expense of employment; the detailed evaluation increased the fraction of workers with any subsequent employment from 53% to 69%.

However, the detailed evaluation treatment did not improve average employment outcomes of inexperienced workers relative to the coarse evaluation treatment. While workers in the detailed evaluation treatment had better outcomes on four out of the five employment measures than workers in the coarse evaluation treatment, the differences are neither large nor significant. There are three potential explanations for the somewhat surprising result that the detailed evaluation treatment had larger effects for experienced than inexperienced workers. First, Proposition 5 says that, conditional on expected ability, the effect of the detailed evaluation relative to the coarse evaluation is not necessarily increasing in uncertainty about worker ability. Second, the detailed evaluation should have the most impact for workers with expected abilities near the hiring threshold because there is the most uncertainty about whether these workers should be hired. Experienced workers may have expected abilities that are closer to the hiring threshold. Finally, a higher fraction of experienced than inexperienced workers in the detailed evaluation treatment actually received a detailed evaluation (74% vs. 66%).

I next consider the robustness of the treatments' effects on the pooled sample of experienced and inexperienced workers. Panel A of Table 4 displays the results of regressing

each employment outcome on an indicator for being in the detailed evaluation treatment, an indicator for being in the coarse evaluation treatment, and indicator for having prior oDesk experience. I control for a number of covariates.¹¹

The table shows that both treatments improved employment outcomes for the pooled sample. Relative to the control group, the detailed evaluation treatment increased the fraction of workers employed in the next two months by 17 percentage points, while the coarse evaluation treatment increased this fraction by 13 percentage points. Both treatments increased the wage workers posted on their profiles (by \$0.26 for the detailed treatment and \$0.15 for the coarse treatment relative to the control group mean of \$2.19). The detailed treatment increased workers' earnings by \$29 on average. The regression indicates that the coarse treatment increased earnings by approximately \$10, 17% of the control group mean, but this coefficient is far from significant. While the treatment effects are large relative to the control group's average employment outcomes, they are smaller than the coefficient on the indicator for prior oDesk experience. This is expected as this latter coefficient includes both the treatment effects of multiple previous jobs and differential selection into employment.

Panel B shows the effect of receiving any experimental job, whether or not it came with the possibility of a detailed evaluation. It displays the results of regressing each employment outcome on an indicator for receiving any treatment job, an indicator for prior oDesk experience, and the same controls as in the previous panel. It shows that obtaining a treatment job increased workers' average earnings by approximately \$20, which exceeds the average amount I spent to hire a worker (\$16.64).

Panel A of Table 5 further probes the impact of the detailed evaluations by estimating their effects on workers who received a score of four or five in my jobs (those who were eligible to receive detailed evaluations). It displays the results of regressing each employment outcome on an indicator for being in the detailed evaluation treatment group. I do not control for prior oDesk experience since, conditional on receiving any treatment, all workers had a 50% chance of being assigned to each treatment. The treatment a worker was assigned to did not affect her rating; the treatments were identical until the evaluation was made public. The table shows that receiving a detailed evaluation increased earnings and posted wages relative to receiving a coarse evaluation. Detailed evaluations increased average earnings by \$24 within two months, 44% of the mean for workers receiving fours and fives in the coarse

¹¹I control for the tests the worker took (dummies for passing one test, two or three tests, and four or more tests, an indicator for having taken the most popular skills test among the sample, and an indicator for passing it), the number of qualifications the worker had (dummies for listing two to four qualifications and five or more qualifications), the number of oDesk applications the worker sent before the experiment (dummies for having sent three to five, six to 15, 16 to 50, 51 to 100, and over 100 applications), the wage proposed for the experimental job (dummies for offering \$1 to \$1.99, \$2 to \$2.99, and exactly \$3), the number of jobs the worker had before the experiment, and an indicator for being in the second experimental wave.

evaluation treatment. They increased workers' average posted wages by \$0.24, 10% of the mean for high-scoring coarse evaluation workers. Both effects are significant at the 10% level. After one month, these effects are significant at the 5% level. Appendix Table 2 shows that the estimates do not change when the controls used in the previous table are added.¹²

The effects of the treatments persisted and even increased over time. I calculate each worker's cumulative employment and earnings outcomes for each week from one to 26 weeks after the treatment jobs ended. I also record each worker's posted wages at the end of each week. Then, I regress these employment outcomes on an indicator for receiving any treatment job, controlling for prior oDesk experience. Panels A and B of Figure 3 plot the coefficients on the treatment indicator by week. The effect of obtaining a treatment job on earnings more than tripled from \$19.72 after two months to \$65.56 after six months. The effect of the treatment job on hours worked (which is not plotted) almost tripled, increasing from 9.7 after two months to 25.3 after six months (the p-value is only 0.11 after six months). It is striking that the treatment effects are so large given that the treatment was not restricted to those who would most benefit from it: some treatment group workers already had so much prior experience and such high expected abilities that they were unlikely to benefit, while others had such low expected abilities that even successfully completing one job would not increase their expected abilities above the hiring threshold.

The effects of receiving a detailed evaluation also appears to persist over time, but they are much noisier. For example, when I regress cumulative earnings on an indicator for receiving a detailed evaluation in the sample of workers earning fours and fives in my jobs, the effect is positive and significant at the 5% level through week seven and significant at the 10% level through week nine. After six months, the measured effect on earnings is \$60.17, 2.5 times the effect after two months, but it has a t-statistic of only 1.1.

Finally, I test Proposition 6, which says that detailed evaluations do not always help workers. It says that detailed evaluations that are more negative than both a worker's initial expected ability and the coarse evaluation she would have received impair subsequent employment outcomes relative to coarse evaluations. Detailed evaluations that are more positive than both her initial expected ability and the coarse evaluation she would have received improve outcomes. Because I know all treatment group workers' performance, I can determine the effect of a particular detailed evaluation by comparing the outcomes of

¹²Another robustness check is to measure the effect of the detailed evaluation treatment for workers scoring below four. Relative to the coarse evaluation treatment, the detailed evaluation treatment should have no effect for these workers. Replicating Panel A of Table 5 for workers rated below four shows that one of the five coefficients, the effect on the number of jobs obtained, is significantly positive. This is driven by a few outliers, one who obtained 24 jobs and one who obtained 15 jobs in the two month period, relative to a maximum of 5 jobs among workers receiving coarse evaluations. None of the other coefficients are significant and the point estimate indicates detailed evaluations had a negative effect on posted wages for this group.

workers who received that evaluation with the outcomes of workers who would have received that evaluation had they been in the detailed evaluation treatment.

I consider the effect of revealing whether a worker met the deadline on workers earning fours and fives in my jobs. Since the vast majority of workers earning fours and fives met the deadline, not meeting the deadline was a very negative signal. It is likely more negative than the coarse rating of four or five and most workers' initial expected abilities. Thus, revealing it should have impaired workers' subsequent employment outcomes. Meeting the deadline was a positive signal, but may not have been more positive than a rating of five or many workers' initial expected abilities. Panel B of Table 5 shows the results of regressing each employment outcome on an indicator for meeting the deadline interacted with an indicator for being in the detailed evaluation treatment, an indicator for missing the deadline interacted with the detailed treatment indicator, and the indicator for meeting the deadline itself.

Revealing that a worker missed the deadline decreased her earnings in the subsequent two months by \$35 on average. Revealing that she met the deadline increased her earnings by \$24. These effects are large relative to mean earnings (\$54 for workers with ratings of four or five in the coarse evaluation treatment). The estimates indicate that revealing whether the worker met the deadline had large effects on the other employment outcomes as well, but these are imprecisely measured.

Whether the worker met the deadline is the only piece of information that appeared to matter to employers. Appendix Table 3 shows that revealing that workers did not follow all the instructions, were in the bottom third of workers in speed, or were in the bottom third of workers in accuracy did not impair their employment outcomes. This could be because these are less negative signals than missing the deadline. Alternatively, it could be because whether the worker met the deadline was the first piece of information revealed in the comment. Because of the comment's length, the parts about speed and accuracy were not immediately visible on most workers' profiles; one had to click on the continuation to see them. This was unintentional.

4.2 Mechanisms

The previous section showed that the treatments' effects were consistent with the model's predictions. Here, I assess whether the experimental results could also be explained by alternative mechanisms. If the model's mechanism is correct, then workers with higher ratings should have better subsequent employment outcomes, all else equal. This is true in the data: after two months, workers who received ratings of five had earned \$34 more than the control group, while workers who received ratings of one and two had earned \$23

less. However, since that these ratings were not randomly assigned, it could simply be that highly-rated workers would have earned more even without the experiment.

The first alternative is that completing a treatment job provided workers with human capital. Many existing programs that employ disadvantaged workers explicitly provide on-the-job training and, even in those that do not, workers could easily gain human capital by working for several months. It is much less likely that workers accumulated human capital in these jobs. Workers worked a maximum of 10 hours; the average hire worked for only 7.6 hours. Given workers' offline experience, this was 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, workers in both treatment groups completed the same task, so human capital accumulation cannot explain the effects of the detailed evaluations.

The second alternative is that the act of hiring a worker caused the market to positively update its belief about the worker's ability. Hiring a worker would cause the market to positively update its belief if different employers received different signals of worker quality. This is less likely on oDesk than in a traditional labor market because, on oDesk, all employers saw exactly the same resume and there were no face-to-face interactions. However, employers could interview workers and they might have valued the same information differently.

An empirical test of this explanation utilizes the fact that the market observed that treatment group workers were hired as soon as they began working, but could not see their ratings until nine to 11 days later. In the week the workers were completing my job (during which the market only observed that they had been hired, not their evaluations), treatment group workers were no more likely to obtain jobs from other employers than control group workers. In contrast, in the week immediately after the rating became public, treatment group workers obtained significantly more jobs and were more likely to be employed than control group workers. This is not simply because treatment group workers were too busy to apply to jobs while they were completing my task: they actually applied to more jobs than control group workers during this period. Moreover, it does not appear that oDesk employers typically penalized job applicants for being currently employed on other jobs. Among oDesk workers with at least twenty previous jobs (whose reputation should not have substantially changed with another evaluation), a given worker's job application was slightly more likely to be successful if she applied while working on another job.

This alternative cannot explain the effects of the detailed evaluations since the market observed that workers in both treatment groups were hired. However, a similar alternative is that the market positively updated its beliefs about a worker's ability based on the fact that she received a detailed evaluation. I think this is unlikely. The detailed comments were often negative. Twenty-seven percent of workers were revealed to be in the bottom third or

bottom 10% of speed, while 28% were revealed to be in the bottom third or bottom 10% of accuracy. (These are less than 33% because workers earning ratings below four were not given detailed comments.) Less than 19% of workers were described as being in the top third or top 10% of both speed and accuracy. Perhaps more importantly, these comments were formulaic, contained no subjective information, and were common in the marketplace. In the two months after the experiment, only 12% of the applications workers with detailed evaluations sent went to firms that did not have another applicant who had received a detailed evaluation (with the exact same formula). Thus, these comments likely appeared to be the hallmark of a particular employer, not a particularly good worker.

A fourth alternative is that receiving a treatment job induced workers to apply to more oDesk jobs for reasons unrelated to the evaluations. For example, treatment group workers may have realized oDesk jobs were more desirable than they had thought or their initial hiring may have led them to believe it was easy to obtain oDesk employment. However, Proposition 7 also predicts that the coarse evaluation treatment should have induced workers to remain in the market by improving their oDesk employment opportunities. It is difficult to distinguish these two explanations because they have the same prediction. This prediction is borne out in the data: treatment group workers did apply to more jobs than control group workers after the experiment. However, the alternative explanation cannot explain the entire effect of the treatments because the treatments significantly increased the probability that a worker obtained a given job she applied to. Appendix B describes these results in much more detail. The explanation similarly cannot explain why receiving a treatment job increased workers' posted wages or the effects of the detailed evaluation, which did not alter workers' application patterns.

