Title
Essays on Platform Markets

Author
Barach, Moshe Adiel

Publication Date
2016

Peer reviewed|Thesis/dissertation

Essays on Platform Markets

by

Moshe Adiel Barach

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor John Morgan, Co-chair
Professor Steven Tadelis, Co-chair
Professor Benjamin Handel
Professor Noam Yuchtman

Spring 2016

Essays on Platform Markets

Abstract

Essays on Platform Markets

by

Moshe Adiel Barach

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor John Morgan, Co-chair

Professor Steven Tadelis, Co-chair

This dissertation is comprised of three studies investigating how the provision of public information facilitates matching in online marketplaces.

The first study examines how employers on oDesk.com, the world's largest online marketplace, use public information in hiring. I show that when employers are searching for someone low-skilled the provision of coarse information is sufficient. When employers are looking for someone high-skilled they will pay fixed screening costs to acquire information beyond what is provided by the platform. When information is not provided by the marketplace all employers will pay to acquire more information. This leads to more matches and hiring quality workers at a lower price. However, the cost savings from hiring these low-cost, but high-quality workers does not outweigh the upfront cost of information acquisition.

The second study investigates how employers on oDesk.com alter their hiring behavior when some applicants, who were selected to be of higher quality by a machine learning algorithm, are covered by a money-back guarantee (MBG). The MBG increases the probability of transacting for employers looking for high-expertise talent. However, conditional on hiring, these employers do not substitute to observably different employees. Employers that are looking for lower expertise applicants substitute to higher quality applicants who are backed by the money-back guarantee. A follow-up experiment found that, given some reasonable assumptions, employers found the MBG to be more useful as a signal of applicant quality than as a risk-shifting mechanism.

The third study uses a series of empirical tests to isolate the mechanism through which the weather affects bidders' willingness to pay (WTP) on eBay. I find that some variation in bidders' WTP across multiple auctions can be explained by weather affecting mood. Changes in mood can alter WTP through two channels. Mood can directly alter the perceived value of the item or mood can influence cognitive processes, changing bidding strategies and indirectly changing the bid amount. I concluded that rain alters mood, which alters bidders' perceived value of an item on eBay leading to bidders increasing their bids.

For My Family

Contents

List of Figures

List of Tables

Acknowledgments

I begin by acknowledging those whose advice, time, patience, and trust helped set me on my way towards a doctorate. Guy Genin, you were the first professor to introduce me to research, and placed me on the path towards this degree. Dan Elfenbein, Lamar Pierce, Todd Zenger and my other professors at Washington University in St. Louis, you suggested I pursue a doctorate in business when I, like most, didn't know such a degree existed.

Choosing a dissertation advisor did not come easy for me. My interests are broad and my intellectual curiosity seldom satiated. As such, I could not limit myself to only one dissertation chair. John Morgan, the first time I sat in your office and watched you take notes on our conversation, you told me that if I wanted to work with you, we would do things the "John Morgan" way. Thank you for having patience to let me find my own interests, and the foresight to see the researcher in me. I will always be a faithful student of the "John Morgan" way. Steve Tadelis, you taught me to search for good questions not good data, and always knew when I needed a good kick in the butt. Your advice was never easy to hear, but often exactly what I needed to hear. I will always seek to replicate your pursuit of exciting questions.

During my time at Berkeley other faculty members provided me with support and advice when I needed it most. Thank you to Ned Augenblick, Paul Gertler, Ben Handel, Ming Leung, Toby Stuart, and Noam Yuchtman. Noam, your name might always appear last according to alphabetical order, but your contributions to my success as an academic are first order. Your enthusiasm and thoughts can't be contained on a notepad and left me with no choice but to record our conversations. Any dissertation by a Business and Public Policy student is not complete without acknowledging his or her colleges. You were the reasons I came to Berkeley, and the reasons I made it through Berkeley.

My family, you are and always will be the most important thing in my life. Abba, you have influenced me more than any other person. The older I get, the wiser I realize you are. You worked relentlessly to make my life easier, to provide me with all the opportunities and advantages that built towards this moment. You deserve recognition in everything I accomplish. Imma, you are the kindest, most warmhearted person I know. You are hardworking and smart, and I am proud to be so much like you. Shoshi, you have always been an inspiration to me. You follow your own path; dressing like a pirate and looking good isn't easy. Ilana, I am so proud of you. I would have never finished this dissertation if your indecisiveness did not keep me in the office late on Fridays. Sheppy, you are my best friend. Thank you for showing me the way by writing a thesis and becoming a university lecturer first. Next, I may even ride my bike to work. Jackson the Dog, you are the youngest and hairiest of the Barachs. Thank you for taking Abba for walks.

The day I began writing this dissertation was the first day of my academic career. On that day I also met my future wife, Dr. Sydnie Lieb, so it was also the first day of our life together. And a very good day it was.

I love you, Girl. You make me the happiest Boy I can be.

Chapter 1

Introduction

Throughout history, there have existed businesses that create value by enabling two or more customers to find each other and exchange value. Such businesses were called "two-sided markets" in the seminal paper by Jean-Charles Rochet and Jean Tirole. These markets are at their core businesses that create value by acting as intermediaries. They create and provide information, act as third-party certifiers and facilitate monetary exchanges.

In the first two chapters of this dissertation, I study how the platform market itself affects matching by controlling the provision of information in the marketplace. In the third and final chapter, I demonstrate that variations in a consumer's mood, as influenced by the weather, affects economic decision-making. This finding has important implications for advertising, marketing, and matching on the platform.

In the first chapter, I study a randomized experiment that I designed and ran on oDesk.com, the worlds largest online labor marketplace, and one of a number of two-sided markets that are playing an increasing role in the hiring process. In this experiment, I altered the information set available to a randomly selected group of employers by hiding all the previous hourly wage rates an applicant had earned previously on the platform. Past wage rates are particularly useful in the screening process, as they provide a snapshot estimate of the worker's relative marginal product. No matter the competitive market environment, a worker's marginal product must be at least his wage, or no employer would pay that wage. A worker's current or past employer may be better informed about a worker's ability than the overall marketplace, and past wages provide at least a glimpse of what others think a worker's marginal product might be.

In conventional labor markets, employers do not directly have access to an applicant's past wages unless the applicant chooses to disclose this information. In this paper, I consider an online marketplace where the status quo is, and always has been, for employers to know the complete work history, including wages and earnings, for all applicants to their job. I observe how altering the publicly available set of information by removing applicants' past wage rates changed the employers' screening strategy and the result this had on match formation, wage negotiation, and job outcomes. When employers do not have the ability to observe an applicant's wage history, they substitute for this informational loss by exerting costly effort to acquire more information about candidates though interviewing. This strategic reaction to a reduced information environment actually leads employers to be more likely to fill an open job posting and to hire cheaper candidates. This treatment effect is limited to employers looking to hire low-expertise workers.

In the second chapter, I describe another experiment also run on oDesk.com. In this experiment, I offered a money-back guarantee to some randomly chosen employers. This guarantee only covered workers who were calculated to be of high quality using a machine learning algorithm. The algorithm was designed to predict contractors who were a "best match" for that employer and that posted job.

I find that this guarantee causes employers who are looking for high-quality applicants to be 3 percentage points more likely to locate and hire an applicant from a baseline of 40% fill rate. However, conditional on filling a job, the money-back guarantee does not influence these employers to substitute to candidates who are any different on observable characteristics. The introduction of the money-back guarantee causes employers that are looking at candidates who are just below the quality threshold to substitute to candidates that are just above the quality threshold and are covered by the money-back guarantee. This leads to a 3% increase in feedback left by employers that substitute to candidates that are just above the threshold, but no observable change in feedback left by employers who specifically targeted higher expertise applicants. While the employers looking to hire expert talent do not report any higher satisfaction, the increase along the extensive margin leads to a 3 percentage point increase in rehiring from a baseline of 14%.

This pattern of results is consistent with the money-back guarantee acting as a de facto certification mechanism that helps the employer to identify top talent, and the money-back guarantee acting as a risk-shifting mechanism that lowers the risk involved in hiring for employers by shifting some risk to the platform. To disentangle these two effects, I ran a follow-up experiment comparing the effects of marking applicants not as "guaranteed," but as "recommended." Guaranteeing applicants as opposed to simply signaling their quality by marking them as "recommended" does not have any effect on either the probability of filling a job posting or the likelihood of employers looking at candidates who are just below the quality threshold to substitute to candidates who are just above the quality threshold. This suggests that a large fraction of the MBG in the setting I study is really an endorsement effect, not the result of risk sharing. This work, along with the work from chapter one of this dissertation, suggests that information possessed by the market platform is valued by employers and shapes the matching process.

The third chapter uses a series of empirical tests to study the effects of the weather on eBay bidding behavior. These empirical tests are meant to help tease out the mechanism through which weather and possibly mood affect bidders' willingness to pay (WTP) for items on eBay. I show that rain increases the WTP for consumer goods by about 1% on eBay. The products I analyze can be separated into two general categories. Consumer goods, which can generally be considered private value goods, and cash-like substitutes, which can be considered public value goods. This differentiation is important as it allows me to unpack the mechanism through which mood affects willingness to pay. The psychology literature has highlighed that mood can directly affect an individual's perceived value of an item. The literature has also detailed that mood influences cognitive processes, changing bidding strategies and indirectly changing the amount bid. If mood were to affect cognitive processes more generally, it seems likely that the effects would also be seen on cash-like substitutes as well on consumer goods. This is not the case. Finally, to help alleviate concerns that the effects of weather on bidding are not due to mood at all, but instead to changes in opportunity costs of bidding, I show that weather has no effect of the timing and number of bids submitted in an auction. Taken all together, the evidence presented seems to suggest that weather does have a significant yet small effect on bidding behavior on eBay. It further suggests that the mechanism through which weather affects bidding behavior on eBay is by altering mood, which alters a bidder's perceived valuation of an item.

Chapter 2

Search, Screening, and Information Provision

2.1 Introduction

Identifying high-quality workers is one of the most important problems a firm faces. According to a survey by Silicon Valley Bank, 90% of startups believe finding talent is their biggest challenge.[1] In fact, talent acquisition is so important to the strategy of firms like Facebook, that its CEO, Mark Zuckerberg, once remarked that, "Facebook has not once bought a company for the company itself. We buy companies to get excellent people."

But, how do firms determine which potential employees are indeed "excellent people?" The answer has to do with what information about job applicants is available and how costly it is to acquire more information (Barney, 1986). This paper studies the effect of publicly provided information in a marketplace on search, screening, and match outcomes. In every marketplace in which agents search for a match, there is some choice to engage in costly private search. Private search does not occur in a vacuum. It occurs in the presence of public information, which is cheap to acquire. How much effort is exerted in costly private search must depend on the provision of public information that exists in the marketplace.

It is generally assumed that the more information publicly provided within a market, the better the matches should be. This paper, however, will question these assumptions by pointing out that agents strategically adjust their use of costly search to the provision of cheap-to-observe information in the market. Thus, the equilibrium amount of information an employer has when hiring an applicant may not be monotonically decreasing in the cost of obtaining match quality information from the marketplace. This is because the equilibrium level of information depends on both market provided information and information acquired via costly search and screening.

More specifically, my focus will be on the impact of worker wage rate history in online

[1] `http://www.svb.com/News/Company-News/Looking-For-a-Job--Try-a-Tech-Startup/`

labor markets.[2] I chose to study these online markets for two reasons: first, their importance and prevalence is growing.[3] Second, the particularities of these markets have allowed me to run an experiment that would be impossible in a traditional market. Identifying the causal impact of publicly provided information in a marketplace is problematic because, in general, the provision of this information is an endogenous choice of the marketplace. In oDesk.com, the online labor market that was studied, I was able to completely randomize which employers observed or did not observe coarse platform provided information. Furthermore, I was able to utilize the market's system of messages and reviews to delve into precisely what sorts of costly information acquisition employers chose to engage in as well as how satisfied they were with their hiring choices after the fact.

Traditionally, firms have relied on resumes and costly interviews to ascertain if a worker is a good fit for the firm. Increasingly, firms are turning to "big data" and workplace analytics to cheaply generate insights into candidates' potential productivity.[4] However, technology that provides agents with cheap but noisy signals reduces incentives to acquire more precise information and consequently might lead to worse decisions. Online labor marketplaces such as oDesk.com are playing an increasing role in the hiring process. Online labor marketplaces, even more so than the traditional labor market, seek to harness the power of big data and algorithms to foster efficient matching.[5] These marketplaces are in the position to provide large quantities of standardized and verified information to facilitate hiring. For example, oDesk.com does not allow workers to delete ratings or comments provided by employers after a job is completed. This information is distinct from what workers include in their resumes and thus valuable to employers. Despite the general view that more information is better when it comes to bilateral matching with asymmetric information, some research has hypothesized limitations to the social gains of the digitization of the labor market. Autor (2001b), proposes that information about workers' attributes can be usefully grouped into "low bandwidth" and "high bandwidth" varieties.[6] Low bandwidth data are objectively verifiable information such as education, credentials, experience, and salaries – the kind of information that is easily provided for employers through online marketplaces. High bandwidth data are attributes such as quality, motivation, and "fit," which are typically

[2]The major players in this industry are oDesk.com, Elance.com, Freelancer.com, and Thumbtack.com. In December of 2013, oDesk and Elance merged to form oDesk-Elance.com, but maintained separate online platforms until May 2015 when the company changed its name to UpWork.com and released one unified platform. My experiment was run post-merger, but solely on the oDesk.com platform.

[3]The online market for labor is already an economically important source of labor supply and demand, and it is estimated to grow substantially in the coming years (Agrawal, Horton, et al., 2013). *The Economist* (2013) predicted that online labor markets will be worth more than $5 billion by 2018.

[4]`https://hbr.org/resources/pdfs/comm/workday/workday_report_oct.pdf`

[5]Horton (2013) and Agrawal, Lacetera, et al. (2013) looks at the effects of algorithmically recruiting applicants to job openings for employers, and Agrawal, Lacetera, et al. (2013) show that verified work experience disproportionally helps contractors from less developed countries.

[6]Rees (1966) makes a similar distinction between formal and informal information networks. Formal networks include state employment services, private fee-charging employment agencies, newspaper advertisements, union hiring halls, and school or college placement bureaus. Informal sources include referrals from employees, other employers, miscellaneous sources, and walk-ins or hiring at the gate.

hard to verify except through direct interactions such as interviews and repeated contact. Hence, the comparatively low cost of obtaining applicants low-bandwidth information might reduce incentives of employers to ascertain applicants' high-bandwidth information.

The experiment I designed, which is described in detail in section 2.3, altered the information set available to a randomly selected group of employers by hiding all the previous hourly wage rates an applicant had earned on the platform. Past wage rates are particularly useful in the screening process, as they provide a snapshot estimate of the worker's relative marginal product. No matter the competitive market environment, a worker's marginal product must be at least his wage, or no employer would pay that wage. A worker's current or past employer may be better informed about a worker's ability than the overall marketplace, and past wages provide at least a glimpse of what others think a worker's marginal product might be. In conventional labor markets, employers do not directly have access to an applicant's past wages unless the applicant chooses to disclose them. In this paper, I consider an online marketplace where the status quo is, and always has been, for employers to know the complete work history, including wages and earnings, for all applicants to their job. I observe how altering the publicly available set of low-bandwidth information by removing applicants' past wage rates changed the employers' screening strategy and the result this had on match formation, wage negotiation, and job outcomes.

This paper has five main findings. First, I find that when employers are unable to observe an applicant's past wage rates they increase their search by viewing more applicants and increase screening by interviewing more applicants more deeply. Employers strategically react to having less information by actively acquiring information through a more costly, but perhaps more precise, source by taking time to ask the applicant more questions and get more information about his quality.

Second, I find that when employers are unable to observe applicants' past wage rates employers are more likely to fill a job posting. This may seem surprising because standard economic theory predicts that removing information from an employer causes the employer to be less able to differentiate between candidates based on quality. However, if there is a fixed cost associated with acquiring information, then a reduction in the exogenous information provision does not necessitate a reduction in the equilibrium level of information. This is because employers substitute for the reduction in market provided information with even more information from costly interviewing.

Third, I show that employers who are unable to observe past wage rates interview and hire cheaper but not observably worse candidates, and are more satisfied with the work output. This selection effect on wages completely dominates a small positive negotiation effect for workers. Hired applicants earn a higher percentage of their original bid.

Fourth, I provide evidence that although employers hire cheaper workers and are ex post more satisfied with the work output when they cannot observe past wage rates, employers are not making a mistake by using platform provided information when available. Using a difference-in-difference approach I show that treated employers who increased their use of costly intensive search during the experimental period reduced their use of costly search once the experiment concluded and they once again could observe past wage rate histories.

This finding indicates that when employers have access to coarse but cheap information, it may be optimal for them to use this information to reduce upfront hiring costs. However, employers need to be aware that this upfront cost savings comes at a price, as they will likely pay slightly more for an equally skilled worker.

Fifth, I demonstrate that not all employers are affected by the loss of the past wage history signal equally. The treatment disproportionally affects the search and screening behavior of employers looking to hire workers with a lower level of expertise.

The remainder of the chapter proceeds as follows: section 2 briefly reviews the relevant literature and describes this paper's contributions; section 3 describes empirical context; section 4 presents the empirical findings; section 5 discusses the results; and section 6 concludes.

2.2 Literature Review

One of the central problems in the personnel economics literature on hiring is that firms and workers cannot costlessly observe all relevant aspects of potential trading partners. This means that search is a common feature of hiring.[7] The traditional economic literature models the labor market as employees searching for wages, which are posted by the hiring firm.[8] Another, much smaller branch of the literature has focused on the demand side of the market. Barron, Bishop, et al. (1985) and Barron, Black, et al. (1989) study employer search by relating the number of applicants or interviews per employment offer and the time spent on recruiting and screening per applicant or per interview to characteristics of the vacancy and the employer. They argue that search along the extensive margin and search along the intensive margin are substitutes. In traditional models of search, a firm (or employee) acquires information simultaneously on the existence of an additional candidate (or job) as well as the value of the match. In fact, hiring procedures in the firm generally consist of two sets of activities. One set involves recruitment of applicants, while the second set involves screening and selection from among these applicants.

In contrast to most of the literature on hiring that takes a search theoretic approach, this paper follows Van Ours and Ridder (1992), who show that most vacancies are filled by a pool of applicants formed shortly after search. I separate the questions of locating an additional candidate from ascertaining the match quality of a candidate. Thus, the important economic question is which of the applicants to the job the firm should choose, or whether the firm should choose none of the available applicants at all.[9] In doing so, I deviate from much of the current literature, and choose to describe the firm's hiring decision as an information acquisition model where a firm must choose how much information to acquire

[7] See Devine and Kiefer (1991) for a detailed summary of the search literature on hiring.

[8] For example, Stigler (1961), Mortensen (1970), and McCall (1970); see Mortensen and Pissarides (1999) for a detailed review of costly applicant search.

[9] Additional support for focusing on selection comes from Van Ours and Ridder (1993) who finds that employers spend far more time on selection than on search.

about applicants before choosing whether to hire an applicant.[10]

The economic literature on endogenous information acquisition usually assumes that agents acquire information when the value of information exceeds its cost (Arrow, 1996; Grossman and Stiglitz, 1980). The information acquisition literature has largely focused on situations where externalities such as the public nature of information or behavioral considerations lead to either an over or under investment in information acquisition.[11]

In contrast to the existing literature, which highlights inefficient investments in information, I focus on how the endowment of information to agents alters incentives to acquire more costly information. Building on evidence that the equilibrium level of information acquisition is fundamentally dependent on market institutions, this chapter empirically demonstrates the effects of the market provision of information on employer hiring strategies, including costly search and screening.[12]

The literature on determinants of hiring focuses on how either firm and job attributes or market attributes affect recruiting and screening behavior. Holzer (1996) argues that employers choose among the various recruitment and screening methods on the basis of their relative costs and benefits.[13] Barron, Black, et al. (1989) shows that employers tend to screen more applicants for more important positions. My paper is closely related to Barron, Berger, et al. (1997), since I also allows the firm to endogenously acquire information about the match quality of the applicant. Barron found that both applicants and vacancy characteristics strongly influence firms' search at both the intensive and extensive margin. My paper, in contrast, focuses on the effect of market provided information on the firms' information acquisition decision, and the interaction with firm and applicant characteristics.[14]

This paper also contributes to a growing literature that details the impact of screening technology on the quality of hires. While information technology and big data signals seem to suggest "efficiency" to many practitioners and academics, I find evidence that they can get in the way of costly but more informative practices like interviewing. Strategies discussed in the literature include the use of job testing, labor market intermediaries, and employee

[10]See figure A.1 which confirms that the applicant pool is generally set before employers begin to message applicants.

[11]See Kübler and Weizsäcker (2004), Kraemer et al. (2006), and Hoffman (2014) as a recent sampling of experimental literature which focuses on behavioral reasons for miss-allocation in information.

[12]Endogenous information acquisition has been used to analyze auctions (e.g., Milgrom and Weber (1982)), voting (e.g., Martinelli (2006) and Persico (2004)), and medical patient decision-making (e.g., Kőszegi (2003)), among many other applications.

[13]Holzer (1987) was the first to detail an employer search model in which firms choose hiring procedures as well as reservation productivity levels. Fallick (1992) details how more expensive recruitment methods must provide more applicants and or better ones.

[14]There also exists a more macro focused literature including Burdett and Cunningham (1998) which advocate that the analysis of employers' search would be greatly improved if the market conditions the firm faced at the time of search could be taken into account. Russo, Rietveld, et al. (2000) details how tightness of the labor market affects employer recruitment behavior. The procyclicality of on-the-job search is mainly driven by the increase in the availability of better jobs in Pissarides (1994) and Schettkat (1996). Russo, Gorter, et al. (2001), adds to this literature by analyzing changes in recruitment behavior at the individual firm level at different points of the business cycle.

referrals. My paper is similar to Hoffman et al. (2015) in that it investigates the effects of verifiable information on employer hiring. Hoffman found that in addition to reducing costs, verifiable signals such as job testing can solve agency problems in hiring. They also show that applicants hired using job testing have longer tenure than those hired through traditional means. My finding that public verifiable information can reduce incentives to conduct costly search adds an interesting confounding factor to their results, as it makes clear the complete costs of using job testing to solve agency problems.[15]

Finally, the paper seeks to tie hiring strategy directly to firm outcomes by measuring the causal impact of the employer's use of public information on costly search and the effect of that costly search on employer satisfaction. Recent theoretical literature in management has highlighted the critical role of individuals in creating and sustaining competitive advantages (Abell et al., 2007; Teece, 2007).[16] Understanding firm-level recruitment and hiring strategies may help explain persistent firm-level conditional differences observed in profitability (Oyer and Schaefer, 2011; Blasco and Pertold-Gebicka, 2013). My paper details the importance of costly search that enables an employer to assess applicants high-bandwith information and hire cheap, high-quality employees.

2.3 Empirical Context

During the past ten years, a number of online labor markets have emerged. In these markets, firms hire workers to perform tasks that can be done remotely, such as computer programming, graphic design, data entry, research, and writing. Online labor markets differ in their scope and focus, but common services provided by the platforms include maintaining job listings, hosting user profile pages, arbitrating disputes, certifying worker skills, and maintaining feedback systems. The experiment reported in this paper was conducted on oDesk, the largest of these online labor markets.

On oDesk, employers write job descriptions, self-categorize the nature of the work, and choose required skills for jobs posted to the oDesk website. Workers learn about vacancies via electronic searches or email notifications, and can submit applications to any publicly advertised job on the platform. After a worker submits an application, the employer can interview and hire the worker on the terms proposed by the worker or make a counteroffer, which the worker can accept or reject, and so on.

In 2014, employers spent $900 million on wages through oDesk. The 2013 wage bill was $285 million, representing 200% year-on-year growth from 2013. As of October 2014, more than 3.5 million employers and 9.7 million contractors had created profiles, though a

[15] See Autor and Scarborough (2008) for other recent work focusing on the effects of job testing on worker performance. Work by Autor (2001a), Stanton and Thomas (2012b), and Horton (2013) focuses on labor market intermediaries. Recent empirical work focusing on the effects of employee referrals includes Burks et al. (2013), Brown, Setren, et al. (2012), and Pallais and Sands (2013).

[16] Also see Barney (1991), Hall (1993), Reed and DeFillippi (1990), Peteraf and Barney (2003), and Coff and Kryscynski (2011).

considerably smaller fraction are active on the site. Approximately 2.8 million job openings were posted in 2014.[17]

Based on dollars spent, the top skills in the marketplace are technical skills, such as web programming, mobile applications development (e.g., iPhone and Android), and web design. Based on hours worked, the top skills are web programming, data entry, search engine optimization, and web research. The top five countries for oDesk employers are: the United States, the United Kingdom, France, Germany, and Israel. The top five countries for oDesk workers are: the United States, Philippines, Russia, Bangladesh, and the United Kingdom.

There has been some research that focuses on the oDesk marketplace. Pallais (2014) uses a field experiment to show hiring inexperienced workers generates information about their abilities. She further shows that because workers cannot compensate firms for this information, inexperienced workers may be underemployed. Stanton and Thomas (2012b) use oDesk data to show that agencies (which act as quasi-firms) help workers find jobs and break into the marketplace. Agrawal, Lacetera, et al. (2013) investigates which factors matter to employers in making selections from an applicant pool, and presents some evidence of statistical discrimination; the paper also supports the view of employers selecting from a more-or-less complete pool of applicants rather than serially screening.

Transacting on oDesk

The process for filling a job opening on oDesk is qualitatively similar to the process in conventional labor markets. First, a would-be employer on oDesk creates a job post and chooses whether to make it public or private. Public jobs can be seen by all workers on the platform, while only invited applicants can see private jobs. Employers choose a job title and describe the nature of the work. Additionally, employers choose a contractual form (hourly or fixed-price), specify what skills the project requires (both by listing skills and by choosing a category from a mutually exclusive list), and specify how much experience they want applicants to have. Employers also estimate how long the project is likely to last. Once the job posting is written, it is reviewed by oDesk and then posted to the marketplace.

Once posted to the marketplace, would-be job applicants can view all this information. Additionally, oDesk also presents verified attributes of the employer, such as their number of past jobs, and average wage rate paid. Applicants can apply to any public job posting on the platform. When they apply, they include a bid and a cover letter. After applying, the applicant immediately appears in the employer's "applicant tracking system," or ATS, with their name, picture, bid, self-reported skills, and a few pieces of oDesk-verified information, such as total hours-worked and average feedback rating from previous projects (if any). Figure 2.1 shows the employer's default view of the ATS prior to viewing the detailed information contained in the job application. By default, employers only view job applicants who are predicted to be a good fit using oDesk's proprietary machine-learning algorithms.

[17] See Agrawal, Horton, et al. (2013) for additional descriptive statistics on oDesk.

To view a complete list of applicants the employer must click on the "Applicant" tab on the right side of the screen.

Employers can click on an applicant's limited listing to view their full profile, which has that applicant's disaggregated work history, with per-project details on feedback received, hours worked, and wage rate earned on all past jobs. Figure 2.2 shows an employer's view of an applicant's disaggregated work history after expanding or viewing an application.

Although all job applications start with the worker applying to a job opening, not all of these applications are initiated by the worker. As in conventional labor markets, oDesk employers may choose to actively recruit candidates to apply for their jobs. Upon completing a job posting employers are shown a pool of ten applicants who report having the skills requested for the posted job. Employers are given the option of inviting some or all of these applicants to apply to the job. Additionally, the employer can search on his own for some skill or attribute they are looking for in candidates. The search tools on oDesk will return lists of workers, and will contain information about that worker's past work history. If they choose, employers can then "invite" a worker they are interested in. These recruiting invitations are not job offers, but rather invitations to apply to the employer's already-posted job opening.

Only 36% of employers choose to recruit applicants on jobs posted on oDesk. Of employers who choose to recruit, on average three out of four recruited applicants are recruited from the oDesk-provided pool of applicants shown to all employers after posting a job. Of course, these recruited applicants are not required to apply to the job opening—about half do apply. These "recruited" applicants and organic applicants (applicants who are not recruited) both appear in the employer's ATS. Employers are free to evaluate candidates at any time after they post their job.

If the employer hires a candidate via oDesk, oDesk mediates the relationship. If the project is hourly, hours worked are measured via custom tracking software that workers install on their computers. The tracking software, or Work Diary, essentially serves as a digital punch clock.

The oDesk marketplace is not the only marketplace for online work (or IT work more generally). One might thus worry that every job opening on oDesk is simultaneously posted on several other online labor market sites *and* in the traditional market. However, survey evidence suggests that online and offline hiring are only very weak substitutes and that posting of job openings on multiple platforms is relatively rare. For example, Horton (2010) found limited evidence of multiple postings when comparing jobs posted on oDesk and its largest (former) rival, Elance.