A final alternative is that I gave workers more positive ratings and comments than they deserved, despite the facts that my ratings matched the distribution of one-to-five ratings in the market (controlling for the relative inexperience of my workers) and that the detailed comment contained some objective information. If this explanation were correct, the treatments should have had a diminished effect over time as the market learned more about the workers' true capabilities. The data do not show this pattern: the effects of both treatments on weekly earnings appears to remain constant over time.

5 Welfare Analysis

In this section, I estimate a lower bound on the experiment's effect on oDesk market surplus in the six months after the experiment. Proposition 4's proof suggests that the experiment increased welfare through two channels: (1) by increasing employment and (2) by allowing

firms that would have hired workers in the absence of the experiment to hire workers with higher expected abilities. In this section, I estimate the effect of the extra employment on market surplus. I compare this increase in surplus to the costs of the experiment: the direct costs of the experimental jobs and workers' and firms' opportunity costs of the extra employment. I consider this a lower bound on oDesk market surplus because I cannot estimate the benefit firms received from hiring workers with higher expected abilities. However, this calculation only includes surplus obtained by oDesk workers and firms, not other agents that were affected by the experiment. For example, when workers increased their oDesk employment, they may have forgone offline jobs. This calculation includes their opportunity costs of not taking the offline jobs, but it does not account for either the lost profit of the firms that would have hired them or the increased earnings of the workers who took the offline jobs in their absence.

I first estimate the effect of the experiment on market employment. While I do not have experimental variation I can use to estimate this, I can compare the change in employment after the experiment in those oDesk sectors more and less affected by the experiment. For each of oDesk's 74 job categories, I calculate a measure of the experiment's effect on the number of experienced workers in the category: "% change experience." First, I estimate the number of workers I hired in each category: 952 (the total number of workers I hired) multiplied by the share of treatment group applications sent to jobs in that category in the month before the experiment. Then, I divide this by the number of experienced workers working in that category before the experiment. This fraction averaged 8.5% for the entire marketplace, ranging from 55% in data entry and 79% in email response handling to less than 1% in 25 job categories, primarily ones that required specific computer skills such as web programming or game development.¹³

Panels A and B of Table 6 show the results of regressing two measures of log employment (log jobs created and log hours worked) in a job category-week on the interaction of "% change experience" and an indicator for a week after the experiment.¹⁴ I control for week fixed effects, job category fixed effects, and job category-specific linear time trends. Each regression includes 26 weeks of data before the experiment, omitting the weeks of the experiment and all jobs I offered. The first three columns use data roughly two, four, and six months (8, 17, and 26 weeks) after the experiment. The fourth column adds job category-specific

¹³The variable "% change experience" should not be interpreted as the actual percentage change in the number of workers with any oDesk experience because 23% of the treatment group already had oDesk experience. However, it provides a simple metric measuring the intensity with which the experiment affected each job category.

¹⁴For the one week between the two waves of the experiment, I set this indicator equal to $\frac{608}{952}$, the fraction of hired workers in the first wave.

quadratic time trends and the fifth drops data entry, which accounts for approximately half of the experimental workers' employment. Observations are weighted by job category size in a preperiod and standard errors are clustered by job category. The key assumption is that, conditional on job category-specific linear and quadratic trends, employment would have changed similarly after the experiment in job categories more and less affected by the experiment.¹⁵

In all three time windows, the effect of the experiment on new jobs was positive and significant. The elasticity ranged from 0.22 after two months to 0.27 after six months. The effect on hours worked was also positive across all three time periods, with an elasticity ranging from 0.10 to 0.14. Including job category-specific quadratic time trends has virtually no effect on the employment estimates and the results persist even when data entry is excluded.

The total increase in worker earnings due to this increased employment is the level change in hours worked multiplied by the average hourly wage in these new jobs. This excludes any increase in fixed wage employment. To calculate this change in hours worked per week, I multiply the most conservative hours coefficient in Table 6 (0.099) by the average hours worked per week and the average value of % change experience. This increase of 950 hours worked per week is smaller than the increase in the number of hours worked by the treatment group (1010), but not much smaller, suggesting that treatment group workers crowded out relatively little other employment.¹⁶

To determine the average wage in these new jobs, I compare the wage distributions of treatment and control group workers after the experiment, adjusting for the stratification.¹⁷ Some of the new jobs obtained by treatment group workers represented aggregate employment increases while the rest replaced jobs that would have been held by lower expected-ability workers. I define the replacement jobs as the fraction $\frac{(1010-950)}{1010}$ of these new jobs

¹⁵In approximately 2% of job category-weeks, there were no jobs created and in approximately 0.4% of job category-weeks, there were no hours worked. For these observations, I set log jobs created and log hours worked equal to zero (the log of 1). This assumption has very little impact on the results. Dropping any job category with any week with zero jobs created does not change any of the coefficients by more than 0.01, while imputing log jobs created as $\log(0.1)$ instead of $\log(1)$ does not change any of the coefficients by more than 0.002. For the hours regressions, neither dropping any job category with a week with zero hours worked nor imputing log hours worked as $\log(0.1)$ changes any of the coefficients by more than 0.001.

¹⁶While the model in Section 3 predicts that experimental workers should displace non-experimental workers, other models would predict that an increase in the number of oDesk workers recognized to be of adequate ability might actually improve other workers' employment outcomes, for example by inducing new employers to join the marketplace or existing employers to remain in the marketplace.

¹⁷I do this by categorizing jobs based on their wages: wages less than \$1, exactly \$1, \$1 to \$2, exactly \$2, \$2 to \$3, exactly \$3, and \$3 or more per hour. I calculate the number of additional hours the treatment group worked over the control group in each category, adjusting for the stratification. I calculate the average wage in each of these categories and use these to form a weighted average of the wage at which the extra hours were worked.

with the lowest wages. This provides the average wages in the new jobs as \$2.11, slightly lower than treatment group workers' average wage after the experiment (\$2.20). Using these estimates, the total increase in worker earnings from the employment expansion was

$$\underbrace{\left(26 \text{ weeks} \times 950 \frac{\text{hours}}{\text{week}}\right)}_{\text{Increase in Employment}} \times \underbrace{\$2.11}_{\text{Wage Level}} \approx \$52,000.$$

Firms also benefitted from the increased employment. In particular, the theory predicts that firms with fixed costs between \bar{c} and c' hired oDesk workers after the experiment. Their average profit was $E[\hat{a}_i - w_{ij} - c_j | \bar{c} < c_j \leq c'] = E[\hat{a}_i - (\hat{a}_i - c') - c_j | \bar{c} < c_j \leq c'] = E[c' - c_j | \bar{c} < c_j \leq c']$. If c_j is uniformly distributed between \bar{c} and c' , this is $\frac{1}{2}(c' - \bar{c})$. I can approximate $c' - \bar{c}$ using the decrease in market-level wages induced by the experiment. This is only an approximation, however, because the change in market wages captures both the change in the marginal firm's hiring cost and the change in hired workers' expected abilities. If hired workers' expected abilities increased after the experiment, the market-level wage change is a lower bound on the change in the fixed cost of the marginal firm.

Panel C of Table 6 estimates the experiment's effect on market-level wages. It regresses the hourly wage for a given job in a given week on interaction of the job category's % change experience with the indicator for a week after the experiment. It controls for week fixed effects, job category fixed effects, and job category-specific linear time trends. Assuming that, conditional on the category-specific trends, wages would have changed similarly in affected and unaffected categories, the results show that the experiment substantially decreased wages. This is true in all three time periods, with elasticities ranging from 0.27 to 0.44. While adding category-specific quadratic time trends does not substantially affect the estimates, the experiment's effect on wages was driven primarily by its effect on the data-entry market. I determine the experiment's effect on the wage level (\$0.22) analogously to the level change in hours using the same column's estimate of the wage change (-0.265). If c_j is uniformly distributed, the increase in firm profit from the additional employment was:

$$\underbrace{\left(26 \text{ weeks} \times 950 \frac{\text{hours}}{\text{week}}\right)}_{\text{Increase in Employment}} \times \underbrace{\frac{1}{2} \times \$0.22}_{\text{Average Profit/Job}} \approx \$2,800.$$

Table 7 estimates the overall effect of the experiment on market surplus in the six months after the experiment by combining the above estimates of the experiment's benefits with estimates of the experiment's costs. The middle column presents the most realistic estimates, while the right- and left-hand columns make more and less conservative assumptions, re-

spectively. Panel A presents the measured benefits of the experiment. Panel B estimates workers' and firms' opportunity costs of the additional employment after the experiment, while Panel C estimates the cost of the experimental jobs themselves.

I assess workers' opportunity costs of working by using their willingness to apply to the jobs I posted three weeks after the experiment. I assume that workers applied to these jobs if and only if their opportunity costs were below the jobs' wages. If there was a time cost of applying, this overstates workers' opportunity costs, particularly for workers with low probabilities of getting the jobs. (Experienced workers were more likely to apply to these jobs than were workers with fewer other options, suggesting that the probability of obtaining a job might be an important determinant of application behavior.) This may also slightly overestimate the average opportunity cost of the work induced by the experiment because workers with lower posted wages saw (insignificantly) larger employment gains as a result of the treatments. Because 49% of the treatment group applied to the \$1 wage job, I use this as the "best-guess" estimate of workers' average opportunity cost. To make a more conservative estimate, I assume that if workers did not apply to my second job, they were not willing to accept any wage below their initial posted wage. I categorize treatment group workers into five groups based on the wages they proposed for the initial treatment job. Then, I calculate an average opportunity cost for each group using the fraction of workers that applied to the \$0.75, \$1, and \$2 jobs and their posted wages.¹⁸ The weighted average of these opportunity costs for the entire sample was \$1.49.

I have less information on firms' opportunity costs. However, as I estimate that the increase in firm profits was only about 5% of the increase in worker earnings, even assuming that firms did not benefit from the increased employment does not affect the welfare conclusions.

Panel C assesses the opportunity cost of the experimental jobs themselves. The opportunity cost of workers' time is calculated as in the previous panel. In a traditional job, these costs would have been offset by the value of the output produced. However, in this experiment, the output had no value. Because workers were not expected to produce usable output, my time cost of employing them was relatively low (five minutes per worker).

Even under the conservative assumptions in the right-most column, the estimates suggest that hiring 952 randomly-selected, relatively inexperienced workers for a meaningless task

¹⁸These five categories are workers proposing wages below \$1, \$1 to \$2, exactly \$2, \$2 to \$3, and exactly \$3. For example, 21.4% of the group proposing wages of \$3 was willing to accept a job with a \$0.75 wage, an additional 19.5% was willing to accept a job with a \$1 wage, and an additional 12.9% was willing to accept a job with a \$2 wage. I assume the remaining 46.2% was not willing to accept any wage below \$3. So, the average opportunity cost of workers offering wages of exactly \$3 was estimated to be $21.4\% \times \$0.75 + 19.5\% \times \$1 + 12.9\% \times \$2 + 46.2\% \times \$3 = \$2$.

increased market surplus by over \$2,700 in the subsequent six months. As the benefits of this experiment for treatment group workers and the overall market appear to have increased steadily over time, extending the time frame over which the benefits were calculated would likely lead to larger estimates of the experiment’s effect on market surplus. These calculations suggest that, before the experiment, novice employment on oDesk was inefficiently low.

6 Conclusion

There is a debate in the literature over whether simply helping young and disadvantaged workers enter the labor market can improve their long-term employment outcomes or whether intensive skills training is required. This paper proposes that merely giving workers a first job benefits them by providing the market with information about their abilities, which in turn, makes them more valuable to firms. However, to the extent this information is public, its benefits accrue to workers, so firms may hire too few inexperienced workers. In particular, firms will hire too few inexperienced workers if hiring is costly, they do not receive the benefits of the information produced, and workers cannot fully compensate them for being hired.

This paper uses a field experiment in an online marketplace to test whether giving workers a chance to demonstrate their abilities improves their labor market outcomes. In the experiment, workers were randomly selected to receive a job with a coarse evaluation, a job with a detailed evaluation, or no job. Simply giving workers a job substantially increased their subsequent employment rates, earnings, and reservation wages. Giving the market more detailed information about their job performance also increased their average earnings and reservation wages. These results are consistent with the paper’s proposed mechanism: information about their abilities made workers more valuable to employers, but are inconsistent with several alternative mechanisms. Despite the fact that the experiment was not designed with this purpose, under plausible assumptions, it increased market surplus by more than its social cost, suggesting that, before the experiment, oDesk firms hired inefficiently few inexperienced workers.

These results come from a particular marketplace and an important question is whether and how they would generalize to other contexts. The oDesk setting is probably most similar to traditional low-wage labor markets. It is characterized by high unemployment rates and its data-entry workers earn wages just above the allowable minimum. However, oDesk employers are likely more uncertain about the abilities of job applicants than offline employers of low-wage workers. They typically have less hiring experience, may be unfamiliar with credentials from foreign schools or employers, and have limited ability to verify these credentials. This suggests that the benefits of performance evaluations may be particularly

large in the oDesk marketplace. On the other hand, oDesk has particularly low hiring costs. There is no positive minimum wage on oDesk and firms do not provide employment benefits or supplies. oDesk workers can be hired with the click of a mouse and fired instantaneously with no penalty. The tasks they complete are typically well-defined, easy to explain, and require no training. This suggests that the labor market inefficiency and the benefits of the treatments might be larger in a traditional labor market.

Assuming these results would generalize to traditional low-wage labor markets, there are several public policy responses that might reduce the inefficiency. First, a government could partially or fully subsidize firms for hiring young workers. Second, the government could itself hire young workers. For this to be maximally effective, the government would have to provide the market with honest measures of worker performance and ensure the jobs would not be negatively stigmatized. Avoiding stigmatization might entail hiring some workers already recognized by the market to be high-ability (as I did in this experiment). This would increase the program's costs but also the output of hired workers. Finally, the government or a subsidized firm could provide employment tests, simulated work experiences designed to reliably measure workers' capabilities and diligence. These tests could be designed (as was the task in this experiment), to require little managerial time, but provide a useful signal of workers' performance to potential hiring firms.

As most oDesk jobs are offshored from U.S. employers to foreign workers, the paper's results 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, in new plants and capital. In contrast, 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), F112–F142.
- Altonji, J. and C. Pierret (2001). Employer Learning and Statistical Discrimination. *Quarterly Journal of Economics* 116(1), 313–350.
- Autor, D. (2001). Why Do Temporary Help Firms Provide Free General Skills Training? *Quarterly Journal of Economics* 116(4), 1409–1448.
- Autor, D. and S. Houseman (2010). Do Temporary-Help Jobs Improve Labor Market Outcomes for Low-Skilled Workers? Evidence from 'Work First'. *American Economic Journal: Applied Economics* 2(3), 96–128.
- Becker, G. (1964). *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education*. Columbia University Press.
- Bell, S. and L. Orr (1994). Is Subsidized Employment Cost Effective for Welfare Recipients? Experimental Evidence from Seven State Demonstrations. *Journal of Human Resources* 29(1), 42–61.
- Bloom, D. (2010). *Transitional Jobs: Background, Program Models, and Evaluation Evidence*. New York, NY: MDRC.
- Bloom, D., S. Rich, C. Redcross, E. Jacobs, J. Yahner, and N. Pindus (2009). *Alternative Welfare-to-Work Strategies for the Hard-to-Employ: Testing Transitional Jobs and Pre-Employment Services in Philadelphia*. New York, NY: MDRC.
- Bloom, H., L. Orr, S. Bell, G. Cave, F. Doolittle, W. Lin, and J. Bos (1997). The Benefits and Cost of JTPA Title II-A Programs: Key Findings from the National Job Training Partnership Act Study. *Journal of Human Resources* 32(3), 549–576.
- Couch, K. (1992). New Evidence on the Long-Term Effects of Employment Training Programs. *Journal of Labor Economics* 10(4), 380–388.
- Dickens, W., L. Katz, K. Lang, and L. Summers (1989). Employee Crime and the Monitoring Puzzle. *Journal of Labor Economics* 7(3), 331–347.
- Farber, H. and R. Gibbons (1996). Learning and Wage Dynamics. *The Quarterly Journal of Economics* 111(4), 1007–1047.
- Gibbons, R. and L. F. Katz (1991). Layoffs and lemons. *Journal of Labor Economics* 9(4), 351–380.

- Greenwald, B. C. (1986). Adverse Selection in the Labour Market. *The Review of Economic Studies* 53(3), 325–347.
- Hollister Jr., R., P. Kemper, and R. Maynard (1984). *The National Supported Work Demonstration*. Madison, WI: University of Wisconsin Press.
- Kahn, L. (2012). Asymmetric Information between Employers. Working Paper.
- Loewenstein, M. and J. Spletzer (1998). Dividing the Costs and Returns to General Training. *Journal of Labor Economics* 16(1), 142–171.
- Redcross, C., D. Bloom, G. Azurdia, J. Zweig, and N. Pindus (2009). *Transitional Jobs for Ex-Prisoners: Implementation, Two-Year Impacts and Cost of the Center for Employment Opportunities (CEO) Prisoner Reentry Program*. New York, NY: MDRC.
- Schönberg, U. (2007). Testing for Asymmetric Employer Learning. *Journal of Labor Economics* 25(4), 651–691.
- Stanley, M., L. Katz, and A. Krueger (1998). Developing Skills: What We Know About the Impact of American Educational and Training Programs on Employment, Earnings and Educational Outcomes. *Report for G8 Economic Summit October*.
- Stanton, C. and C. Thomas (2012). Landing the First Job: the Value of Intermediaries in Online Hiring. Working Paper.
- Tervio, M. (2009). Superstars and Mediocrities: Market Failure in the Discovery of Talent. *Review of Economic Studies* 76(2), 829–850.
- Waldman, M. (1984). Job Assignments, Signalling, and Efficiency. *The RAND Journal of Economics* 15(2), 255–267.
- Weiss, A. (1980). Job Queues and Layoffs in Labor Markets with Flexible Wages. *Journal of Political Economy* 88(3), 526–538.

Figure 1. Example oDesk Profile

Josefina R. - "Web developer/ web design/ wordpress/ php/ SEO - Freelance Programmer / Developer, Argentina"

Permalink: <http://www.odesk.com/users/~~> **\$13.33/hr** [Contact](#)

Overview | [Résumé](#) | [Work History & Feedback \(4\)](#) | [Tests \(5\)](#) | [Portfolio \(1\)](#)

My goal in oDesk is to gain more experience in designing rich web applications using web standards, and improving content with SEO tools. I always worked as a web developer both back-end and front-end so I am opened to do interesting things and learn in the process.

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

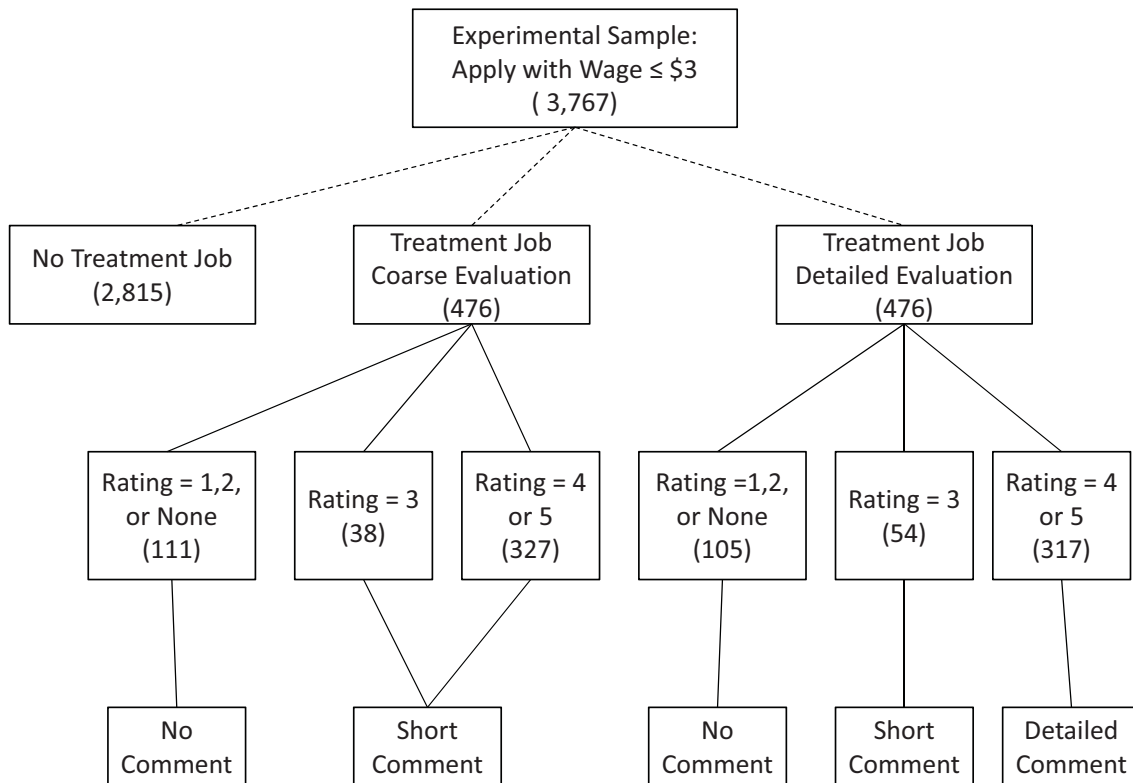
Employer ID	From/To	Job Title	Paid	Feedback
193850	05/2010 - 05/2010	Facebook Fan Page Developer - FBML	\$15 (1.33 hrs @ \$11.11/hr)	★★★★★ 5.0 Awesome provider! Helped me with a fanpage, thanks a lot! ? Contractor-to-Employer Feedback: ★★★★★ 5.0 ?
118136	04/2010 - 04/2010	PHP expert required for OsTicket	\$22 (fixed-price)	★★★★★ 5.0 Finished the job very fast and worked really hard. ? Contractor-to-Employer Feedback: ★★★★★ 5.0 ?
110612	03/2010 - 03/2010	Text integration	\$122 (11 hrs @ \$11.11/hr)	★★★★★ 5.0 ? Contractor-to-Employer Feedback: ★★★★★ 5.0 ?
168070	02/2010 - 02/2010	Form Submission Testing - Cross Browser	\$117 (10.00 hrs @ \$11.67/hr)	★★★★★ 5.0 Josefina is fantastical! I'm looking forward to working with her again! ? Contractor-to-Employer Feedback: ★★★★★ 5.0 ?