Description of the Experiment

In September 2014, I conducted an experiment on oDesk that altered the information set available to employers by hiding applicants' past wage rates. All employers operating on the oDesk platform, both new and experienced, were eligible for the experiment, and were randomized for the duration of the experiment into either a treatment or control group when they posted their first job during the experimental period. Once an employer was assigned to the treatment group, the treatment affected all aspects of the site. Thus, past wage information was completely hidden on workers' profiles both in search results and when viewing the detailed profile after clicking on an applicant in the ATS. Table 2.1 shows that randomization was effective and the experimental groups were well balanced.

Figure 2.3 shows the changes implemented in both search and the employer's ATS for treated employers. When employers click an applicant's limited listing and view his or her full profile, which has that applicant's disaggregated work history, the information presented to the employers allocated to the treatment group differs from that presented to the set of control employers. Specifically, in the per-project details view both the number of hours worked on the job and the hourly wage rate earned on any job were removed. Employers allocated to the control group can view both the price, p, the applicants worked for as well as the quantity of work performed, q. Employers allocated to the treatment group can only observe the $p * q$ or the total earnings on each job.

When reviewing past work history, which according to a February 2015 client survey was the second most important piece of information in hiring after the bid amount, employers can only observe the title of the job, the total earnings on the job, and any feedback left (if available) by the past employer. Other than the lack of job specific past wage information, there were no other changes in the user interface or in information available between treated and control employers. Treated employers still knew the number of past jobs the applicants worked, the applicants' total number of hours worked on the platform, and the applicants' total past earnings. Thus, it is possible for employers to calculate applicants' average past wages, but not applicants' current wage rate.

Overview of the Data

Over the course of the experiment, 2,948 employers were assigned to the control group and posted a total of 4,661 job openings; 2,975 employers where assigned to the treatment group and posted 4,815 job openings. Table 2.2 presents summary statistics of the baseline hiring behavior on oDesk.com.[18] Beginning at the top of the hiring funnel and proceeding downwards, it is clear employers narrow down the set of applicants until they arrive at the candidate(s) they wish to hire. On average, 35 applicants apply to each job posting on oDesk. On average 1.3 of these applicants are invited to apply to the job by the employer leaving 33.6 applicants who apply to a job without being invited. I refer to these applicants as *organic applicants*.[19] On average, employers only view about seven of the applications submitted to the job by organic applicants. It is here that employers become much more selective, choosing only to send messages to about two applicants on average. Employers specifically ask at least one question to about 60% of the applicants they message, or about one applicant on average. Finally, in order to conduct "face-to-face" interviews, about half

[18]To show baseline hiring behavior, I show average statistics for the control group of employers.

[19]For the analysis on screening, I look only at the messaging behavior of organic applicants. When an applicant is invited, oDesk automatically begins a message thread between the applicant and the employer. Thus, it becomes difficult to identify invited applicants the employer actually screens.

of applicants who are messaged are asked for a Skype ID by the employer. On average, this hiring process leads to about 40% of job openings posted on oDesk being filled.

2.4 Empirical Results

The first question I seek to answer using this experiment is, how do employers shift their hiring strategies to make use of publicly available low-bandwith information such as applicants' past hourly wage rate? Hiring can conveniently be separated into search and screening. Employers search to identify possible candidates, and then screen those candidates to identify quality or fit for the job. Table 2.3 presents differences in mean search and screening measures across treatment groups.

To measure employer search, I track whether an employer "views" an application by clicking to expand the applicant's application within the applicant tracking system. Prior to "viewing" an application an employer only sees an applicant's bid, country, average feedback score, and total hours worked. Thus, "viewing" an application is a strong signal of interest in a candidate. Treated employers on average view another 0.4 applications from a baseline of 6.7 applications per opening. While this results seems small on a per-job basis, there are over 100,000 jobs posted per month. On a per-month basis, being unable to observe applicants' past wages leads employers to consider an additional 20,000 applicants, who previously were not even in the employers' consideration set.

Employers on oDesk acquire information through a costly screening process. This process involves messaging an applicant using the platform-provided messaging system. Employers message applicants, and are messaged back by applicants. All of the messages back and forth between one employer and one applicant are considered a message thread. In these messages, employers ask applicants questions, and/or exchange Skype IDs in order to set up a "face-to-face" meeting with the applicant. The second panel of Table 2.3 looks at four measures of employer screening or information acquisition behavior: (i) the number of applicants "messaged" by the employer by formally contacting the applicant for further discussion; (ii) the number of applicants an employer "questions" as evidenced by use of a question mark in at least one message; (iii) the number of applicants and employer "questions" as evidenced by use of one or more question words: *who, what, where, when, why, or how*; and (iv) the number of applicants an employer "Skypes" by asking the applicant to exchange Skype IDs. Table 2.3 shows that employers who are unable to observe applicants' past wage rates increase the number of applicants they question as evidenced by both the number of applicants messaged that include a question mark, and the number of applicants messaged that include a question word. Employers ask questions to an additional 0.15 applicants per job opening. The raw differences in mean counts do not report any statistical difference in the number of applicants messaged or the number of applicants asked for Skype IDs.

Since the four measures of screening magnitude are all rough proxies for information acquisition behavior, I use a principal component analysis to create an aggregate measure of screening behavior. This measure explains 75% of the variation in the four measures.

Figure 2.4 shows that the difference in the means between the first principal component is significantly different between employers who could and could not observe applicants' wage history.

I collected the data from a randomized experiment; therefore the simple correlations reported here can be interpreted as causal relationships. However, to further allay concerns about omitted variables, and to potentially increase precision, in the next sections I estimate the effect of hiding applicants' past wages on search and screening in a multivariate regression framework.

Employer Search

Due to the count nature of the outcome variable, I follow Budish et al. (2015) and show estimates from a quasi-maximum likelihood Poisson model with heteroskedasticity-robust standard errors. The regressions in this section are derived from a version of the following model:

$$\text{VIEWED}_j = \beta_0 + \beta_1 \textsc{Trt}_j + \textsc{Employer Covariates}_j + \textsc{Opening Covariates}_j + \epsilon_j \tag{2.1}$$

In table 2.4, the dependent variable is VIEWED_j which is the number of applicants viewed by the employer on job opening j. $\textsc{Trt}_j$ is an indicator for treatment assignment of the employer of job opening j. Model (1) shows results with no covariates; Model (2) adds both employer and job opening covariates, which include: job subcategory dummies, the number of prior jobs filled, the employer's past spending, the number of applications to the job, the number of recommended applicants to the job, the average bid by applicants on the job, and if the employer specifically requested a particular skill for the job.

Treated employers view 1.07 times the number of applications from a baseline average of seven applications per opening in the control group.[20] After observing an applicant's application and noticing that wage rate history information that was previously used in ascertaining applicant quality is missing, employers view additional candidates.[21]

Characteristics of Searched Applicants

In addition to changing the number of applicants employers search for and screen, not being able to observe applicants' past wage history could alter the characteristics of the applicants employers choose to view, message, and hire. According to a 2015 client survey on oDesk.com, the top three features, in order of importance, used when making a decision over whom to contact and eventually hire are: the applicant's hourly rate, the applicant's feedback rating

[20]The coefficient on TRT from the Poisson model is interpreted as the difference in the logs of expected count of viewed application; thus, I interoperate the coefficient as being associated with viewing exp(.069)=1.07 times as many applicants.

[21]Table A.1 in appendix A.2 shows that the results of log-linear regression models are similar to those of the quasi-Poisson models.

(specifically does the applicant have "good enough" feedback), and the applicant's experience. Thus, I will analyze three main groups of applicant characteristics: wage characteristics, experience characteristics, and feedback characteristics.

Before treatment and control group employers choose to view an application, they can only observe applicants' basic information. Thus, when deciding to view or not view an application, the information set is identical for both treatment and control employers. Therefore, I do not expect the treatment to have any effect on the characteristics of viewed applicants. It is useful to view table 2.5 as a placebo test, which confirms balance in the experiment.

Employer Screening

The differences in means presented above suggest that removing employers' ability to observe applicants' past wage rates increase screening. The addition of controls and structure by analyzing the data using quasi-maximum likelihood Poisson model with heteroskedasticity-robust standard errors confirms these results. The regressions in this section are derived from a version of the following model:

$$\mathrm{Y}_j = \beta_0 + \beta_1 \textsc{Trt}_j + \textsc{Employer Covariates}_j + \textsc{Opening Covariates}_j + \epsilon_j \tag{2.2}$$

In column (1) of table 2.6, the dependent variable, Y_j, is the number of applicants messaged in job j. Controlling for job opening characteristics as well as employer characteristics, employers in the treatment group message 1.08 times the number of applicants as control employers. The mean number of applicants messaged per job is 1.7. Thus, employers message an additional 0.14 applicants messaged per job, or over 21,000 additional applicants per month.

In column (2), the dependent variable, Y_j, is the number of applicants messaged where Skype IDs are exchanged in job j. Controlling for job opening characteristics as well as employer characteristics, employers in the treatment group ask 1.08 times the number of applicants for Skype IDs as control employers. This result is not significant at a 10% level, but is still consistent with the other results presented on employers screening behavior of applicants.

In column (3) the dependent variable, Y_j, is a count of the number of applicants messaged in job j, where at least one message exchanged with the employer contained a question mark.[22] A question mark is taken as evidence of the employer asking the applicant a question, and thus, as evidence of seeking additional information about that applicant. Controlling for the number of applicants messaged, treated employers use question marks in 1.15 times as many message threads as control employers.

In column (4) the dependent variable, Y_j, is a count of the number of applicants messaged in job j, where at least one message contains at least one of the following question words:

[22] Using a count of total question marks over all messages sent in a job opening instead of a count of message threads with question marks gives similar results. I report a count of message threads to maintain consistency with the specification in column (1) and column (2). The alternate specifications are located in the appendix.

who, what, where, when, why, or how. This result confirms that employers are substantially increasing the number of applicants questioned per job. Here, a coefficient on the treatment indicator of .178 means that treated employers use at least one question word in 1.19 times the number of message threads than control employers per job opening.

Taken all together, these four results in table 2.6, as well as the results of the principle component analysis, provide evidence that removing past wage information causes employers to not only to widen their choice set and contact more applicants, but also to attempt to get more information from the applicants they contact via message.[23]

Characteristics of Screened Applicants

Does altering the information publicly available to employers alter employers' preferences over characteristics of the applicants? I compare the observable characteristics of applicants who are messaged along three main dimensions: their wages, their experience, and their feedback.

The top panel of table 2.7 indicates that treated employers choose to message applicants with bids that are on average 6% lower than the group of applicants messaged by control employers. These are applicants who clearly value their work product lower as evidenced by a lower profile wage rate and lower historical average accepted wages. There is some slight evidence that employers also seem to message applicants who have less work experience as indicated by a 10% reduction in the average dollar amount of previous hourly earnings, and a 15% reduction in fixed-price earnings of the applicants messaged. However, this is not due to having worked fewer hours, but having been willing to previously work for a lower wage. These do not seem to be applicants who are newer to the platform or applicants with substantially worse work experience. Additionally, messaged applicants do not differ on their historical feedback scores. These results provide evidence that employers that are unable to use an applicant's past wages as a signal of quality seem to locate a subset of applicants who are substantially cheaper to an employer without being observably worse applicants as measured by market-level signals.

Hiring Outcomes

I turn now to analyzing effects of employers substituting from cheap platform provided information to information obtained through costly search and screening on the probability of a job being filled, the wage negotiation that occurs after a contract is offered, the wages paid on the job, and the employers' ex post job satisfaction.

[23]In appendix A.2, in table A.4, I present the results of application-level regressions. These regressions indicate that the odds of asking a question conditional on sending a message is higher for treated employers than control employers; although, not significantly so in all specifications. Interestingly, the results on odds of asking for a Skype ID are a tightly measured zero, indicating that employers have a tendency to ask either all or none of the applicants they message for Skype IDs.

Probability of Hiring

According to a 2013 oDesk client survey, the primary reason for not hiring a worker for an open job posting is that the employer "couldn't identify a contractor with the right skills." An employer will only fill a job when he can adequately identify an applicant who is a proper match for the position at a wage the employer is willing to pay. The treatment unambiguously reduces the number of available signals the employer can use to identify if an applicant has the right skills. Thus, according to standard search theory I would presume that fewer matches would occur. However, as an employer's ability to ascertain quality is diminished, his incentives to acquire information increase. Employers that are unable to observe past wage information increase the number of applicants they view, the number of applicants they message, and the number of applicants from whom they acquire additional information. Thus, in equilibrium, the effect of hiding past wages may actually increase the precision of an employer's signal of an applicants' quality, and the probability of hiring an applicant for an open job posting. Table 2.8 shows that the treatment increases the probability of hiring.

The regressions in this section are derived from a version of the following linear model:

$$\mathrm{Y}_j = \beta_0 + \beta_1 \textsc{Trt}_j + \epsilon_j \tag{2.3}$$

The independent variable in table 2.8 is $\textsc{AnyHired}_j$ is an indicator for whether the employer made a hire on job j, and $\textsc{Trt}_j$ is an indicator for treatment assignment on job j. In column (1) the sample is limited to only the first job posting by each employer after the start of the experiment. Further job postings were dropped from this specification to prevent against possible, but highly unlikely, employer selection into additional jobs posted to the platform. The coefficient on the treatment indicator is positive and highly significant, with the treatment increasing hiring by about 3 percentage points, from a baseline hire rate in the control group of only 40%. This is a percent increase of 7%, which is extremely high and economically important to a platform that generates revenue by taking a percentage of contractor earnings. Column (2) includes both employer-level and job opening-level covariates and shows that adding covariates has little effect on the outcome of the regression. Column (3) adds job postings that are not an employer's first during the course of the experiment. An indicator variable for the order of the job posting is included and standard errors are clustered at the employer level. Running the fill rate regression on the full sample only slightly reduces the coefficient from 0.029 to 0.026 and increases the precision of the estimates slightly. Adding covariates in column (4) increases precision slightly and has little effect on the magnitude of the coefficients. Removing a signal of an applicants' relative marginal value does not reduce the employer's ability to identify and hire quality applicants. In fact, the opposite is true, removing a cheap-to-observe signal of applicant quality substantially increases the probability of a job being filled.