Statistics:
 Last 6 mos. Feedback: ★★★★★ (5.00) 3 feedbacks
 All-time Feedback: ★★★★★ (5.00) 4 feedbacks
 Hours: 1 (Last 6 mos), 22 (All-time)
 Contracts: 2 (Last 6 mos), 4 (All-time)
[See all Work History & Feedback](#)

Location: Buenos Aires, Argentina (UTC+02:00)
English Skills: (self-assessed) ██████████ 4.0
Member Since: January 29, 2010
Last Worked: May 9, 2010
oDesk Ready: [?](#) Yes

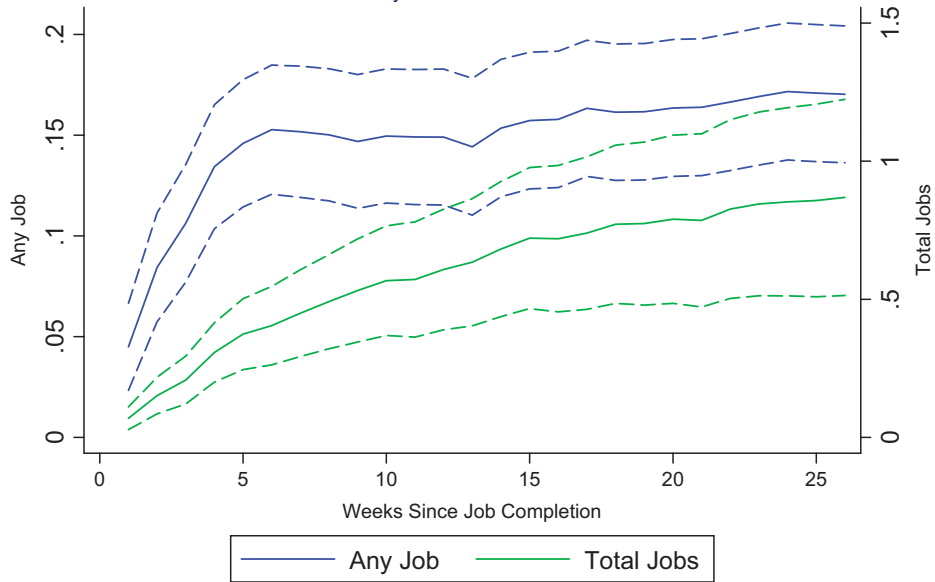
Related links:
 ■ Trends for [PHP Developers](#)
 ■ Trends for [Zend Developers](#)
 ■ Trends for [CSS Designers](#)
 ■ Trends for [Smarty Developers](#)
 ■ Trends for [Javascript Developers](#)

Figure 2. Experimental Design



Note: Dashed lines indicate random assignment.

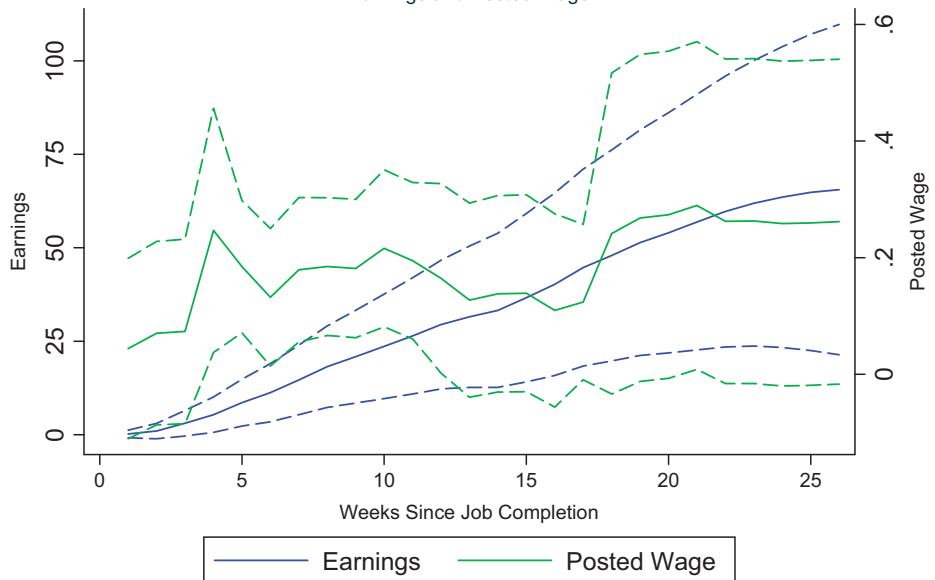
Figure 3a. Effect of Treatment Job by Week
Any Job and Total Jobs



Note: Dashed lines indicate 95% confidence intervals.

STATA™

Figure 3b. Effect of Treatment Job by Week
Earnings and Posted Wage



Note: Dashed lines indicate 95% confidence intervals.

STATA™

Table 1. Randomization Assessment: Treatment vs. Control Groups

	A. No Previous Job		B. With Previous Jobs		C. All Workers
	Treatment	Control	Treatment	Control	
	Group	Group	Group	Group	
Posted Wage	2.18	2.16	2.23	2.27	1.98
Days Since Joining oDesk	137	126	251	257	179
Number of Applications Sent	25*	22*	160	167	27
Proposed Wage for Treatment Job	2.18	2.16	2.29	2.32	2.01
Number of Tests Passed	2.7	2.7	4.5	4.7	3.5
Number of Qualifications	2.9	3.0	4.6	4.8	3.7
Percent with Portfolio	7%	6%	26%	25%	14%
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 indicated characteristic for the indicated group of workers immediately before the experiment. "Qualifications" are certifications from entities other than oDesk that are posted to the worker's profile. A "portfolio" is where a worker posts examples of her prior work. One asterisk indicates the difference between treatment and control group workers is significant at the 10% level and two asterisks indicate the difference is significant at the 5% level.

Table 2. Randomization Assessment: Detailed vs. Coarse Evaluation Treatment Groups

	A. All Treatment Group Workers		B. Workers with Ratings of 4 and 5	
	Detailed Treatment	Coarse Treatment	Detailed Treatment	Coarse Treatment
Posted Wage	2.17	2.22	2.21	2.25
Days Since Joining oDesk	161	164	163	164
Number of Applications Sent	58	53	55	53
Proposed Wage for Experiment Job	2.19	2.22	2.22	2.24
Number of Tests Passed	3.1	3.2	3.2	3.4
Number of Qualifications	3.2	3.4	3.4	3.6
Percent with Portfolio	12%	11%	13%	12%
Philippines	63%	62%	67%	67%
India	10%	10%	9%	8%
Bangladesh	12%	11%	11%	9%
Pakistan	6.1%	5.9%	4.7%	4.9%
United States	2.1%	3.4%	1.6%	2.8%
Fraction with Previous Job	23%	23%	25%	25%
Workers with Previous Jobs Only				
Number of Previous Jobs	7.3	6.9	6.2	6.7
Average Feedback Score	4.6**	4.2**	4.6*	4.3*
Observations	476	476	317	327

Notes: Each cell presents the mean value of the indicated characteristic for the indicated group of workers immediately before the experiment. "Qualifications" are certifications from entities other than oDesk that are posted to the worker's profile. A "portfolio" is where a worker posts examples of her prior work. One asterisk indicates the difference between workers in the two treatment groups is significant at the 10% level and two asteriks indicate the difference is significant at the 5% level.

Table 3. The Effects of the Treatments on Employment Outcomes
During the Two Months After the Experiment

	A. Workers with No Previous Jobs			B. Workers with Previous Jobs		
	Detailed Treatment	Coarse Treatment	Control	Detailed Treatment	Coarse Treatment	Control
Total Jobs	0.883	0.731	0.284	2.889	2.037	1.958
p-value: Equal to Control	(0.000)	(0.000)		(0.005)	(0.809)	
p-value: Equal to Coarse Treatment	(0.307)			(0.079)		
Any Job	0.302	0.296	0.117	0.694	0.528	0.545
p-value: Equal to Control	(0.000)	(0.000)		(0.003)	(0.729)	
p-value: Equal to Coarse Treatment	(0.872)			(0.012)		
Hours Worked	11.32	13.49	5.36	74.46	47.52	47.80
p-value: Equal to Control	(0.002)	(0.000)		(0.006)	(0.977)	
p-value: Equal to Coarse Treatment	(0.482)			(0.074)		
Posted Wage	2.31	2.25	2.03	2.68	2.32	2.38
p-value: Equal to Control	(0.000)	(0.000)		(0.010)	(0.565)	
p-value: Equal to Coarse Treatment	(0.573)			(0.043)		
Earnings	29.72	27.14	10.06	186.84	101.19	120.60
p-value: Equal to Control	(0.000)	(0.000)		(0.008)	(0.423)	
p-value: Equal to Coarse Treatment	(0.750)			(0.018)		
Observations	368	368	1,562	108	108	1,253

Notes: Each statistic not in parentheses is the mean of the indicated employment outcome for workers in the indicated experimental group. Employment outcomes are calculated for the two months after the experiment; all experimental jobs and earnings are excluded. Each statistic in parentheses is a p-value from a test that the means for the groups indicated by the row and column are equal.

Table 4. Regression Estimates of the Effects of the Treatments with Controls
During the Two Months After the Experiment

	Total Jobs (1)	Any Job (2)	Hours Worked (3)	Posted Wage (4)	Earnings (5)
<u>A. Treatments Separately</u>					
Detailed Treatment	0.617** (0.127)	0.165** (0.022)	11.36** (3.42)	0.26** (0.08)	29.37** (8.74)
Coarse Treatment	0.342** (0.102)	0.126** (0.022)	7.54** (2.98)	0.15** (0.04)	9.80 (6.76)
Previous Job	0.321** (0.104)	0.199** (0.019)	20.22** (3.23)	0.20** (0.04)	40.78** (7.65)
p-value: Detailed = Coarse	0.07	0.19	0.34	0.21	0.05
Control Group Mean	1.029	0.308	24.25	2.19	59.27
Observations	3,767	3,767	3,767	3,767	3,767
<u>B. Treatments Combined</u>					
Treatment Job	0.480** (0.087)	0.145** (0.016)	9.45** (2.50)	0.20** (0.05)	19.59** (6.03)
Previous Job	0.321** (0.104)	0.199** (0.019)	20.21** (3.23)	0.20** (0.04)	40.76** (7.66)
Control Group Mean	1.029	0.308	24.25	2.19	59.27
Observations	3,767	3,767	3,767	3,767	3,767

Notes: Each column in Panel A displays the results of regressing the dependent variable indicated by the column on indicators for being in the detailed evaluation treatment, being in the coarse evaluation treatment, and having prior oDesk employment. Each column in Panel B displays the results of regressing the same dependent variable on indicators for being in any treatment group and having prior oDesk experience. All regressions control for the tests the worker took (dummies for passing one test, two or three tests, and four or more tests, an indicator for having taken the most popular skills test among the sample, and an indicator for passing it), the number of qualifications the worker had (dummies for listing two to four qualifications and five or more qualifications), the number of oDesk applications the worker sent before the experiment (dummies for having sent three to five, six to 15, 16 to 50, 51 to 100, and over 100 applications), the wage proposed for the experimental job (dummies for offering \$1 to \$1.99, \$2 to \$2.99, and exactly \$3), the number of jobs the worker had before the experiment, and an indicator for being in the second experimental wave. Employment outcomes are calculated for the two months after the experiment; all experimental jobs and earnings are excluded. Huber-White standard errors are in parentheses. P-values in Panel A are from a test that the coefficients on the detailed evaluation treatment and coarse evaluation treatment indicators are equal. Two asterisks indicate the coefficient is significant at the 5% level.

Table 5. The Effects of Receiving a Detailed Evaluation for Workers Earning Ratings of 4 or 5 During the Two Months After the Experiment

	Total Jobs (1)	Any Job (2)	Hours Worked (3)	Posted Wage (4)	Earnings (5)
<u>A. Main Effect of Detailed Evaluation</u>					
Detailed Treatment	0.210 (0.215)	0.042 (0.039)	3.86 (5.85)	0.24* (0.13)	24.06* (14.07)
Coarse Evaluation Mean Observations	1.346 644	0.416 644	26.95 644	2.34 644	54.36 644
<u>B. Differential Effect of Detailed Evaluation</u>					
Met Deadline × Detailed Treatment	0.204 (0.219)	0.043 (0.040)	4.17 (5.93)	0.24* (0.13)	24.39* (14.27)
Missed Deadline × Detailed Treatment	-0.238 (0.388)	-0.095 (0.331)	-27.40** (13.98)	-0.28 (0.60)	-34.88** (15.51)
Met Deadline	0.791** (0.311)	-0.013 (0.190)	-0.47 (14.53)	0.06 (0.37)	19.90 (17.61)
p-value: Interaction Terms Equal	0.321	0.679	0.038	0.395	0.005
Coarse Evaluation Mean Observations	1.346 644	0.416 644	26.95 644	2.34 644	54.36 644

Notes: Each column in Panel A displays the results of regressing the dependent variable indicated by the column on an indicator for being in the detailed evaluation treatment. Each column in Panel B displays the results of regressing the same dependent variable on an indicator for meeting the deadline interacted with an indicator for being in the detailed evaluation treatment, an indicator for missing the deadline interacted with the detailed evaluation treatment dummy, and an indicator for meeting the deadline. Only workers who obtained a rating of at least four in an experimental job are included. Employment outcomes are calculated for the two months after the experiment; all experimental jobs and earnings 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. The "Coarse Evaluation Mean" rows present the mean of the dependent variables for workers in the coarse evaluation treatment who received a rating of at least four.

Table 6. Market-Level Effects of the Experiment

	<u>A. Log Jobs</u>				
% Change Experience x After	0.220** (0.108)	0.246** (0.105)	0.273** (0.091)	0.272** (0.092)	0.352* (0.201)
Mean of Dependent Variable	3.11	3.20	3.25	3.25	3.21
	<u>B. Log Hours</u>				
% Change Experience x After	0.142* (0.073)	0.153** (0.072)	0.099 (0.061)	0.102* (0.061)	0.107 (0.196)
Mean of Dependent Variable	6.45	6.55	6.63	6.63	6.59
	<u>C. Log Wages</u>				
% Change Experience x After	-0.418** (0.112)	-0.437** (0.097)	-0.265** (0.095)	-0.315** (0.096)	-0.034 (0.149)
Mean of Dependent Variable	1.86	1.86	1.85	1.85	1.94
Weeks After Experiment	8	17	26	26	26
Category-Specific Time Trends	Yes	Yes	Yes	Yes	Yes
Category-Specific Quadratic	No	No	No	Yes	Yes
Data-Entry Included	Yes	Yes	Yes	Yes	No

Notes: Panels A, B, and C present the results of regressing the natural logarithm of jobs created in a job category-week (Panel A), the natural logarithm of hours worked in a job category-week (Panel B), and the natural logarithm of a job's hourly wages (Panel C), on "% Change Experience" interacted with a dummy for being after the experiment. All regressions include week fixed effects, job category fixed effects, and job category linear time trends. The last two columns additionally include job category quadratic time trends. These regressions contain outcomes from 26 weeks before the experiment to 8, 17, or 26 weeks afterwards, excluding the weeks of the experiment. In Panels A and B, observations are job category-weeks and are weighted by the number of jobs created in the category in a pre-period. In Panel C, observations are individual job-weeks and are unweighted. All standard errors are clustered by job category. 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. Estimated Effect of the Experiment on Market Welfare

	High Benefit, Low Cost	Medium Benefit, Medium Cost	Low Benefit, High Cost
<u>A. Increased Market Surplus</u>			
Increased Worker Surplus (Excluding the Cost of Effort)	\$52,036	\$52,036	\$52,036
Increased Firm Profit (From Increased Employment)	\$2,757	\$2,757	\$2,757
<u>B. Opportunity Cost of Increased Work</u>			
Workers' Opportunity Cost (Alternative Hourly Wage)	\$12,351 (\$0.50/hour)	\$24,701 (\$1 per hour)	\$36,884 (\$1.49 per hour)
Firms' Opportunity Cost (Fraction of Increased Profits)	\$689 (25%)	\$1,378 (50%)	\$2,757 (100%)
<u>C. Cost of Experimental Jobs</u>			
Workers' Opportunity Cost (Alternative Hourly Wage)	\$3,622 (\$0.50/hour)	\$7,245 (\$1 per hour)	\$10,818 (\$1.49 per hour)
Fixed Cost of Employing (\$10 per hour spent)	\$476 3 min/worker	\$793 5 min/worker	\$1,587 10 min/worker
<u>D. Overall Welfare Change</u>			
Total Market Surplus - Cost	\$37,655	\$20,675	\$2,747

Note: Each cell in Panel D is the difference of the column's entry in Panel A and its entries in Panels B and C.

7 Appendix A: Proposition Proofs

7.1 Proof of Proposition 1

Any perfect Bayesian equilibrium must have a single-crossing property. That is, if a worker with expected ability \hat{a}_i^1 is hired, all workers with expected abilities $\hat{a}_i^2 > \hat{a}_i^1$ will also be hired. All workers weakly prefer working at any non-negative wage to their outside options and, at the same wage, firms strictly prefer to hire workers with higher expected abilities. Similarly, if a firm with fixed cost c_j^1 hires a worker, all firms with $c_j^2 < c_j^1$ also hire a worker.

Some agents accept their outside options: every firm prefers its outside option of zero to hiring a worker with negative expected ability at a non-negative wage. However, not all agents take their outside options: firms with c_j arbitrarily close to zero and workers with arbitrarily high expected abilities prefer employment relationships with each other to their outside options. Thus, there exists thresholds \hat{a}_w and c_f such that only workers with $\hat{a}_i \geq \hat{a}_w$ and firms with $c_j \leq c_f$ are in employment relationships. These inequalities are weak because of the assumption that indifferent agents enter the market.

Market wages must equal $w_{ij} = \hat{a}_i - b$ for some b . If $\hat{a}_i - w_{ij}$ were not constant for all hired workers, for some $\varepsilon > 0$, there would exist a hiring firm that would benefit by offering a wage of $w_{ij} + \varepsilon$ to a worker employed by another firm at a wage of w_{ij} . Because the marginal firm and worker must earn zero profit and wages, $c_f = b = \hat{a}_w \equiv \bar{c}$. This \bar{c} exists such that the market clears: as c increases on $[0, \alpha)$, the mass of firms with $c_j \leq c$ increases monotonically from 0 to 1, while the mass of workers with $\hat{a}_i \geq c$ decreases monotonically to 0.

There are no profitable deviations from this equilibrium. Hiring firms earn expected profit $\bar{c} - c_j$ from hiring any worker with $\hat{a}_i \geq \bar{c}$, which weakly exceeds their outside option. Firms with $c_j > \bar{c}$ would earn negative profit from paying market wages. All firms would earn lower profit from hiring workers with $\hat{a}_i < \bar{c}$ at non-negative wages. No workers with $\hat{a}_i \geq \bar{c}$ strictly prefer their outside option to market work. Veteran workers' wages weakly exceed their outside options. If novices work, they earn $\hat{a}_i - \bar{c}$ in their novice period which weakly exceeds their total maximum earnings over two periods if they take their outside options: $0 + \beta(\hat{a}_i - \bar{c})$.

7.2 Proof of Proposition 2

Let $l \in [0, 1]$ index workers and firms in each period such that only agents with $l \leq \bar{l}$ participate in the market. The utilitarian social planner maximizes

$$\sum_{t=0}^{\infty} \beta^t \int_0^{\bar{l}} (\hat{a}_l - c_l) dl$$

Here, $\bar{l} \in (0, 1)$. If $\bar{l} = 0$, having firms with arbitrarily low fixed costs hire novices with arbitrarily high expected abilities would increase social surplus. If $\bar{l} = 1$, having veterans with negative expected abilities take their outside options would increase expected surplus. There is a single-crossing property where if a firm with c_j^1 hires a worker, firms with $c_j^2 < c_j^1$ also hire a worker. Thus, there exists a c^* , such that all firms with $c_j \leq c^*$ and only those firms hire workers. (This relies on the assumption that if entering the market and taking their outside options generate equal social surplus, agents enter the market.)

The current-period expected social surplus from hiring the additional worker i is $\hat{a}_i - c^*$. Veterans will be hired if and only if $\hat{a}_{i1} \geq c^*$. The social planner will employ novices if and only if

$$\underbrace{(\hat{a}_{i0} - c^*)}_{\text{Novice-Period Surplus}} + \underbrace{\beta \Pr(\hat{a}_{i1H} \geq c^*) \times E[\hat{a}_{i1H} - c^* | \hat{a}_{i1H} \geq c^*]}_{\text{Veteran-Period Surplus}} \geq 0. \quad (1)$$

The social planner employs all novices with $\hat{a}_{i0} \geq c^*$. (Equation (1) holds for these workers.) If unemployed, these workers generate a maximum surplus of $0 + \beta(\hat{a}_{i0} - c^*)$ over their lives, which is less than the novice-period surplus generated from hiring them: $\hat{a}_{i0} - c^*$. If novices with $\hat{a}_{i0} < c^*$ are unemployed, they will have $\hat{a}_{i1} = \hat{a}_{i0} < c^*$ and be unemployed as veterans, generating zero social surplus over their lives. If employed as novices, the left-hand side of Equation (1) gives the lifetime surplus produced: they will be hired as veterans with probability $\Pr(\hat{a}_{i1H} \geq c^*)$, in which case, they will generate expected surplus $E[\hat{a}_{i1H} - c^* | \hat{a}_{i1H} \geq c^*]$.

A c^* exists that clears the market: as c increases on $[0, \infty)$, the mass of firms with $c_j \leq c$ increases monotonically from 0 to 1 while the mass of veterans with $\hat{a}_{i1} \geq c$ and the mass of novices for whom Equation (1) holds decrease monotonically to 0.

It must be that $c^* > \bar{c}$. For any hiring threshold c , more novices and veterans are hired in the social planner's solution than in the market equilibrium. (The social planner's solution employs some novices with $\hat{a}_{i0} < c$, which leads to more veterans with $\hat{a}_{i1} \geq c$.) Thus, if $c^* \leq \bar{c}$, the social planner's solution would have more employed workers than hiring firms. The social planner's solution must employ more novices than the market equilibrium. If it did not, it would employ fewer veterans than the market equilibrium. Both the social planner's

solution and the market equilibrium hire novices with $\hat{a}_{i0} \geq c^*$. A higher fraction of these workers will be hired as veterans in the market equilibrium, since $c^* > \bar{c}$. The remaining novices hired by the market have $\hat{a}_{i0} \geq \bar{c}$, so they have more than a 50% chance of being rehired as veterans, while the remaining novices hired by the social planner have $\hat{a}_{i1} < c^*$, so they have less than a 50% chance of being rehired as veterans. The social planner's solution cannot hire fewer veterans and no more novices than the market equilibrium while having more hiring firms ($c^* > c$).

7.3 Proof of Proposition 3

Relative to remaining in the market equilibrium, being hired by Employer C does not affect the expected veteran-period employment outcomes of novices with $\hat{a}_{i0} \geq \bar{c}$. These workers will be hired and receive an output signal with the same distribution whether Employer C hires them or they remain in the market. In contrast, if not hired by Employer C, novices with $\hat{a}_{i0} < \bar{c}$ will not be hired. Thus, they will have $\hat{a}_{i1} = \hat{a}_{i0} < \bar{c}$ and will be unemployed with no earnings and with reservation wages equal to their outside option of zero in their veteran periods. However, if hired by Employer C, some will receive sufficiently high output signals that their veteran-period expected abilities will exceed \bar{c} . They will be employed in their veteran periods with positive earnings and reservation wages.