Wage Negotiation and Wages

Employers who are unable to observe workers' wage rate history message workers who bid nearly 10% lower than the workers messaged by control employers, but the overall effect on wages is ambiguous as the treatment may also affect bargaining. A survey by Hall and Krueger (2010) found that only about 30% of workers report that there was some bargaining in setting the wage for their current job. The bargaining rate is especially low for blue-collar workers (5%) but much higher for knowledge workers (86%). On oDesk, at least 14% of jobs on the platform participate in some type of wage bargaining prior to signing a contract, as evidenced by agreeing to a contract wage that is not equal to the winning applicants' wage bid.[24]

Once an employer has chosen an applicant, hiding that applicant's past wage rate may alter negotiations over pay with the employer. Table 2.9 shows that there is a small but positive effect on wages, which comes from workers being offered a contract on which they make a higher percentage of their bid. Column (1) reports an estimate of the regression:

$$\log\left(\text{WAGE-TO-BID RATIO}_{ja}\right) = \beta_0 + \beta_1 \mathrm{TRT}_{ja} + \epsilon_{ja} \tag{2.4}$$

where $\text{WAGE-TO-BID RATIO}_{ja}$ is calculated as the ratio of the wages paid to the winning applicant to the hourly bid submitted by the winning applicant in assignment, a, which came from job posting j. An assignment is the oDesk-specific word for a job once both parties have signed the contract. In this model, there is always only one assignment to each job posting, as I limit the analysis to only the first job assignment which is created from a job posting. I subset the data in this way, as negotiation effects on follow-up job assignments, which stem from an job posting in the experiment, cannot be directly attributable to the hiding of the applicant's wage history. A wage-to-bid ratio of 1.0 means the employer paid the employee exactly the employee's bid amount. A wage-to-bid ratio below 1.0 indicates that the employer is paying the employee an amount less than the employee's bid. The coefficient on the treatment indicator is positive and highly significant, with the treatment increasing the wage to bid ratio by about 1.2, from a baseline ratio of 0.973.

Table 2.10 demonstrates that despite the positive effect on negotiations, there is an overall negative effect of the treatment on wages. Employers who are unable to observe applicants' wage rate history pay wages that are 10% lower. These results indicate the treatment has two very interesting affects on wages. Firstly, there is a very large negative selection effect. Employers who are unable to observe applicants' past wage rates screen more and more deeply, and identify cheaper workers. There is also a positive information rent gained by the hired workers, as employers no longer know the applicants' past wage rate history and thus agree to pay the chosen applicant a higher percentage of his bid.

[24] This is a lower bound estimate, as it's possible that there is wage negotiation, but the amount settled on was exactly the contractor's bid. Getting data on wage negotiation and delving into this phenomenon is beyond the scope of this paper.

Feedback

Employers in the treatment group are hiring cheaper workers, but not workers who appear worse in measurable quality ex ante. Thus, employers who do not know past wage rate information tend to get their work completed cheaper relative to employers who make hiring decisions conditional on applicants' past wage rate histories. Although treatment employers do not choose applicants who are ex ante of lower quality than those chosen by control employers, the applicants' true quality is still unknown. Thus, I turn my attention to determining if applicants hired through the fundamentally different search process caused by removing past wage rate information do a better job completing the job.

Specifically, conditional on an employer leaving feedback on the hired applicant, do treated employers leave different feedback from control employers? Table 2.11 shows there is no measurable difference in public feedback left between treatment and control employers, but that employers who could not view an applicant's past wage rates leave better private feedback on the first job they hired from a job opening. The regressions in this section are derived from a version of the following model:

$$\log\left(\textsc{Feedback Variable}_{ja}\right) = \beta_0 + \beta_1 \textsc{Trt}_{ja} + \textsc{Controls}_{ja} + \epsilon_{ja} \mid \textsc{Feedback Left}_{ja} \tag{2.5}$$

where the sample is again limited to only first job openings posted by each employer and only the first of the assignments spawned by that job opening. The reason the analysis is limited to only the first assignment spawned from a job opening is that including follow-up assignments biases the results as employers are more likely to create a follow-up assignment when the worker is of high quality. Controls for the employers' prior experience, the employer's prior experience with the hired worker, and category indicators are added to both models.

In column (1), Feedback Variable is the publicly viewable 1-star to 5-star feedback rating that an employer can leave for a worker on assignment a from job opening j after the job is complete, $\textsc{Trt}_{ja}$ is an indicator for treatment assignment of assignment a, which came from job posting j. The coefficient on the treatment indicator in column (1) is not significant, which indicates that employers who hire without knowing an applicant's past wage history do not leave better or worse feedback than employers who know this information.

There has been substantial research, including Nosko and Tadelis (2015) and Horton and Golden (2015) that details the limits of public reputation in platform markets. On oDesk, both employers and workers also have the option of leaving private feedback that is never displayed on the platform. They are told that this feedback will only be used internally by oDesk.

In Column (2), Feedback Variable is equal to a 0-10 rating, which is viewable only to oDesk. The coefficient on the treatment indicator is positive and significant. Employers who were not able to view applicants' past wages when hiring leave feedback scores that are 5% higher than the score left by employers who hired applicants with knowlage of their past wage rate history.

Are Employers Making a Mistake?

The evidence presented in the previous sections demonstrates that when employers cannot observe wage rate history they choose to acquire information about applicants' quality though more costly intensive search. This intensive search allows employers to better identify, at an upfront cost, high-quality low-cost employees, which can have long-term positive effects on firm outcomes. Thus, should employers ignore market provided low-bandwidth information and rely on costly screening all the time?

To assess this, I take advantage of the fact that I observe the behavior of treated and control employers after the experiment was concluded, and look for persistent treatment effects. I plot the difference in mean levels of screening for treatment and control employers for each employer's first job posted during the experimental period, and each employer's first job posted after the experimental period ended. The top panel of figure 2.5 shows that employers search more than control employers during the experimental period. The bottom panel of figure 2.5 shows that there is no difference in the screening behavior of employers treated during the experiment and employers not treated during the experiment once the experiment was completed.

I also examine the impact of the treatment on treated employers in the post period using difference-in-difference methodology. The models presented are quasi-maximum likelihood Poisson regression with heteroskedasticity-robust standard errors:

$$\begin{aligned} \mathrm{Y}_j = \beta_0 + \beta_1 \textsc{Trt}_j + \textsc{Post Period}_j + \textsc{Trt x Post Period}_j + \textsc{Employer Covariates}_j \\ + \textsc{Opening Covariates}_j + \epsilon_j \end{aligned} \tag{2.6}$$

where Y_j is one of the measures of costly intensive search used previously on job j. Table 2.12 shows that employers for which past wages were hidden reduce their use of costly search in their first job posted after the experiment ended. These results are consistent with a story that indicates that when employers have access to imprecise signals of worker quality they take advantage of these signals and use them instead of other more costly information acquisition strategies. However, it is extremely important to understand that these results are generated in a market for task-based labor where the long-run effects of hiring a slightly less skilled employee or an employee at a slightly higher cost are minimized relative to traditional labor markets due to lower costs of firing and shorter work durations.

Heterogenous Screening Effects

There is no one strategy that fits all when it comes to screening strategies. Generally, personnel economics assumes there are complementarities between certain firms and certain employees, such that firms should tailor their hiring to attract the employees that generate the most match specific productivity.[25] On oDesk, jobs range from data entry to complicated

[25] The assumption of such a complementarity underlies the large literature on assortive matching in labor markets (see Rosen, 1982 and Sattinger, 1993).

legal work, and the quality and experience of workers who compete to complete these jobs ranges from novice to expert. The platform offers employers the opportunity when posting a job opening on oDesk to choose if they would like to see more beginner, intermediate, or expert workers. This expressed preference over contractor type gives insight into how important a high-quality worker is to the employer on each job opening.

Figure 2.6 shows the differences in screening levels for treatment and control employers separated by the requested expertise level of the applicants. This split sample indicates that employers interested in hiring an expert level worker do not increase their level of costly search when the employer cannot observe applicants' wage rate history. This result is further confirmed by use of an interaction model that looks for differential treatment effects on intensive search by preferred contractor tier using using a quasi-maximum likelihood Poisson regression with heteroskedasticity-robust standard errors:

$$\begin{aligned} \mathrm{Y}_j = \beta_0 + \beta_1 \textsc{Trt}_j + \textsc{Contractor Tier}_j + \textsc{Trt x Contractor Tier}_j \\ + \textsc{Employer Covariates}_j + \textsc{Opening Covariates}_j + \epsilon \end{aligned} \tag{2.7}$$

The results of table 2.13 indicate that the treatment significantly increases the level of costly intensive screening for employers looking for beginner level workers. Employers that cannot observe wage rate history and who are looking to hire beginner applicants exchange Skype IDs with 1.25 times as many applicants as employers that can observe wage history and are interested in hiring beginner applicants. The significant negative coefficient on the interaction term for "TRT x Expert" indicates that the treatment effects are significantly different for employers looking to hire experienced workers than for employers looking to hire beginner level workers. The effects on employers interested in hiring intermediate expertise workers are smaller than the effects on employers hiring for beginner jobs, but still significantly different from zero. Removing workers wage rate history only alters employer screening behavior for those employers not interested in hiring expert level workers.

Robustness Check

If the employer were to have other high-quality signals of applicant type available, such as having previously observed the applicant's productivity, I would expect the employer to rely on this information and, thus, hiding past wage information would have no effect. In addition to posting public jobs on oDesk, if an employer wishes to work with a specific applicant, he is able to create a private job posting. Only applicants expressly invited by the employer may apply to private job postings. In the sample, private job postings have a median number of applicants of one. Employers therefore usually have a much deeper knowledge of the quality of the applicants to private job postings. Table 2.14 presents the results of running the main screening regression subsetted to include only private jobs. The coefficients on the $\textsc{Trt}_j$ indicator are not significantly different from zero for all four measures of costly screening, demonstrating that there is no treatment effect on screening for private jobs.

2.5 Discussion

I find that when employers on oDesk are unable to observe workers' past wage rates, they seek to acquire more information through costly search and screening. This increase in acquired information allows employers to fill more jobs and hire applicants who are cheaper but not worse on other observables. These employers also report being more satisfied with the work output. However, even after observing that acquiring information through costly search and screening leads to more and better matches, employers return to lower levels of search and screening behavior post-experiment. I show that these results are primarily driven by employers looking to hire workers with lower levels of expertise.

This pattern of results is consistent with a particular story of information acquisition in markets where costly search and screening is available. Specifically, the results indicate that there is a fixed price to acquiring information through search and screening. If the employer is searching for someone low skilled then the provision of coarse information from the market is good enough and the employer is not going to pay a fixed cost to acquire more information. If the employer is looking for someone high skilled he may be willing to pay the fixed cost to acquire more information about the applicants. If coarse information is not provided by the marketplace, then even when the employer is looking for someone unskilled he is willing to pay the fixed cost and acquire more information. This leads to the employer hiring a cheaper worker and being ex post more satisfied. But is the employer's increased satisfaction enough to indicate that they were previously making a mistake by relying on platform provided information when looking for low-skilled workers?

The evidence in this paper suggests that employers who increase their search and screening behavior when they are unable to observe past wage rates decrease their use of costly search and screening when the experiment is completed and they can once again observe workers' past wage rates. This result indicates that employers do not make a mistake by relying on platform provided information at least when hiring beginner level workers.

To get a better perspective on why employers who want to hire unskilled workers choose to rely on platform provided information instead of costly screening although they know this will lead to hiring a more expensive worker, I conduct some back-of-the-envelope calculations.

Figure 2.7 shows that employers hiring a beginner applicant expect to pay substantially less than those hiring intermediate or advanced applicants. The average job that hires a beginner level contractor costs about \$500 compared to about \$1700 for hiring an expert level contractor. Based on the finding that employers who cannot observe applicants' past wage rates offer contracts to applicants that bid about 9.5% lower, I can estimate that costly screening saves an employer hiring a beginner applicant about \$50 on a job in wages. This is compared to about \$170 in wages when hiring for a job that is interested in expert level workers. While estimating the upfront cost of increased search and screening behavior is a bit more difficult, I do know from survey data contained in figure 2.8 that 80% of employers spend fewer than eight hours on search and screening. Most likely the costs of hiring are not linear. It seems fair to assume that a majority of hiring costs come not on posting or viewing, or messaging candidates on platform, but on organizing and screening the candidates off platform. Thus, it seems logical that employers prefer to pay an extra $50 in wages and save an hour or two in time upfront costs, but might not be willing to pay an extra $170 in wages to save a similar amount of time.

2.6 Conclusion

I designed and implemented an experiment that seeks to understand how employers on oDesk.com make use of large quantities of standardized and verified publicly available information. This information is usually not available to employers in the traditional labor market. When employers do not have the ability to observe an applicant's wage history, they substitute for this informational loss by exerting costly effort to acquire more information about candidates though interviewing. This strategic reaction to a reduced information environment actually leads employers to be more likely to fill an open job posting and to hire cheaper candidates. This treatment effect is limited to employers looking to hire low level and intermediate level workers.

It is important to note the limitations of the analysis in this paper. One very important limitation of my analysis is that all firms in this market are hiring task-based labor. Thus, incentives to locate a top-notch employee are lower than in the traditional labor market, since both the expected tenure of an employee and the firing costs are much lower than in the traditional labor market. Additionally, my experiment studies the effects of removing information, not adding information. Firm institutions are extremely important in hiring. For example, some firms recruit every year at a fixed set of colleges regardless of changes in academic rankings. Perhaps basing hiring decisions on wage history is as much a function of tradition as optimal information use. If this were the case, I would expect to observe a larger treatment effect when compared to adding wage history information to a market that traditionally did not have this information.

My results are potentially relevant for understanding under which circumstances firms might seek to use online labor marketplaces. Online labor marketplaces can only reduce asymmetric information when the signals they provide are useful in matching. For example, if a firm is interested in hiring highly skilled labor, my findings suggest that the market provided signals are of little use. Instead, the firm must rely on costly screening by acquiring information about applicants. This might make an online labor platform a less attractive option for this type of labor, as one of the platform's main advantages, reduced hiring costs, cannot be fully harnessed. On oDesk, costly screening of applicants by asking questions or conducting Skype interviews is generally associated with more experienced employers. Less experienced employers may be relying too heavily on market provided signals, which could reduce the quality of their matches and slow their access to the platform. By removing market provided signals like past wage rate information, the platform may be able to "push" first time users of the platform into relying less on the vague market provided signals and more heavily on costly techniques such as asking questions and conducting Skype interviews.