Now, I consider the effect of being hired by Employer D relative to Employer C. The market uses Bayesian updating to determine its expectation of workers' abilities. The market's expectations of worker i 's ability if she is hired by Employer C (\hat{a}_{i1C}) and Employer D (\hat{a}_{i1D}), respectively, are:

$$\begin{aligned}\hat{a}_{i1C} &= \frac{\sigma_M^2 \hat{a}_{i0} + \sigma_{ai}^2 \hat{y}_{iM}}{\sigma_M^2 + \sigma_{ai}^2} \\ \hat{a}_{i1D} &= \frac{\sigma_D^2 \hat{a}_{i0} + \sigma_{ai}^2 \hat{y}_{iD}}{\sigma_D^2 + \sigma_{ai}^2}\end{aligned}$$

The terms \hat{y}_{iM} and \hat{y}_{iD} can be rewritten as

$$\begin{aligned}\hat{y}_{iM} &= a_i + \varepsilon_{iM} = (\hat{a}_{i0} - \varepsilon_{ia}) + \varepsilon_{iM} \\ \hat{y}_{iD} &= a_i + \varepsilon_{iM} = (\hat{a}_{i0} - \varepsilon_{ia}) + \varepsilon_{iD}\end{aligned}$$

Substituting these expressions into the above equations gives

$$\begin{aligned}
\hat{a}_{i1C} &= \frac{\sigma_M^2 \hat{a}_{i0} + \sigma_{ai}^2 (\hat{a}_{i0} - \varepsilon_{ia} + \varepsilon_{iM})}{\sigma_M^2 + \sigma_{ai}^2} \\
\hat{a}_{i1C} &= \hat{a}_{i0} + \frac{\sigma_{ai}^2 (-\varepsilon_{ia} + \varepsilon_{iM})}{\sigma_M^2 + \sigma_{ai}^2} \\
\hat{a}_{i1C} &\sim N \left(\hat{a}_{i0}, \frac{\sigma_{ai}^4}{\sigma_M^2 + \sigma_{ai}^2} \right) \\
\\
\hat{a}_{i1D} &= \frac{\sigma_D^2 \hat{a}_{i0} + \sigma_{ai}^2 (\hat{a}_{i0} - \varepsilon_{ia} + \varepsilon_{iD})}{\sigma_D^2 + \sigma_{ai}^2} \\
\hat{a}_{i1D} &= \hat{a}_{i0} + \frac{\sigma_{ai}^2 (-\varepsilon_{ia} + \varepsilon_{iD})}{\sigma_D^2 + \sigma_{ai}^2} \\
\hat{a}_{i1D} &\sim N \left(\hat{a}_{i0}, \frac{\sigma_{ai}^4}{\sigma_D^2 + \sigma_{ai}^2} \right)
\end{aligned}$$

Since $\sigma_D^2 < \sigma_M^2$, the distribution of \hat{a}_{i1D} is a mean-preserving spread of the distribution of \hat{a}_{i1C} . Worker's expected veteran-period earnings and reservation wages are $E[\max(\hat{a}_{i1} - \bar{c}, 0)]$, a convex function of \hat{a}_{i1} . Thus, by Jensen's inequality, $E[\max(\hat{a}_{i1D} - \bar{c}, 0)] > E[\max(\hat{a}_{i1C} - \bar{c}, 0)]$: being hired by Employer D relative to Employer C strictly increases workers' expected earnings and reservation wages.

A worker is employed in her veteran period with probability $\Pr[\hat{a}_{i1} \geq \bar{c} | \hat{a}_{i0}]$. The probabilities that a worker will be employed in her veteran period after being employed by Employers C and D, respectively are

$$\begin{aligned}
\Pr[\hat{a}_{i1C} \geq \bar{c} | \hat{a}_{i0}] &= \Phi \left(\frac{(\hat{a}_{i0} - \bar{c}) \sqrt{\sigma_M^2 + \sigma_{ai}^2}}{\sigma_{ai}^2} \right) \\
\Pr[\hat{a}_{i1D} \geq \bar{c} | \hat{a}_{i0}] &= \Phi \left(\frac{(\hat{a}_{i0} - \bar{c}) \sqrt{\sigma_D^2 + \sigma_{ai}^2}}{\sigma_{ai}^2} \right)
\end{aligned} \tag{2}$$

where Φ is the standard normal cumulative distribution function. Since $\sigma_D^2 < \sigma_M^2$, a worker with $\hat{a}_{i0} < \bar{c}$ has a higher probability of veteran-period employment after being hired by Employer D, while a worker with $\hat{a}_{i0} > \bar{c}$ has a higher probability of veteran-period employment after being hired by Employer C.

7.4 Proposition 5

Proposition 5 *Conditional on novice expected ability, \hat{a}_{i0} , the effect of being hired by Employer C (relative to remaining in the market equilibrium) on expected veteran-period employment, earnings, and reservation wages is weakly increasing in the market's initial uncertainty about the worker's ability, σ_{ai}^2 . The effect is strictly increasing in σ_{ai}^2 for workers with $\hat{a}_{i0} < \bar{c}$. However, conditional on \hat{a}_{i0} , the effect of being hired by Employer D (relative to being hired by Employer C) on expected veteran-period employment, earnings, and reservation wages is not monotonic in σ_{ai}^2 .*

Proof. First, I consider the effect of being hired by Employer C relative to remaining in the market equilibrium. Being hired by Employer C has no effect on expected employment outcomes for workers with $\hat{a}_{i0} \geq \bar{c}$. Thus, all that remains to prove is that the effect of being hired by Employer C is strictly increasing in σ_{ai}^2 for workers with $\hat{a}_{i0} < \bar{c}$. The probability that a worker with $\hat{a}_{i0} < \bar{c}$ who was hired by Employer C is hired in her veteran period is

$$\begin{aligned} \Pr(\hat{a}_{i1C} \geq \bar{c}) &= \Pr(\hat{y}_{iM} > \hat{a}_{i0}) \times \Pr(\hat{a}_{i1C} \geq \bar{c} | \hat{y}_{iM} > \hat{a}_{i0}) \\ &\quad + \Pr(\hat{y}_{iM} \leq \hat{a}_{i0}) \times \Pr(\hat{a}_{i1C} \geq \bar{c} | \hat{y}_{iM} \leq \hat{a}_{i0}) \\ \Pr(\hat{a}_{i1C} \geq \bar{c}) &= \Pr(\hat{y}_{iM} > \hat{a}_{i0}) \times \Pr(\hat{a}_{i1C} \geq \bar{c} | \hat{y}_{iM} > \hat{a}_{i0}) + 0. \end{aligned}$$

The second line follows because if $\hat{y}_{iM} \leq \hat{a}_{i0} < \bar{c}$, then $\hat{a}_{i1C} < \bar{c}$. When, $\hat{y}_{iM} > \hat{a}_{i0}$, \hat{a}_{i1C} is strictly increasing in σ_{ai}^2 . Since $\Pr(\hat{y}_{iM} > \hat{a}_{i0})$ does not depend on σ_{ai}^2 , $\Pr(\hat{a}_{i1C} \geq \bar{c})$ is strictly increasing in σ_{ai}^2 .

Similarly, workers' expected earnings and reservation wages are

$$E[1[\hat{a}_{i1C} \geq \bar{c}] \times (\hat{a}_{i1C} - \bar{c})]$$

where $1[\cdot]$ is an indicator for the expression being true. This can be rewritten as

$$\begin{aligned} &\Pr(\hat{y}_{iM} > \hat{a}_{i0}) \times E[1[\hat{a}_{i1C} \geq \bar{c}] \times (\hat{a}_{i1C} - \bar{c}) | \hat{y}_{iM} > \hat{a}_{i0}] \\ &\quad + \Pr(\hat{y}_{iM} \leq \hat{a}_{i0}) \times E[1[\hat{a}_{i1C} \geq \bar{c}] \times (\hat{a}_{i1C} - \bar{c}) | \hat{y}_{iM} \leq \hat{a}_{i0}] \\ = &\Pr(\hat{y}_{iM} > \hat{a}_{i0}) \times E[1[\hat{a}_{i1C} \geq \bar{c}] \times (\hat{a}_{i1C} - \bar{c}) | \hat{y}_{iM} > \hat{a}_{i0}] + 0. \end{aligned}$$

As above, the equality follows because when $\hat{y}_{iM} \leq \hat{a}_{i0}$, $\hat{a}_{i1C} < \bar{c}$. When $\hat{y}_{iM} > \hat{a}_{i0}$, \hat{a}_{i1C} is strictly increasing in σ_{ai}^2 . Since $\Pr(\hat{y}_{iM} > \hat{a}_{i0})$ does not depend on σ_{ai}^2 , the right-hand side of the equation is strictly increasing in σ_{ai}^2 .

Now I consider the effect of being hired by Employer D relative to being hired by Employer C on expected earnings and reservation wages. Given the lack of closed form solutions

to integrating the normal distribution, it is easiest to see the result by example. The distributions of \hat{a}_{i1C} and \hat{a}_{i1D} are given in Section 7.3. Let $\sigma_M^2 = 2$, $\sigma_D^2 = 1$, $\hat{a}_{i0} = 1$, and $\bar{c} = 2$. Then, expected earnings after being hired by Employers D and C are as follows for different levels of σ_{ai}^2 :

σ_{ai}^2	Earnings after D	Earnings after C	Earnings Difference: D - C
1	0.025	0.010	0.015
10	0.768	0.720	0.048
20	1.286	1.247	0.039

As σ_{ai}^2 increases from 1 to 10, the earnings and reservation wage benefit from being hired by Employer D relative to Employer C increases from 0.015 to 0.048, but it decreases to 0.039 as σ_{ai}^2 increases to 20. Thus, the earnings and reservation wage increase from being hired by Employer D relative to Employer C is not monotonic in σ_{ai}^2 .

Finally, I consider the effect of being hired by Employer D relative to Employer C on the expected probability of employment. Equation (2) provides the probabilities of employment after being hired by Employers C and D. For any \hat{a}_{i0}^1 , let \hat{a}_{i0}^2 be defined such that $\hat{a}_{i0}^1 - \bar{c} = \bar{c} - \hat{a}_{i0}^2$. Then, by the symmetry of the normal distribution

$$\Pr [\hat{a}_{i1D} \geq \bar{c} | \hat{a}_{i0}^1] - \Pr [\hat{a}_{i1C} \geq \bar{c} | \hat{a}_{i0}^1] = - (\Pr [\hat{a}_{i1D} \geq \bar{c} | \hat{a}_{i0}^2] - \Pr [\hat{a}_{i1C} \geq \bar{c} | \hat{a}_{i0}^2]).$$

Thus, if the effect of being hired by Employer D on employment is increasing in σ_{ai}^2 at $\hat{a}_{i0} = \hat{a}_{i0}^1$, then it is decreasing at $\hat{a}_{i0} = \hat{a}_{i0}^2$ and vice versa. ■

7.5 Proposition 6

Proposition 6 *Being hired by Employer D (relative to Employer C) weakly decreases workers' veteran-period probability of employment, earnings, and reservation wages when $\hat{y}_{iD} \leq \min(\hat{a}_{i0}, \hat{y}_{iM})$ and weakly increases these outcomes when $\hat{y}_{iD} \geq \max(\hat{a}_{i0}, \hat{y}_{iM})$.*

Proof. Since expected earnings, reservation wages, and employment probabilities are weakly increasing functions of \hat{a}_{i1} , all I need to show is that when $\hat{y}_{iD} \leq \min(\hat{a}_{i0}, \hat{y}_{iM})$, $\hat{a}_{i1D} \leq \hat{a}_{i1C}$ and when $\hat{y}_{iD} \geq \max(\hat{a}_{i0}, \hat{y}_{iM})$, $\hat{a}_{i1D} \geq \hat{a}_{i1C}$. The difference in expected abilities after having been hired by Employers D and C is:

$$\begin{aligned} \hat{a}_{i1D} - \hat{a}_{i1C} &= \frac{\sigma_D^2 \hat{a}_{i0} + \sigma_{ai}^2 \hat{y}_{iD}}{\sigma_D^2 + \sigma_{ai}^2} - \frac{\sigma_M^2 \hat{a}_{i0} + \sigma_{ai}^2 \hat{y}_{iM}}{\sigma_M^2 + \sigma_{ai}^2} \\ \hat{a}_{i1D} - \hat{a}_{i1C} &= \sigma_{ai}^2 \times \frac{[\sigma_D^2 (\hat{a}_{i0} - \hat{y}_{iM}) + \sigma_M^2 (\hat{y}_{iD} - \hat{a}_{i0}) + \sigma_{ai}^2 (\hat{y}_{iD} - \hat{y}_{iM})]}{(\sigma_D^2 + \sigma_{ai}^2)(\sigma_M^2 + \sigma_{ai}^2)}. \end{aligned}$$

Since variances are always positive, the sign of this expression will be determined by the sign of the terms in brackets. If $\hat{y}_{iD} \leq \min(\hat{a}_{i0}, \hat{y}_{iM})$

$$\begin{aligned} & \sigma_D^2(\hat{a}_{i0} - \hat{y}_{iM}) + \sigma_M^2(\hat{y}_{iD} - \hat{a}_{i0}) + \sigma_{ai}^2(\hat{y}_{iD} - \hat{y}_{iM}) \\ \leq & \sigma_D^2(\hat{a}_{i0} - \hat{y}_{iM}) + \sigma_M^2(\hat{y}_{iD} - \hat{a}_{i0}) + \sigma_{ai}^2(\hat{y}_{iD} - \hat{y}_{iM}) + \sigma_D^2(\hat{y}_{iM} - \hat{y}_{iD}) \\ = & (\sigma_M^2 - \sigma_D^2)(\hat{y}_{iD} - \hat{a}_{i0}) + \sigma_{ai}^2(\hat{y}_{iD} - \hat{y}_{iM}) \leq 0. \end{aligned}$$

So, $\hat{a}_{i1D} - \hat{a}_{iC} \leq 0$. If $\hat{y}_{iD} \geq \max(\hat{a}_{i0}, \hat{y}_{iM})$

$$\begin{aligned} & \sigma_D^2(\hat{a}_{i0} - \hat{y}_{iM}) + \sigma_M^2(\hat{y}_{iD} - \hat{a}_{i0}) + \sigma_{ai}^2(\hat{y}_{iD} - \hat{y}_{iM}) \\ \geq & \sigma_D^2(\hat{a}_{i0} - \hat{y}_{iM}) + \sigma_M^2(\hat{y}_{iD} - \hat{a}_{i0}) + \sigma_{ai}^2(\hat{y}_{iD} - \hat{y}_{iM}) + \sigma_D^2(\hat{y}_{iM} - \hat{y}_{iD}) \\ = & (\sigma_M^2 - \sigma_D^2)(\hat{y}_{iD} - \hat{a}_{i0}) + \sigma_{ai}^2(\hat{y}_{iD} - \hat{y}_{iM}) \geq 0. \end{aligned}$$

So, $\hat{a}_{i1D} - \hat{a}_{iC} \geq 0$. ■

7.6 Proposition 7

I extend the model as follows: before each period, a worker observes her period-specific outside option. With probability $\kappa > 0$, her outside option is $w_1 > 0$; with probability $1 - \kappa$, it is $w_0 = 0$. She then decides whether to accept her outside option. If she does, she exits the market. Otherwise, she remains in the market. The market clears after workers make their exit decisions. There is some randomness in hiring: after the market clears, with probability ε , each worker who has not received a job receives a wage offer of $\delta > 0$ such that $w_1 > \varepsilon\delta$. This randomness reflects the fact that some firms hire workers without reading workers' resumes. It affects the model only by ensuring that workers with outside option $w_0 = 0$ do not exit the market and thereby removes the multiplicity of equilibria.

Proposition 7 *Being hired by Employer C (relative to remaining in the market equilibrium) weakly increases all workers' probability of remaining in the market in their veteran periods and strictly increases this probability for workers with $\hat{a}_{i0} < \bar{c}$. Being hired by Employer D (relative to being hired by Employer C) increases this probability for workers with $\hat{a}_{i0} < w_1 + \bar{c}$ and decreases this probability for workers with $\hat{a}_{i0} > w_1 + \bar{c}$.*

Proof. No worker with $w_0 = 0$ will exit the market. If she exits, she receives 0, while if she remains in the market, her expected gain is at least $\varepsilon\delta$. Workers offered outside option w_1 will remain in the market if and only if their market earnings $(\hat{a}_{i1} - \bar{c})$ weakly exceed their outside option: $\hat{a}_{i1} \geq w_1 + \bar{c}$. Being hired by Employer C only affects novices with $\hat{a}_{i0} < \bar{c}$.

If these workers were not hired by Employer C, they would have $\hat{a}_{i1} = \hat{a}_{i0} < \bar{c} < w_1 + \bar{c}$ and would exit the market if offered outside option w_1 . On the other hand, if hired by Employer C, some would have sufficiently positive output signals that $\hat{a}_{i1C} > w_1 + \bar{c}$ and they would remain in the market even with outside option w_1 .

I now consider the effect of being hired by Employer D relative to Employer C. The probabilities that a worker remains in the market after being hired by Employers C and D are, respectively

$$\begin{aligned} \Pr[\hat{a}_{i1C} \geq w_1 + \bar{c} | \hat{a}_{i0}] &= \Phi\left(\frac{(\hat{a}_{i0} - w_1 - \bar{c})\sqrt{\sigma_M^2 + \sigma_{ai}^2}}{\sigma_{ai}^2}\right) \\ \Pr[\hat{a}_{i1D} \geq w_1 + \bar{c} | \hat{a}_{i0}] &= \Phi\left(\frac{(\hat{a}_{i0} - w_1 - \bar{c})\sqrt{\sigma_D^2 + \sigma_{ai}^2}}{\sigma_{ai}^2}\right) \end{aligned}$$

Since $\sigma_M^2 > \sigma_D^2$, when $\hat{a}_{i0} > w_1 + \bar{c}$, the probability of remaining in the market is higher after being hired by Employer C, while when $\hat{a}_{i0} < w_1 + \bar{c}$ the opposite is true. ■

7.7 Proof of Proposition 4

In the period after Employer C hires, there will be a larger mass of veterans with $\hat{a}_{i1} \geq \bar{c}$. Employer C directly hires novices with $\hat{a}_{i0} < \bar{c}$. Because Employer C removes novices with $\hat{a}_{i0} \geq \bar{c}$ from the market, the market hiring threshold adjusts from \bar{c} to $\tilde{c} < \bar{c}$ and novice workers with $\tilde{c} \leq \hat{a}_{i0} < \bar{c}$ are hired. Some of these novices will have $\hat{a}_{i1} \geq \bar{c}$. In the subsequent period, the mass of workers with $\hat{a}_i \geq \bar{c}$ will exceed the mass of firms with fixed costs $c_j \leq \bar{c}$, so the hiring threshold will increase to c' . Total employment increases with the mass of hiring firms. Wages fall, conditional on workers' expected abilities, from $w_{ij} = \hat{a}_i - \bar{c}$ to $w'_{ij} = \hat{a}_i - c'$.

To define the change in market surplus, I define $k \in [0, 1]$ which, in each period, indexes firms from lowest to highest hiring cost and workers from highest to lowest expected ability. That is, $k \equiv F_c(c_k) \equiv 1 - F_{\hat{a}}(\hat{a}_k)$ where $F_c(\cdot)$ and $F_{\hat{a}}(\cdot)$ are the cumulative distribution functions of c_j and \hat{a}_i , respectively. I define k_1 such that in the standard market equilibrium only workers and firms with $k \leq k_1$ are in employment relationships ($\hat{a}_{k_1} = c_{k_1} = \bar{c}$). In the period after Employer C hires, employment expands, so I define $k_2 > k_1$ such that agents with $k \leq k_2$ are in employment relationships in that period. Employer C's hiring also changes the distribution of \hat{a}_k to \hat{a}'_k . For all $c \geq \bar{c}$, Employer C increases the mass of workers with $\hat{a}_i \geq c$ in the subsequent period. That implies that $\hat{a}'_k > \hat{a}_k$ for all $k \leq k_1$.

The change in the subsequent period’s market surplus due to Employer C is

$$\begin{aligned}
& \underbrace{\int_0^{k_2} (\hat{a}'_k - c_k) dk}_{\text{After Employer C}} - \underbrace{\int_0^{k_1} (\hat{a}_k - c_k) dk}_{\text{No Employer C}} \\
= & \underbrace{\int_{k_1}^{k_2} (\hat{a}'_k - c_k) dk}_{\text{Increased Employment}} + \underbrace{\int_0^{k_1} (\hat{a}'_k - \hat{a}_k) dk}_{\text{Higher Expected Ability Workers Hired}} > 0.
\end{aligned}$$

Both terms in the second line are strictly positive. The first is the increased surplus due to employment expanding from k_1 to k_2 . It is positive because $\hat{a}'_k > c_k$ for all $k < k_2$. The second term is the increased surplus from higher expected ability workers displacing lower expected ability workers at firms with $k \leq k_1$. It is positive since $\hat{a}'_k > \hat{a}_k$ for $k \leq k_1$.

8 Appendix B

This appendix analyzes the treatments’ effects on workers’ application patterns. First, it assesses the treatments’ effects on the number and types of job applications workers sent. Then, it analyzes their effects on the probability that a given job application was successful.

Panel A of Appendix Table 4 displays the results of regressing (1) an indicator for sending at least one application in the two months after the experiment and (2) the number of applications sent in these two months on a dummy for being in either treatment group and an indicator for having prior oDesk experience. The table shows that the average treatment group worker was 24 percentage points more likely to send at least one application and sent 24 more applications than the average control group worker. Panel B displays the results of regressing the same dependent variables on a dummy for being in the detailed evaluation treatment. The regressions are limited to workers earning ratings of at least four in my treatment job. The results show that, relative to the coarse evaluation, the detailed evaluation did not affect the probability that a worker sent any application or the average number of applications she sent.

Appendix Table 5 analyzes the treatments’ effects on the types of jobs workers applied to. The first column of the table lists several objective features of jobs. Employers had to indicate whether each opening was an hourly (as opposed to fixed wage) job, its job category, and whether they had preferences for workers with certain credentials such as a given self-assessed English ability or level of oDesk feedback. I also observe the number of applicants to each job. Each characteristic is correlated with a job being more competitive. That is, when I consider the sample of applications experimental workers sent in the month

before the experiment and separately regress a dummy for an application’s success on each characteristic, controlling for worker fixed effects, each coefficient is statistically significantly negative.

Each cell in Panels A and B of Appendix Table 5 shows the results of a separate regression. Observations in these regressions are applications sent by experimental workers in the two months after the experiment. Panel A shows the results of regressing the indicated job characteristic on a dummy for being in either treatment group, controlling for a dummy for prior oDesk experience. It shows that being in a treatment group induced workers to apply to more-competitive jobs. Panel B shows the results of regressing the indicated job characteristic on a dummy for being in the detailed evaluation treatment, limiting the sample to applications sent by workers earning fours and fives in my treatment jobs. Similar to the results in the previous table, these regressions show that, relative to the coarse evaluation treatment, the detailed evaluation treatment did not change workers’ application patterns.