Finally, my results have more general implications about the relationship between coarse public information and costly private information acquisition. Firms need to consider the effects of providing their hiring agents with coarse data especially when long and short term incentives of agents and firms are not completely aligned. Hiring agents may use this data as a substitute for more costly but more precise information acquisition strategies. This increased use of coarse information might help to explain persistant discrimination towards underrepresented populations. These populations often appear to have lower coarse signals, but might be identified as diamonds in the rough if screened using more costly intensive interviewing.

Figure 2.1: Default View of the Applicant Tracking System (ATS)

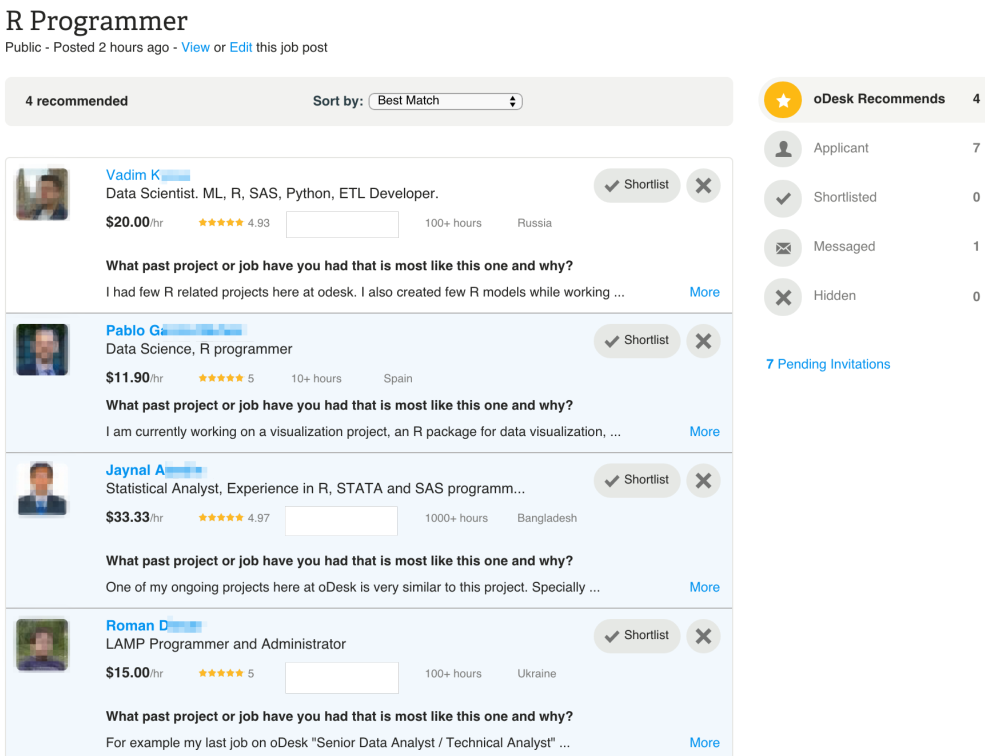

Note: This is the default listing of applications as observed by the employer after posting a job and having applicants apply. Only applicants recommended by oDesk's proprietary matching algorithm are displayed by default. Notice there are seven total applications submitted at the time of the screenshot but only four are displayed by default. Employers can directly contact applicants from this page, directly hire applicants from this page, or click on a listing to expand it and view the applicants' complete application and work history.

Figure 2.2: Expanded View of Disaggregated Work History

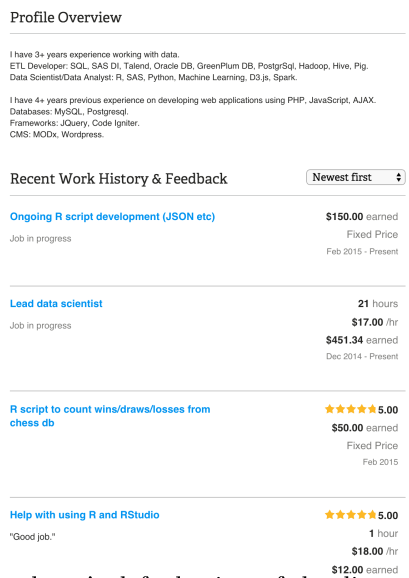

Note: This is the employer's default view of the disaggregated work history of an applicant which is viewable after expanding (viewing) an application. In the treatment group, the 21 hours and \$17.00/hr would be hidden for the "lead data scientist" job. Only the \$451.34 would be observable. The employer also sees an expanded profile (see A.2).

Figure 2.3: Changes to the UI for treated employers.

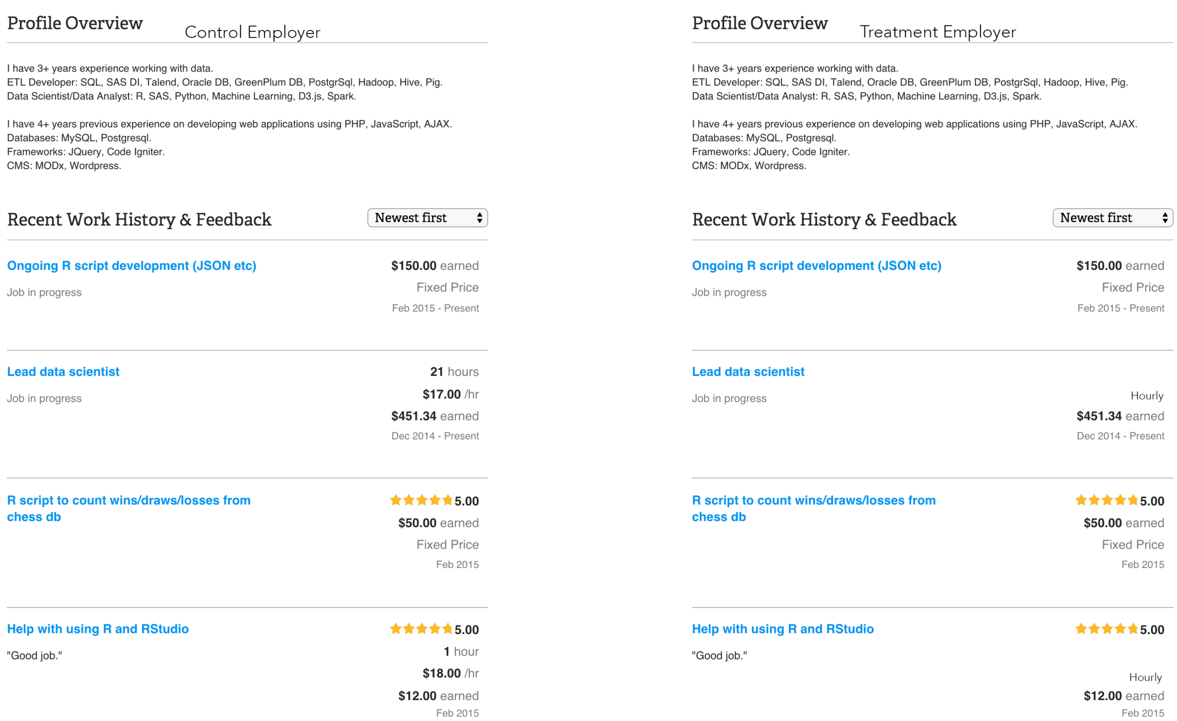

Note: This screenshot details the changes made to workers' profile pages.

Figure 2.4: First Principle Component of Screening

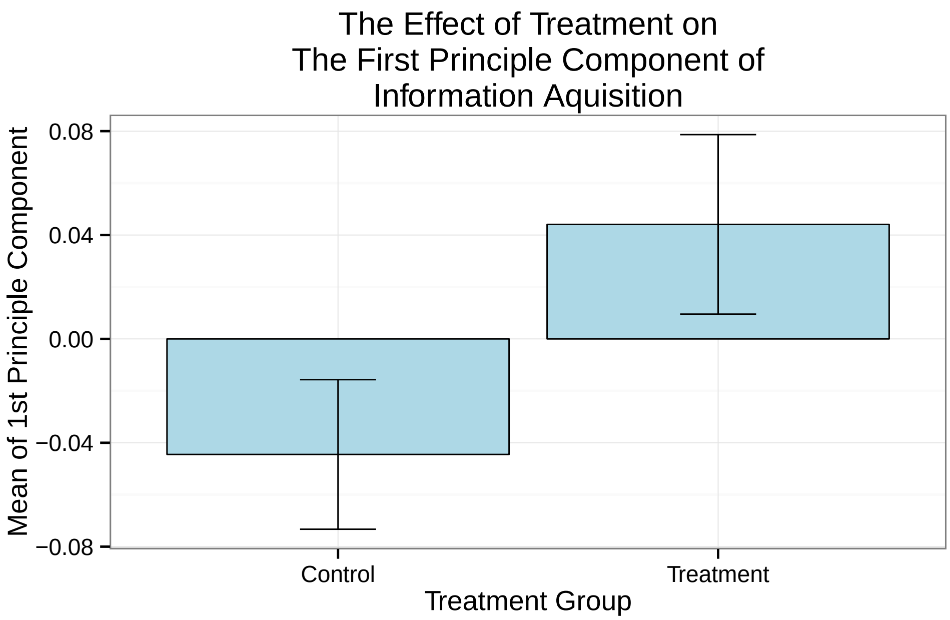

Note: The analysis was limited to the first hourly job posting after the start of the experiment. The four inputs to the principle component analysis are: number of applicants messaged, number of applicants asked a question as indicated by use of a question mark, number of applicants asked a question as indicated by use of a question word, and the number of applicants asked for a Skype ID.

Figure 2.5: Screening by Experimental Period

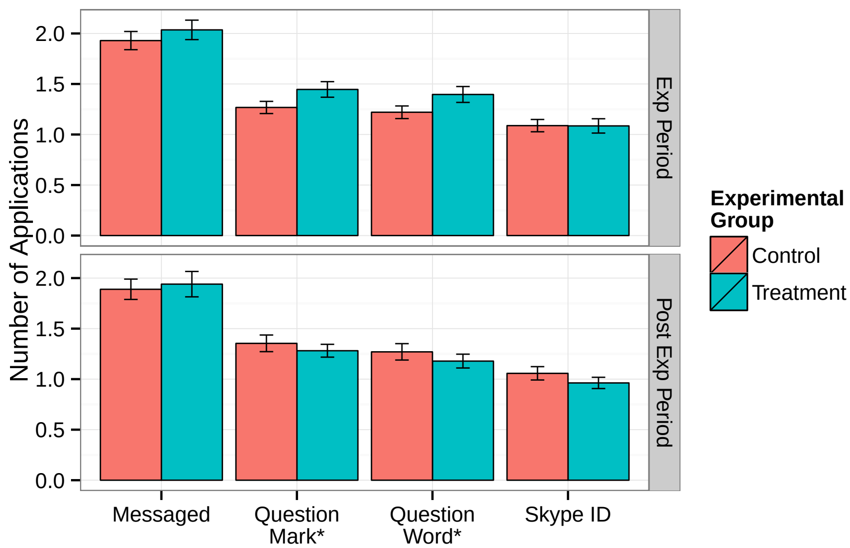

Note: This figure shows difference in raw counts and standard errors of screening outcomes between treatment and control employers separated by the jobs posted during the experimental period and directly after the experiment ended.

Figure 2.6: Screening by Requested Contractor Tier

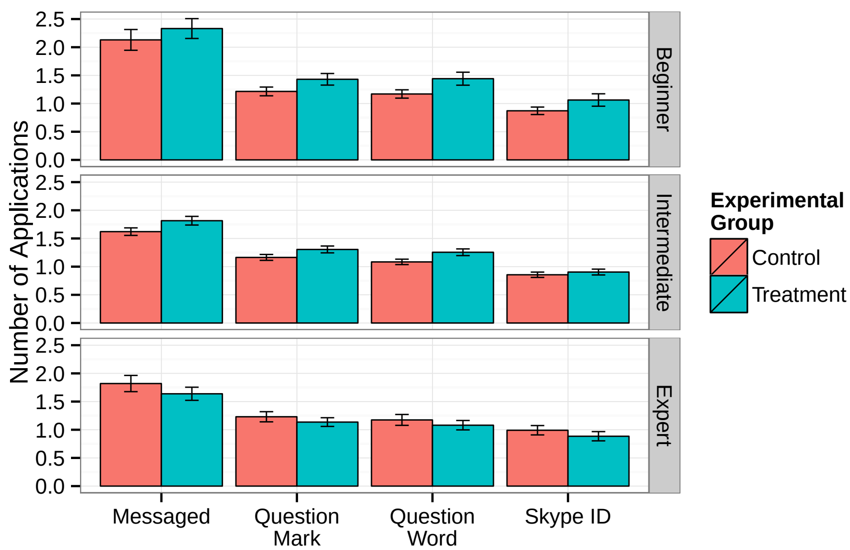

Note: This figure shows difference in raw counts and standard errors of screening outcomes between treatment and control employers separated by the employer requested skill level of applicants.

Figure 2.7: Job Characteristics by Contractor Tier

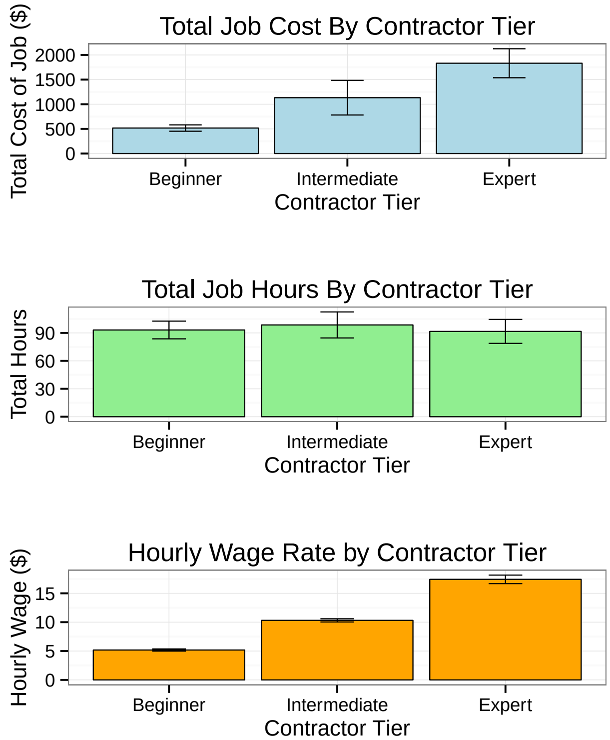

Note: This figure is limited to hourly job postings that were filled and billed more than $0.01.