Finally, Appendix Table 6 evaluates the treatments’ effects on workers’ application success. Panels A and B consider the same sample of applications as the corresponding panels in the previous table. Panel A displays the result of regressing an indicator that an application was successful (multiplied by 100 for ease of viewing) on an indicator for being in one of the treatment groups and a dummy for having prior oDesk experience. The first column includes no worker or job controls. Because receiving a job affected both the set of workers who applied for jobs and the types of jobs workers applied to, Column 2 adds controls for the job characteristics listed in Appendix Table 5 and employer fixed effects and Column 3 additionally includes controls for worker characteristics and the wage the worker proposed for the job.¹⁹ Once these job and worker characteristics are included, receiving a treatment job is estimated to have increased the probability that a given worker’s job application was successful by approximately 10%. While this regression controls for the treatments’ effect on selection into the regression on observable characteristics, this estimate will be biased if the treatments change selection on unobservable characteristics as well. In particular, if the selection on unobservables biases the coefficient in the same direction as the selection on observables, the regression underestimates the true effect of obtaining a treatment job on application success.

Panel B shows the results of regressing the same dependent variable on an indicator that a worker was in the detailed evaluation treatment. Columns 2 and 3 add the same controls as in Panel A. The estimates suggest that the detailed evaluation had a positive effect on the probability that a worker obtained a given job she applied to, but it is very imprecise.

¹⁹These are dummies for requesting a wage less than \$1 per hour, exactly \$1, between \$1 and \$2, exactly \$2, between \$2 and \$3, exactly \$3, and more than \$3.

This is similar to the positive, imprecise effect the detailed evaluation had on employment.

Appendix 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%
In Treatment Group	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 who actually applied. The third row is the experimental sample: workers who applied for the job requesting a wage less than or equal to \$3. The final row counts the workers who were randomly selected to be in either treatment group. Unless otherwise indicated, percentages refer to the percentage of workers contacted.

Appendix Table 2. The Effects of Receiving a Detailed Evaluation with Controls
During the Two Months After the Experiment

	Total Jobs (1)	Any Job (2)	Hours Worked (3)	Posted Wage (4)	Earnings (5)
<u>A. Main Effect of Detailed Evaluation</u>					
Detailed Treatment	0.219 (0.199)	0.040 (0.037)	3.09 (5.29)	0.20 (0.12)	22.66* (12.94)
Coarse Evaluation Mean Observations	1.346 644	0.416 644	26.95 644	2.34 644	54.36 644
<u>B. Differential Effect of Detailed Evaluation</u>					
Met Deadline × Detailed Treatment	0.205 (0.201)	0.040 (0.037)	3.20 (5.37)	0.21* (0.12)	22.63* (13.12)
Missed Deadline × Detailed Treatment	-0.129 (0.638)	-0.076 (0.302)	-25.67 (25.59)	-0.47 (0.38)	-37.39 (55.67)
Met Deadline	1.556** (0.624)	0.092 (0.178)	15.55 (22.68)	-0.14 (0.30)	55.92 (46.51)
p-value: Interaction Terms Equal	0.621	0.703	0.038	0.395	0.005
Coarse Evaluation Mean Observations	1.346 644	0.416 644	26.95 644	2.34 644	54.36 644

Notes: Each column in Panel A displays the results of regressing the dependent variable indicated by the column on an indicator for being in the detailed evaluation treatment. Each column in Panel B displays the results of regressing the same dependent variable on an indicator for meeting the deadline interacted with an indicator for being in the detailed evaluation treatment, an indicator for missing the deadline interacted with the detailed evaluation treatment dummy, and an indicator for meeting the deadline. Only workers who obtained a rating of at least four in an experimental job are included. All regressions control for the tests the worker took (dummies for passing one test, two or three tests, and four or more tests, an indicator for having taken the most popular skills test among the sample, and an indicator for passing it), the number of qualifications the worker had (dummies for listing two to four qualifications and five or more qualifications), the number of oDesk applications the worker sent before the experiment (dummies for having sent three to five, six to 15, 16 to 50, 51 to 100, and over 100 applications), the wage proposed for the experimental job (dummies for offering \$1 to \$1.99, \$2 to \$2.99, and exactly \$3), the number of jobs the worker had before the experiment, and an indicator for being in the second experimental wave. Employment outcomes are calculated for the two months after the experiment; all experimental jobs and earnings 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. The "Coarse Evaluation Mean" rows present the mean of the dependent variables for workers in the coarse evaluation treatment who received a rating of at least four.

Appendix Table 3. Differential Effects of Detailed Evaluations: 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 Treatment	0.345 (0.250)	0.064 (0.045)	2.60 (6.25)	0.36** (0.16)	26.48* (15.10)
Not Follow All Instructions × Detailed Treatment	-0.234 (0.415)	-0.032 (0.080)	8.20 (14.49)	-0.17 (0.20)	15.82 (34.57)
Follow All Instructions	0.022 (0.316)	0.024 (0.065)	4.43 (9.77)	-0.31* (0.18)	-6.95 (21.95)
p-value: Interaction Terms Equal	0.232	0.293	0.723	0.040	0.777
<u>B. Speed</u>					
Top Two Thirds × Detailed Treatment	0.183 (0.241)	0.029 (0.046)	-1.76 (6.80)	0.16 (0.16)	14.52 (16.02)
Bottom Third × Detailed Treatment	0.288 (0.457)	0.071 (0.074)	18.47 (11.89)	0.46** (0.23)	49.34* (29.78)
Top Two Thirds	-0.017 (0.340)	0.067 (0.059)	12.09 (7.14)	0.047 (0.13)	16.80 (16.02)
p-value: Interaction Terms Equal	0.838	0.630	0.140	0.291	0.304
<u>C. Accuracy</u>					
Top Two Thirds × Detailed Treatment	0.245 (0.238)	0.032 (0.046)	5.89 (6.91)	0.12 (0.10)	29.93* (16.94)
Bottom Third × Detailed Treatment	0.111 (0.467)	0.067 (0.074)	-1.44 (11.14)	0.54 (0.40)	8.97 (25.51)
Top Two Thirds	-0.274 (0.360)	0.022 (0.061)	-5.73 (10.16)	-0.02 (0.11)	-9.95 (22.29)
p-value: Interaction Terms Equal	0.798	0.689	0.577	0.312	0.494
Coarse Evaluation Mean	1.346	0.416	26.95	2.34	54.36
Observations	644	644	644	644	644

Notes: Each column in each panel presents the results of a separate regression of the dependent variable indicated by the column on the variables listed in the first column. Only workers who obtained a rating of at least four in an experimental job are included. Employment outcomes are calculated for the two months after the experiment; all experimental jobs and earnings 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. The "Coarse Evaluation Mean" row presents the mean of the dependent variables for workers in the coarse evaluation treatment who received a rating of at least four.

Appendix Table 4. The Effects of the Treatments on Application Patterns
During the Two Months After the Experiment

	Sent Any Application (1)	Applications Sent (2)
<u>A. Treatment Job</u>		
Treatment Job	0.238** (0.014)	23.96** (2.18)
Previous Job	0.337** (0.013)	30.37** (1.89)
Control Group Mean	0.703	25.90
Observations	3,767	3,767
<u>B. Detailed Evaluation</u>		
Detailed Treatment	-0.007 (0.016)	0.57 (5.30)
Coarse Evaluation Mean	0.963	53.12
Observations	644	644

Notes: Panel A presents the results from regressing an indicator for whether the worker sent at least one application in the two months after the experiment and the number of applications she sent on indicators for receiving a treatment job and having prior oDesk experience. Panel B presents the results of regressing the same dependent variables on an indicator for being in the detailed evaluation treatment. The regressions in Panel B only include workers who obtained a rating of at least four in an experimental job. Huber-White standard errors are in parentheses. No applications to experimental jobs are included in any regression. Two asterisks indicate the coefficient is significant at the 5% level. The "Control Group Mean" row presents the mean of the dependent variables for the entire control group; the "Coarse Evaluation Mean" row presents the mean of the dependent variables for workers in the coarse evaluation treatment who received a rating of at least four.

Appendix Table 5. The Effects of the Treatment on the Types of Jobs Applied to
During the Two Months After the Experiment

	A. Independent Variable = Indicator for Either Treatment Group	B. Independent Variable = Indicator for Detailed Evaluation Treatment	C. Mean of Dependent Variable for Control Group
Hourly Job	0.015 (0.011)	-0.005 (0.018)	0.754
Number of Applicants	14.670** (5.115)	-2.153 (9.342)	146
Data Entry	0.087** (0.019)	0.008 (0.033)	0.428
Preference for English Ability	0.020** (0.005)	-0.003 (0.007)	0.181
Preference for Number of oDesk Hours	0.014** (0.004)	-0.003 (0.005)	0.123
Preference for Level of oDesk Feedback	0.018** (0.004)	-0.002 (0.006)	0.146
Preference for Maximum Wage below \$5	0.014** (0.003)	-0.005 (0.004)	0.094
Preference for Minimum Wage above \$3	0.008** (0.002)	0.000 (0.003)	0.058
Observations	114,082	34,389	72,919

Notes: Each cell in Panels A and B presents the results from a separate regression of the variable in the first column on an indicator for being in either treatment group (Panel A) or an indicator for being in the detailed evaluation treatment (Panel B). Observations are applications sent in the two months after the experiment; applications to experimental jobs are excluded. The regressions in Panel A include applications sent by all experimental workers and control for whether workers have prior oDesk experience. The regressions in Panel B only include applications sent by workers who obtained a rating of at least four in an experimental job and add no controls. Standard errors are clustered by worker. Two asterisks indicate the coefficient is significant at the 5% level. Each cell in Panel C provides the mean of the variable in the first column for the control group.

Appendix Table 6. The Effects of the Treatments on Application Success
 Dependent Variable: Indicator that an Application was Successful $\times 100$

	(1)	(2)	(3)
	<u>A. Treatment Job</u>		
Treatment Job	-0.330 (0.244)	0.164 (0.187)	0.361** (0.178)
Previous Job	2.026** (0.236)	0.900** (0.180)	0.774** (0.202)
Control Group Mean	4.04	4.04	4.04
Observations	114,082	114,082	114,082
	<u>B. Detailed Evaluation</u>		
Detailed Treatment	0.295 (0.339)	0.312 (0.245)	0.227 (0.220)
Coarse Evaluation Mean	2.63	2.63	2.63
Observations	34,389	34,389	34,389
Job Characteristics	No	Yes	Yes
Employer Fixed Effects	No	Yes	Yes
Worker Characteristics	No	No	Yes

Notes: Each column in Panel A presents the results of regressing an indicator for a given job application being successful multiplied by 100 on indicators for the worker being in any treatment group and having prior oDesk experience. Each column in Panel B presents the results of regressing the same dependent variable on an indicator for the worker being in the detailed evaluation treatment. Observations are applications sent in the two months after the experiment; applications to experimental jobs are excluded. The regressions in Panel B only include applications sent by workers who obtained a rating of at least four in an experimental job. The regressions in Column 1 include no controls. The regressions in Column 2 control for the variables in the first column of Appendix Table 5 and employer fixed effects. The regressions in Column 3 add worker characteristics measured before the experiment: control for the tests the worker took (dummies for passing one test, two or three tests, and four or more tests, an indicator for having taken the most popular skills test among the sample, and an indicator for passing it), the number of qualifications the worker had (dummies for listing two to four qualifications and five or more qualifications), the number of oDesk applications the worker sent before the experiment (dummies for having sent three to five, six to 15, 16 to 50, 51 to 100, and over 100 applications), the wage proposed for the experimental job (dummies for offering \$1 to \$1.99, \$2 to \$2.99, and exactly \$3), the number of jobs the worker had before the experiment, and an indicator for being in the second experimental wave. Column 3 also controls for the wage the worker proposed for the current job (dummies for offering \$1 to \$1.99, \$2 to \$2.99, exactly \$3, and more than \$3). Standard errors are clustered by worker. Two asterisks indicate the coefficient is significant at the 5% level. The "Control Group Mean" row presents the mean of the dependent variable for the entire control group; the "Coarse Evaluation Mean" row presents the mean of the dependent variable for workers in the coarse evaluation treatment who received a rating of at least four.