Figure 2.8: Interviewing Survey

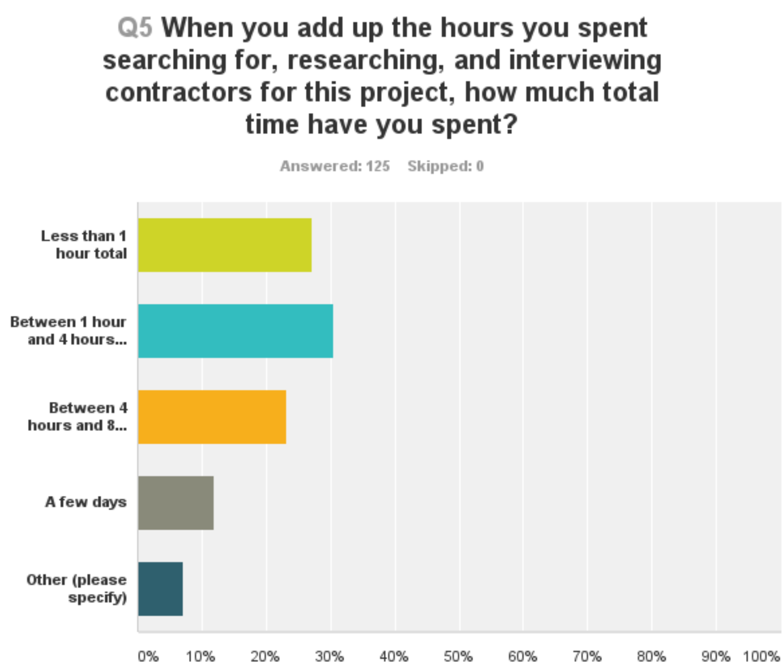

Note: This data came from a 2013 oDesk client survey.

Table 2.1: Balance Table of Employer, Job Posting, and Applicant Characteristics

	Control mean: $\bar{X}_{CTL}$	Treatment mean: $\bar{X}_{TRT}$	Difference In Means	p-value	
Employer Attributes					
Prior Job Postings	23.49	23.98	0.49 (1.31)	0.71	
Prior Billed Jobs	10.71	11.29	0.58 (0.63)	0.35	
Prior Spend by Employers	5643.10	6053.35	410.25 (453.99)	0.37	
Num Prior Contractors	10.84	11.73	0.89 (0.80)	0.27	
Avg Feedback Score of Employer	4.81	4.79	-0.03 (0.02)	0.09	*
Num of Reviews of Employer	8.05	8.84	0.79 (0.71)	0.27	
Job Posting Attributes					
Number non-invited Applicants	33.62	33.44	-0.18 (1.09)	0.87	
Avg Best Match Score	0.36	0.36	0.00 (0.00)	0.41	
Avg Bid	12.76	12.60	-0.16 (0.24)	0.52	
Preferred Experience in Hours	33.69	34.25	0.56 (3.40)	0.87	
Estimated Job Duration in Weeks	17.19	16.93	-0.26 (0.55)	0.63	
Applicant Attributes					
Tenure in Days	895.32	884.86	-10.46 (6.64)	0.12	
Hours Worked to Date	1406.69	1375.84	-30.85 (29.07)	0.29	
Num Past Jobs Worked	29.82	30.45	0.63 (0.55)	0.25	
Past Hourly Earnings	11510.29	11018.87	-491.42 (313.64)	0.12	
Past Fixed Wage Earnings	1799.65	1745.09	-54.55 (54.93)	0.32	
Num Prior Employers	23.44	23.88	0.43 (0.39)	0.27	
Min Feedback Rating	2.84	2.84	0.00 (0.01)	0.84	
Avg Feedback Rating	4.59	4.59	-0.00 (0.00)	0.27	
Max Feedback Rating	4.97	4.97	-0.00 (0.00)	0.01	
Wage Bid	10.73	10.40	-0.33 (0.24)	0.18	
Profile Wage	10.88	10.56	-0.31 (0.22)	0.15	
Min Hr. Wage (6 months)	7.14	6.89	-0.26 (0.16)	0.11	
Avg Hr. Wage (6 months)	8.67	8.42	-0.25 (0.18)	0.17	
Max Hr. Wage (6 months)	10.77	10.52	-0.25 (0.22)	0.26	

Notes: This table reports means and standard errors across experimental groups of employer, job posting, and applicant characteristics. The top panel reports characteristics of employers allocated to treatment and control. The middle panel reports characteristics of job postings by treatment and control groups for the first job posted after allocation to the experiment for each employer. The bottom panel reports characteristics of employers at the time they were allocated to treatment or control groups. The bottom and top 1% by average historical wage were dropped. Reported p-values are for two-sided t-tests of the null hypothesis of no difference in means across groups. In the bottom panel, standard errors are clustered at the employer level. Significance indicators: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

Table 2.2: Baseline Screening Behavior

Statistic	N	Mean	St. Dev.	Min	Median	Max
Number of Applicants	2,948	35.105	43.296	0	22	639
Number of Organic Applicants	2,948	33.691	43.036	0	20.5	639
Number of Applications Viewed	2,948	7.321	9.257	0	5	122
Number of Organic Applicants Messaged	2,948	1.797	3.684	0	1	91
Number of Organic Applicants Questioned	2,948	1.121	2.050	0	0	36
Number of Organic Applicants Skyped	2,948	0.890	1.890	0	0	23
Number of Hires	2,948	0.580	1.039	0	0	26
Pct of Jobs Filled	2,948	0.403	0.491	0	0	1

Notes: This table provides baseline (control) statistics of hiring on oDesk. The statistics reported are for first job openings of employers assigned to the control group. Organic applicants are applicants who were not invited to apply to the job.

Table 2.3: Search and Screening Behavior

	Control Mean: $\bar{X}_{CTL}$	Treatment Mean: $\bar{X}_{TRT}$	Difference In Means	p-value	
Measures of Search					
Num. Viewed Applications	6.671	7.122	0.451 (0.242)	0.062	†
Measures of Screening					
Num. Messaged Applicants	1.797	1.925	0.128 (0.097)	0.188	
Num. Skyped Applicants	0.890	0.946	0.056 (0.057)	0.322	
Num. Questioned Applicants (Q Word)	1.121	1.271	0.150 (0.062)	0.015	*
Num. Questioned Applicants (Q Mark)	1.187	1.308	0.120 (0.060)	0.046	*

Notes: This table reports means and standard errors across experimental groups for the number of applicants searched, and intensely screened by treatment and control group. The top panel reports one measure of search behavior, and the bottom panel reports four measures of screening behavior. Reported p-values are for two-sided t-tests of the null hypothesis of no difference in means across groups. Significance indicators: $p \leq 0.10$:† , $p \leq 0.05$: $*$, $p \leq 0.01$: $**$ and $p \leq .001$: $***$.

Table 2.4: Effect of Hiding Past Wages on Search

	Dependent variable:	
	Num Viewed	
	Poisson	Poisson
	(1)	(2)
Hidden Wages Treatment Group (TRT)	0.065*	0.069**
	(0.035)	(0.034)
Constant	1.898***	1.773***
	(0.025)	(0.134)
Opening Level Covariates	No	Yes
Observations	5,922	5,855

Notes: The sample is restricted to hourly first job posts by an employer. Models (2) include covariates including: category indicators, prior jobs billed by the employer, employer's prior spendings on the platform, the number of applications to the job openings, the number of recommended applications to the job opening, and the average applicant's bid, and an indicator if requested skills. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table 2.5: Characteristics of Viewed Applicants

	Control Mean: $\bar{X}_{CTL}$	Treatment Mean: $\bar{X}_{TRT}$	Difference In Means	p-value
Viewed Applications				
Bid Amount	12.59	12.35	-0.24 (0.34)	0.49
Profile Wage Rate	12.56	12.34	-0.23 (0.32)	0.47
Avg 6 Month Wage	10.62	10.23	-0.40 (0.28)	0.16
Min 6 Month Wage	8.66	8.27	-0.39 (0.25)	0.11
Max 6 Month Wage	13.24	13.03	-0.21 (0.37)	0.57
Previous Hours Worked	899.18	932.30	33.12 (31.20)	0.29
Prior Billed Openings	22.72	22.92	0.19 (0.66)	0.77
Avg Feedback	4.66	4.66	-0.00 (0.01)	0.97
Min Feedback	3.23	3.24	0.01 (0.02)	0.75
Max Feedback	4.96	4.96	-0.00 (0.00)	0.83

Notes: This table reports outcome means and standard errors across experimental groups of applicants who are viewed by employers. The treatment group are employers who are unable to observe past wage history information of the applicants. The unit of randomization was employer. The standard error for the difference in means is in parentheses next to the estimate. Reported p-values are for two-sided t-tests of the null hypothesis of no difference in means across groups. Standard errors are clustered at the employer level. Significance indicators: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : * * *$.

Table 2.6: Effect of Hiding Past Wages on Information Acquisition

	Dependent variable:			
	Number of Apps Messaged	Number of Apps Skyped	Number of Apps with "?"	Number of Apps with Question Words
	Poisson	Poisson	Poisson	Poisson
	(1)	(2)	(3)	(4)
Hidden Wages				
Treatment Group (TRT)	0.085*	0.078	0.114**	0.145***
	(0.051)	(0.059)	(0.048)	(0.051)
Constant	0.607***	−0.627**	0.010	0.028
	(0.176)	(0.260)	(0.208)	(0.224)
Opening Level Covariates	Yes	Yes	Yes	Yes
Observations	5,855	5,855	5,855	5,855

Notes: This table shows the relationship between measures of information acquisition and the treatment. The level of observation is the job posting. Estimates are from quasi-maximum likelihood Poisson models. Heteroskedasticity-robust standard errors are reported. The sample is restricted to hourly first job posts by an employer. The DV in Model (1) is a count of the number of applicants that where messaged. The DV in Model (2) is a count of the number of applicants messaged that included the word "Skype" in a message with the employer. The DV in Model (3) is a count of the number of applicants messaged that exchanged messages including a question mark with the employer. The DV in Model (4) is a count of the number of applicants messaged that included at least one of the following question words: *who, what, where, when, why, or how.* All models include covariates: category indicators, prior jobs billed by the employer, employer's prior spending on the platform, the number of applications to the job opening, the number of recommended applications to the job opening, the average applicants bid, and if the employer requested a specific skill. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table 2.7: Characteristics of Messaged Applicants

	Control mean: $\bar{X}_{CTL}$	Treatment mean: $\bar{X}_{TRT}$	Difference In Means	p-value	
Messaged Applicants					
Bid Amount	13.96	13.12	-0.85 (0.51)	0.10	*
Profile Wage Rate	13.49	12.74	-0.75 (0.44)	0.09	*
Avg 6 Month Wage	11.70	10.87	-0.82 (0.40)	0.04	**
Min 6 Month Wage	9.38	8.68	-0.70 (0.34)	0.04	**
Max 6 Month Wage	14.91	13.96	-0.96 (0.55)	0.08	*
Previous Hours Worked	1189.01	1167.15	-21.86 (56.59)	0.70	
Prior Billed Openings	30.63	29.02	-1.61 (1.25)	0.20	
Avg Feedback	4.71	4.71	-0.00 (0.01)	0.62	
Min Feedback	3.20	3.23	0.03 (0.04)	0.44	
Max Feedback	4.98	4.97	-0.01 (0.00)	0.18	

Notes: This table reports outcome means and standard errors across experimental groups of applicants who are messaged. The treatment group are employers who are unable to observe wage history information of the applicants. The unit of randomization was employer. The standard error for the difference in means is in parentheses next to the estimate. Reported p-values are for two-sided t-tests of the null hypothesis of no difference in means across groups. Standard errors are clustered at the employer level. Significance indicators: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : * * *$.

Table 2.8: Effect of Hiding Past Wages on Job Fill Rate

	Dependent variable:			
	I(Job Filled)			
	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Hidden Wages Treatment Group (TRT)	0.029**	0.027**	0.026**	0.023**
	(0.013)	(0.013)	(0.012)	(0.011)
Constant	0.403***	0.392***	0.404***	0.384***
	(0.009)	(0.061)	(0.009)	(0.055)
Job Order Dummy	No	No	Yes	Yes
Opening Level Covariates	No	Yes	No	Yes
Employer Level Covariates	No	Yes	No	Yes
Observations	5,922	5,855	9,476	8,973

Notes: The results are limited to hourly job posts. Columns (1) and (2) are limited to first job postings by employers. Columns (3) and (4) use the full sample with standard errors clustered at the employer level. Columns (2) and (4) add covariates including: category indicators, prior jobs billed by the employer, employer's prior spendings on the platform, the number of applications to the job opening, the number of recommended applications to the job opening, and the average applicant's bid. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table 2.9: Effect of Treatment on Hired Worker Wage Negotiation

	Dependent variable:
	log(Wage to Bid Ratio)
Hidden Wages Treatment Group (TRT)	0.012^{*}
	(0.006)
Constant	-0.027^{***}
	(0.005)
Observations	1,500

Notes: The sample is restricted to assignments originating from an hourly first job post by an employer who hire exactly one applicant. The top and bottom .5% of ratios were dropped. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table 2.10: Effect of Hiding Past Wages on Hourly Wage Paid

	Dependent variable:
	log(Wage Rate)
Hidden Wages Treatment Group (TRT)	-0.102^{**}
	(0.046)
Constant	2.165^{***}
	(0.033)
Observations	1,514

Notes: The sample is restricted to hourly job posts by an employer. The sample is further restricted to first job posts by an employer who hired exactly one applicant. Heteroskedasticity-robust standard errors are reported. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table 2.11: Effect of Treatment on Job Feedback Score

	Dependent variable:	
	log(Public Feedback)	log(Private Feedback)
	(1)	(2)
Hidden Wages Treatment Group (TRT)	0.012	0.056*
	(0.018)	(0.032)
Constant	1.427***	2.180***
	(0.058)	(0.067)
Employer Level Covariates	Yes	Yes
Observations	1,042	1,212

Notes: The sample is restricted to assignments originating from an hourly first job post by an employer. The sample includes only public job openings. The sample is further limited to include only the first job post spawned from each assignment. Covariates included are category indicators, the total value of the job, the total number of hours of the job, and the number of prior jobs the employer and worker completed. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : * * *$.

Table 2.12: Intensive Search By Experimental Period

	Dependent variable:		
	Number of Apps Skyped	Number of Apps with "?"	Number of Apps with Question Words
	(1)	(2)	(3)
Hidden Wages			
Treatment Group (TRT)	0.079	0.112**	0.142***
	(0.060)	(0.048)	(0.051)
Post Treatment Period	0.232***	0.169**	0.175**
	(0.077)	(0.074)	(0.077)
TRT x Post Period	−0.202*	−0.176*	−0.247**
	(0.109)	(0.100)	(0.109)
Constant	−0.782***	0.015	0.036
	(0.243)	(0.180)	(0.196)
Opening Level Covariates	Yes	Yes	Yes
Observations	7,920	7,920	7,920

Notes: Estimates are from quasi-maximum likelihood Poisson models. Heteroskedasticity-robust standard errors are reported. All models include covariates: category indicators, prior jobs billed by the employer, employer's prior spending on the platform, the number of applications to the job openings, the number of recommended applications to the job opening, the average applicant's bid, and a skill requested indicator. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : * * *$.

Table 2.13: Information Acquisition By Requested Contractor Tier

	Dependent variable:			
	Number of Apps Messaged Poisson	Number of Apps Skyped Poisson	Number of Apps with "?" Poisson	Number of Apps with Question Words Poisson
	(1)	(2)	(3)	(4)
Hidden Wages Treatment Group (TRT)	0.109	0.225*	0.184*	0.219**
	(0.112)	(0.121)	(0.097)	(0.101)
Intermediate Contractor	−0.186**	−0.019	−0.046	−0.091
	(0.086)	(0.096)	(0.077)	(0.077)
Expert Contractor	−0.032	0.130	0.025	0.009
	(0.108)	(0.114)	(0.096)	(0.104)
TRT x Intermediate	0.025	−0.153	−0.051	−0.049
	(0.127)	(0.144)	(0.117)	(0.120)
TRT x Expert	−0.190	−0.344**	−0.248*	−0.278*
	(0.153)	(0.169)	(0.136)	(0.149)
Constant	0.647***	−0.646**	0.034	0.085
	(0.207)	(0.273)	(0.226)	(0.239)
Opening Level Covariates	Yes	Yes	Yes	Yes
Observations	5,754	5,754	5,754	5,754

Notes: This table shows the relationship between measures of information acquisition and the treatment by requested contractor tier. The level of observation is the job posting, and the baseline group is beginner level contractors. Estimates are from quasi-maximum likelihood Poisson models. Heteroskedasticity-robust standard errors are reported. The sample is restricted to hourly first job posts by an employer. The DV in Model (1) is a count of the number of applicants messaged. The DV in Model (2) is a count of the number of applications that included the word "Skype" in a message with the employer. The DV in Model (3) is a count of the number of applications that exchanged messages including a question mark with the employer. All models include covariates: category indicators, prior jobs billed by the employer, employer's prior spending on the platform, the number of applications to the job openings, the number of recommended applications to the job opening, the average applicant's bid, and an indicator for requested skills. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : * * *$.

Table 2.14: Information Acquisition, Private Job Postings

	Dependent variable:			
	Number of Apps Messaged	Number of Apps Skyped	Number of Apps with "?"	Number of Apps with Question Words
	Poisson	Poisson	Poisson	Poisson
	(1)	(2)	(3)	(4)
Hidden Wages				
Treatment Group (TRT)	−0.186	−0.227	−0.224	−0.177
	(0.241)	(0.259)	(0.216)	(0.249)
Constant	−0.750	-1.897^{**}	−0.427	−0.623
	(0.521)	(0.783)	(0.741)	(0.776)
Opening Level Covariates	Yes	Yes	Yes	Yes
Observations	2,392	2,392	2,392	2,392

Notes: This table shows the relationship between measures of information acquisition and the treatment. The level of observation is the job posting. Estimates are from quasi-maximum likelihood Poisson models. Heteroskedasticity-robust standard errors are reported. The sample is restricted to hourly first job posts by an employer.The DV in Model (1) is a count of applicants messaged. The DV in Model (2) is a count of the number of applications that included the word "Skype" in a message with the employer. The DV in Model (3) is a count of the number of applications that exchanged messages including a question mark with the employer. The DV in Model (4) is a count of the number of applications that included at least one of the following question words: *who, what, where, when, why, or how.* All models include covariates: category indicators, prior jobs billed by the employer, employer's prior spending on the platform, the number of applications to the job opening, the number of recommended applications to the job opening, the number of applications that exchanged messaged with the employer, and the average applicant's bid. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Chapter 3

Guarantees and Online Labor

3.1 Introduction

Consumers are frequently uncertain about the quality of the product they are purchasing, and are thus unable to properly ascertain the product's worth. This uncertainty can cost firms potential sales, and, in some extreme circumstances, lead to a complete market failure. The economic rationale for the use of money-back guarantees by firms seems obvious. A money-back guarantee provides insurance against unsatisfactory product performance, and can simultaneously act as a signal of a high-quality product provided it is more costly to guarantee low-quality products than it is to guarantee high-quality products. Money-back guarantees have been a central component to firm strategies at firms such as L.L. Bean, Publix, Trader Joe's, Stew Leonard's, Costco, Aldi, and Nordstrom. L.L. Bean's 100 percent satisfaction guarantee has been part of the firm's strategy since 1912, when the firm refunded money for a hunting boot whose poor design lead to the boot's rubber bottom separating from the leather upper. Marketplaces such as Amazon.com or eBay.com have turned to money-back guarantees to combat informational problems, and increase the efficiency of trade.[1]

One example of a money back guarantee in a labor market is offered by the United States federal government. The federal government offers a money-back bonding program for hard-to-place employees. The program offers to repay up to six months of salary if a firm is not satisfied after hiring one of the program's employees.[2]. Theoretically, there is tremendous reason for either sellers of labor or other interested parties such as universities or labor marketplaces to provide money-back guarantees.[3] Universities compete to attract the most talented students. If a university were to offer recruiting firms a money-back guarantee covering 6 months of wages for graduates who have above a 3.5 GPA, this would lead firms,

---

[1]Online platforms often choose to use money back guarantees as opposed to upfront screening mechanisms as up front screening increases sign up costs, which many limit traffic to the site which is in direct opposition to platforms goal of growing volume quickly and reaching critical scale (Katz and Shapiro, 1994).

[2]http://www.bonds4jobs.com/program-background.html

[3]Some have argued that apprenticeships and internships are de facto guarantees by the novice worker.

for whom the cost of of firing is quite substantial (Dube et al., 2010; O'Connell and Kung, 2007), to prefer guaranteed applicants to equivalent non-guaranteed applicants. This would in turn lead to the best and brightest students selecting into attending the university that offered the money-back guarantee leading to a higher quality student base reducing the cost of actually paying the guarantee.[4]

Most empirical studies of money-back guarantees must evaluate the value of a money-back guarantee specifically on the firm that explicitly chooses to offer the guarantee, or focus on markets in which all firms are covered by a guarantee. The field experiment I study is akin to a 3rd party verifier, like *consumer reports*, evaluating the quality of each firm in a market and then offering a money-back guarantee on those firms that are above some threshold quality. While this situation seems quite contrived, it provides a fantastic laboratory to understand along what margins some firms offering of a money-back guarantee changes behavior of buyers.

Using a randomized experiment, I empirically investigate the effects of the adoption of a money-back guarantee by a two-sided marketplace for short-term labor. Specifically, I study the effects of a third-party platform offering a guarantee on only higher quality sellers of labor. This guarantee by a third party can both reduce the risk of transacting, and provide an endorsement effect. This paper looks at the causal effect of this bundle and traces out the effects on different margins of buyers. Finally, in a follow-up experiment, I attempt to separate the risk-shifting effects of a limited money-back guarantee with the signaling effects of offering the guarantee on only high-quality applicants.

In this paper, I describe an experiment in which a random group of employers were told that certain applicants—if hired—were guaranteed. If one of these randomly treated employer hired a MBG-eligible worker (a worker who would be eligible for the MBG since his quality is measured to be above some threshold), they could, at their discretion, receive a full refund for two weeks worth of hourly billings. The decision to guarantee an applicant was based on a machine learning model designed to predict contractors who were a "best match" for that employer.

I find that when a an MBG covering sellers of labor deemed to be above some quality threshold is introduced to a market for labor, buyers for whom quality (hiring high-expertise applicants) is most important are 3 percentage points more likely to locate and hire an applicant from a baseline of 40% fill rate. However, conditional on filling a job, the MBG does not influence these employers to substitute to candidates who are any different on observable characteristics. The introduction of the MBG causes employers who are looking at candidates just below the quality threshold to substitute to candidates who are just above the quality threshold covered by the guarantee. This leads to a 3% increase in feedback left by employers who substitute to candidates who are just above the threshold, but no observable change in feedback left by employers who specifically targeted higher expertise applicants. While the employers looking to hire expert talent do not report any higher satisfaction, the

[4]this example assumes away a competitive reaction by other universities. In actuality, the market implications of a money back guarantee are much more interesteding but beyond the scope of this paper.

increase along the extensive margin leads to a 3 percentage point increase in rehiring from a baseline of 14%.

This pattern of results is consistant with both a story of the guarantee acting as a de facto certification mechanism and simply helping employers to identify top talent, and a story of the guarantee acting as a risk-shifting mechanism, which lowers the risk envolved in hiring for employers by shifting some risk to the platform. To disentangle these two effects, I ran a follow-up experiment comparing the effects of marking applicants not as "guaranteed" but as "recommended." Guaranteeing applicants as opposed to simply signaling their quality by marking them as "recommended" does not have any effect on either the probability of filling a job posting or the likelihood of employers looking at candidates just below the quality threshold to substitute to candidates just above the quality threshold. This suggests that a large fraction of the MBG in the setting I study is really an endorsement effect, not the result of risk sharing. This work along with Barach (2015) suggests that information possessed by market platforms is valued by employers and shapes the matching process. The remainder of the chapter proceeds as follows: section 2 briefly reviews the relevant literature and describes this paper's contributions; section 3 describes the empirical context; section 4 discusses the experimental design; section 5 presents the empirical findings for the first experiment; section 6 presents the empirical findings for the second experiment; section 7 discusses the results; and section 8 concludes.

3.2 Literature Review

The literature on money-back guarantees has generally focused on answering one question: Why is it so common for sellers to provide guarantees of "satisfaction guaranteed or your money back"? In fact, McWilliams (2012) finds that nearly all the top online clothing retailers offer a fairly comprehensive money-back guarantee policy.

The literature has answered this question by highlighting two mechanisms through which MGBs can solve information problems: risk sharing and signaling.[5] The risk-sharing approach was pioneered by Heal (1977) while the literature focusing on signaling can be traced back to Spence (1977) and Grossman (1981).[6] Generally, high quality comes with a cost. To prevent low-quality firms from mimicking high-quality firms, high-quality firms need to move away from the first-best full information solution. According to the literature, this is why MBGs are so common. Shieh (1996) finds that MBGs and price together completely reveal a monopoly firm's private information about product quality. Moreover, the private infor-

[5]See Grossman and Hart (1980) and Jovanovic (1982) for early examples of the conditions under which truthful disclosure break down.

[6]While Spence (1977) and Grossman (1981) focused on warranties, there is a related literate that focuses more generally on the signaling of product quality. Milgrom and Roberts (1986), Bagwell and Riordan (1991), Bagwell (1992), and Daughety and Reinganum (1995) focus on price as a signal, while Kihlstrom and Riordan (1984) highlights the signaling aspects of advertising.

mation is revealed without incurring the cost of signaling, and the optimal level of monetary compensation specified by a guarantee is the price of the product.

Building off the literature focusing on why firms use guarantees, the literature has also explored which kinds of good and firms are more likely to be associated with a MBG. McWilliams (2012) shows that although much of the traditional theoretical literature on MBGs hypothesized that firms will only use a money back guarantee when the probability of a consumer returning a good is low, in actuality major retailers offer MBGs whether they have high or low return rates. Beales et al. (1981) and Wiener (1985) explain this empirical and theoretical disconnect by exploring how the signals of warranties can be misleading. Balachander (2001), for example, shows that in a product market where a new entrant competes with an established product, signaling behavior can lead to an outcome where the less reliable product carries the longer warranty.[7] Uncertainty is hypothesized to grow when buying by catalog or Internet. Thus, MBGs should have even greater effects for online marketplaces. Moorthy and Srinivasan (1995) and others note that MBGs can be particularly relevant for experience goods. To ascertain fit, which is independent of quality, consumers must use products to learn if the product fits their own idiosyncratic needs. This is because product testing before purchase is inherently limited since personal fit uncertainty remains at the time of purchase. Although the literature theorizes that MBGs are particularly relevant for experience goods, there is little evidence of guarantees in the context of a market for labor, which is a particular type of experience good.

Since the role of a warranty as providing insurance against unsatisfactory product performance is obvious, the literature has traditionally focused on the use of warranties as signals. However, the empirical evidence on consumer product warranties appears to be inconsistent with the use of warranties as signals. Bryant and Gerner (1978), Garvin (1983), Gerner and Bryant (1981), and Priest (1981) all report that most warranties offer only partial insurance, and high-quality products are not always (or even usually) sold with higher quality items. Grossman (1981) and Lutz (1989) present models where warranties act as both signals and insurance and show that pooling equilibriums exist where all products are under warranty regardless of their quality. But little empirical work has attempted to separate the signaling effects of a guarantee from the risk reducing effects, and detail situations where one mechanism might dominate another. This paper adds to the literature by attempting to empirically isolate the risk reducing effects of a money-back guarantee by comparing the effect of marking high-quality applicants as "recommended" as opposed to "guaranteed."

There has been some research focusing on the oDesk marketplace. Much of this research has highlighted that in these sorts of marketplaces, one roadblock to trade is asymmetric information. One way to resolve asymmetric information is to develop a reputation (which is extremely difficult). **pallais2010inefficient** shows that the market-level benefits that come from hiring inexperienced workers and generating information about their abilities is great enough to overcome the costs of hiring these workers. Thus, it might be in a platform's

[7]Other examples of misleading signals include Bar-Isaac et al. (2010), which demonstrates that if quality is multidimensional and only some dimensions are disclosed, firms may shirk on unreported quality.

interest to generate information via paying to develop inexperienced worker's reputations. Another, much cheaper way to overcome asymmetric information in this type of market environment may be to use a money-back guarantee.[8]

3.3 Empirical Context

During the last ten years, a number of online labor markets have emerged. In these markets, firms hire workers to perform tasks that can be done remotely, such as computer programming, graphic design, data entry, and writing. The markets differ in their scope and focus, but common services provided by the platforms include maintaining job listings, hosting user profile pages, arbitrating disputes, certifying worker skills, and maintaining reputation systems. On oDesk, would-be employers write job descriptions, self-categorize the nature of the work and required skills, and then post the vacancies to the oDesk website. Workers learn about vacancies via electronic searches or email notifications.

Workers submit applications, which generally include a wage bid (for hourly jobs) or a total project bid (for fixed-price jobs) and a cover letter. In addition to worker-initiated applications, employers can also search worker profiles and invite workers to apply. After a worker submits an application, the employer can interview and hire the applicant on the terms proposed by the worker or make a counteroffer, which the worker can counter, and so on. The process is not an auction and neither the employer nor the worker are bound to accept an offer.

To work on hourly oDesk contracts, workers must install custom tracking software on their computers. The tracking software, or Work Diary, essentially serves as a digital punch clock that allows for remote monitoring of employees. When the worker is working, the software logs the count of keystrokes and mouse movements; at random intervals, the software also captures an image of the worker's computer screen. All of this captured data is sent to the oDesk servers and then made available to the employer for inspection. This monitoring makes hourly contracts and hence employment relationships possible, which in turn makes the oDesk marketplace more like a traditional labor market than project-based online marketplaces where contracts are usually arm's-length and fixed price.

In the first quarter of 2012, $78 million was spent on oDesk. The 2011 wage bill was $225 million, representing 90% year-on-year growth from 2010. As of October 2012, more than 495,000 employers and 2.5 million workers had created profiles (though a considerably smaller fraction are active on the site). Approximately 790,000 vacancies were posted in the first half of 2012.[9]

[8] Other research that details the importance of reputation and information on oDesk include Stanton and Thomas (2012a) which uses oDesk data to show that agencies (which act as quasi-firms) help workers find jobs and break into the marketplace. Additionally, Agrawal, Lacetera, et al. (2013) investigates what factors matter to employers in making selections from an applicant pool and present some evidence of statistical discrimination, which can be ameliorated by better information.

[9] See Agrawal, Horton, et al. (2013) for additional descriptive statistics on oDesk.

3.4 Experimental Design

On oDesk, jobs are posted and workers apply to those jobs. For each applicant, oDesk computes a “BestMatch” $\in [0, 1]$ prediction about that worker’s “quality” using a proprietary machine learning algorithm trained on historical oDesk data. Workers with a predicted quality above a certain threshold (0.5) were eligible to be marked as guaranteed. Figure 3.1 shows the distribution of BestMatch scores for eligible and non-eligible workers.

When employers posted a job, they were randomized to either a control group—in which the MBG-eligible applicants were not marked as special in any way—or a treatment group, in which MBG-eligible workers were marked with a “guaranteed” badge in the employer’s list of applicants. Table 3.1 contains details showing that the randomization was effective. The experiment started in November 2013 and ended in August 2014. Figure 3.2 shows the interface presented to employers in either the treatment group (left panel) or the control group (right panel). Figure 3.3 shows how an MBG-eligible worker would appear to an employer in the treated group when the employer moused over the MBG indicator.

Terms of the Money-back Guarantee

After making a hire, employers had two weeks to request a refund from oDesk. During the course of the experiment, oDesk refunded $110,000 to about 600 users resulting from the money-back guarantee. The minimum refund was 56 cents and the maximum was $2,700.

Experiment Design for Guaranteed Versus Recommended Badge

Following the main experiment, a second experiment compared the effects of “guaranteeing” workers with the money-back guarantee from the initial experiment with merely “recommending” them, without the money-back guarantee (and the risk-shifting properties it entails). The rules for determining an applicant’s eligibility for being guaranteed remained the same. Clients were randomized into two groups: those who saw MBG-eligible workers as guaranteed, and those who saw them as recommended. In both cells, the initial view of the user interface presented to employers changed from the one presented in the initial experiment. In the initial experiment, employers in the treatment group saw both guaranteed and non-guaranteed workers in the default view of the user interface. In the second experiment, employers only saw MBG-eligible workers in the initial view of the user interface, but could view all workers by selecting a different view. Figure 3.4 shows the interface for the two different cells of the experiment.

3.5 Empirical Results

The money-back guarantee studied in this experiment is algorithmically calculated to cover only providers of labor who are above some threshold of quality. Thus, similar to a firm in a marketplace that decides to offer an MBG, the MBG on labor in this online labor marketplace serves two purposes. It first serves to reduce the risk associated with purchasing a product from the guaranteed firm. Secondly, it serves as a certification or costly signal of which firms are purveyors of higher quality goods. Either by reducing risk associated with a purchase, or making it easier for buyers to identify high-quality sellers, the MBG is expected to cause buyers who may have previously been on the margin when it came to making a purchase to now purchase the good. I first explore along what margins the MBG increased the probability of a buyer of labor making a hire.

Effects on the Probability of the Employer Filling the Job

In table 3.2, the dependent variable is whether at least \$0 was spent by the employer on one or more hired workers. The sample is restricted to hires for non-private hourly jobs. The sample is restricted to the first job post by an employer following their allocation to the experiment. The model is fit using OLS, with the single independent variable being the treatment indicator:

$$\text{Filled} = \beta_0 + \beta_1 \text{TRT} + \epsilon_{ij}. \tag{3.1}$$

The treatment indicator $\hat{\beta}_1$ in table 3.2 is clearly not significantly different than zero. Thus, the introduction of a money-back guarantee by a third party that covers all firms above an acceptable quality threshold does not seem to increase market participation on average. However, the value of the MBG is not equal to prospective buyers of labor in the marketplace. Hiring applicants for positions that require high-expertise workers is riskier than hiring applicants for positions that require low-expertise workers. Employers must pay high-expertise applicants more money, so they lose more money if the worker turns out to be of low quality, but more importantly, higher expertise roles are usually harder to monitor and are much more vital to the success of the firm than low-expertise roles. If a firm hires a data entry technician who turns out to be bad at entering data, the downside risk of missing deadlines and lost reputation is much less than for a marketing manager who interacts with downstream clients. Additionally, identifying high-quality applicants is in itself harder. Holding the risk of hiring fixed across all jobs, a higher proportion of the available applicants will be able to complete a less vital task such as data entry. Thus, even if the MBG only acts as a signal, it is more useful to firms looking to identify high-expertise applicants.

I therefore expect that both the signaling effects as well as the risk-reducing effects of the MBG to be more likely to bind for firms seeking to hire high-expertise applicants. Table 3.3, shows that the money-back guarantee increases the probability of an employer filling a job posting for employers interested in hiring high-expertise applicants. The existence of the money-back guarantee on high-quality applicants increases the probability of filling a job by about 3 percentage points for employers looking to hire high-expertise applicants. This indicates that the MBG is effective at increasing the probability of employers hiring an applicant only at the margin for employers who are either taking on more risk, or for whom a signal of a high-quality applicant quality is most important. This result is consistent with the MBG acting as a signal and indicating which applicants are worth hiring. The result is also consistent with the MBG reducing the risk of hiring applicants.

To further understand exactly which employers are most affected by the MBG, I next seek to identify differential effects by the employers' experience on the platform. Employers who have previous experience using the oDesk platform better understand both the full risks associated with hiring task based labor, and the difficulty associated with locating the high-quality talent. Thus there is reason to expect experienced employers to be more affected by the MBG. On the other hand, those employers who have previously hired labor on oDesk must by definition be willing to accept the risk associated with hiring online labor. Thus, if the MBG either reduces the risk threshold slightly, or provides a slightly more salient signal of quality, it might at the margin cause employers who previously were not willing to hire via an online labor market to now transact in the marketplace.

Turning to table 3.4, the treatment indicator is significant. This indicates that the effect of the MBG on experienced employers is significantly different from zero. While the coefficient of the interaction term is negative, it is not significant, indicating that the treatment effect on a first time employer is not significantly different from the effect on more experienced employers. Thus, while the MBG does seem to slightly increase experienced employers' likelihood of filling a job post, I cannot conclude that the effect is significantly different for experienced employers compared to first time users of oDesk.

Effects on the Hired Worker Composition

The money-back guarantee might not only alter matching behavior along the extensive margin by causing some marginal employers to participate in the labor market, but also alter the characteristics of those applicants the employers choose to employ. I begin by considering whether the treatment was effective at increasing the number of MBG-eligible workers that are hired, and increasing the overall ex ante quality of applicants being selected by employers. To do this, we prepare a sample of all workers hired by an employer on the employer's first job opening if that job opening had at least one applicant who was eligible for treatment, regardless of treatment assignment. We condition these models on hired applicants, as we are interested in understanding how having the option of hiring applicants by the money-back guarantee change the characteristics of the applicant(s) selected.

In model (1) of table 3.5 we observe whether MBG-eligible workers are overrepresented among hires of treated employers. To do this, we can regress a dummy for the worker being MBG-eligible, MBGEligible, on the treatment indicator, where the treatment indicator is for the applied-to opening:

$$\text{MBGEligible} = \beta_0 + \beta_1 \text{TRT} + \epsilon | HiredApplicant. \tag{3.2}$$

In model (3) of table 3.5 we observe to what extent the average BestMatch score of the hire applicants increased. This provides us with a measure of the extent to which the average algorithmically calculated ex ante quality of hired labor increased due to the MBG incentivizing employers to select applicants eligible for the guarantee and thus are of higher ex ante observable quality.

Table 3.5, column (1) shows that workers hired by employers in the treatment group were substantially more likely to be MBG-eligible.

Theoretically, the predicted direction of the result of model (3) is ambiguous. Employers might substitute from applicants whose BestMatch scores are just below the threshold to applicants just above the threshold, increasing the average BestMatch score of hired applicants. However, the money-back guarantee could have an adverse impact on the average BestMatch score of hired applicants as employers who previously would have hired only the most experienced (and highest BestMatch score) applicant might be willing to substitute to a much less experienced (and lower BestMatch score) applicant who is still above the threshold.

Table 3.5, column (3) shows that average measured quality of workers hired by employers in the treatment group increases by only about 1%.

As theorized above, it is possible that the effects of the MBG on the hired applicants BestMatch could be nonlinear. Table 3.5, column (2) and column (4) look for heterogeneous effects by the type of applicant the employer reports he or she is looking for. The results in both column (2) and column (4) show that there is no statistical difference in the propensity to hire higher BestMatch applicants conditional on hiring an applicant for employers looking to hire higher expertise applicants. To get a better picture of the distributional effects of the money-back guarantee on the BestMatch score of hired applicants, figure 3.5 compares the distribution of best-match scores of hired applicants between treatment and control employers. This figure seems to indicate that a majority of the action is coming from substitutions from employers who previously hired applicants just below the threshold to employers now hiring applicants just above the MBG-eligible threshold.

To further explore how the money-back guarantee affects the composition of applicants being hired, table 3.6 looks at how seeing certain applicants covered by an MBG alters the characteristics of applicants chosen by the employer. Recall that on average, MBG-eligible applicants generally have more prior hours worked, more prior jobs on the platform, more earnings, and charge a higher rate. The applicants hired by employers who see MBG-backed applicants have 6.7% more previous hours worked, and 3.6% more jobs worked, and 6.9% more earnings (this result is borderline insignificant). There does not seem to be an increase in the bid of the applicant selected. This result is surprising as theory might predict that the MBG makes buyers less price sensitive, but this does not seem to be the case at all.

Hired, MBG-eligible Worker Attributes

The treatment makes MBG-eligible applicants both more salient and more attractive to employers, and thus, treated employers are more likely to hire a MBG-eligible worker than non-treated employers. As these MBG-eligible workers were chosen to be of higher quality, they tend to have greater experience, higher hourly wages, better BestMatch scores and so on—table B.1 shows this. However, since the treatment shifts default risk from the employer to oDesk, *conditional* upon choosing to hire a MBG-eligible worker, the employer might decide to hire an MBG-eligible worker who is riskier, as the employer might decide to take a shot on a riskier applicant and should that worker not work out, fire him or her and hire a less risky applicant. To investigate this issue, table 3.7 looks to see, *conditional* on hiring a MBG-eligible applicant, how applicants hired by employers who saw and hired MBG-eligible applicants that were marked as backed by the guarantee differ from applicants hired by employers who hired MBG-eligible applicants that were not marked in any special way. There does not seem to be any evidence of, *conditional* on hiring, MBG-eligible treated employers hiring applicants who look any different from applicants hired by control employers.

The Effects of the MBG on Job Outcomes

Table 3.8 investigates if the substitution effect of the money-back guarantee which leads employers to hire slightly more experienced and higher quality applicants, at least as measured by oDesk's BestMatch score, leads to employers having better job outcomes. While there is no effect on the public feedback left, employees who saw and hired MBG-eligible applicants that were marked as backed by a money-back guarantee leave about 1% higher private feedback on average. Also, while the result is borderline insignificant, these treated employers are also about .6 percentage points from a baseline of 14% more likely to rehire an applicant. These results can be thought of as limited evidence that the selection effect of the MBG led to hiring an applicant who is a better employee.

It is possible that the treatment itself causes employers to perceive the labor as higher quality simply because its backed by a money back guarantee. To address this issue, table 3.9 runs the same models as in table 3.8 limiting the sample to jobs that hired a MBG-eligible applicant. Table 3.9 shows that *conditional* on hiring an MBG-eligible applicant, having that applicant backed by a guarantee did not affect the feedback left by the employer or the probability of rehiring a worker. This indicates that the effect is most likely not driven by employers simply perceiving that the quality of labor is better, but by actually hiring higher quality labor.

The expertise of the applicant the employer is looking to hire has been show to affect the probability of filling a job posting. Table 3.10 looks at the effect on private feedback (model (1)) and probability of rehire (model (2)) left by the employer based on the type of applicant the employer is looking to hire. The results indicate that the increase in private feedback is driven by employers hiring low-expertise applicants, but that the increase in rehires is driven by employers hiring high-expertise applicants. This makes sense since the substitution effect of hiring slightly better applicants (at least based on ex ante characteristics) is limited to substitutions from hiring applicants just below the quality threshold to applicants just above the quality threshold. The MBG increases the probability of filling a job only for employers interested in hiring high-expertise applicants, and these employers who previously did not transact in the market are the drivers of the increase in rehiring. Conditional on hiring, the employers looking for high-expertise talent are more likely to hire, but they do not hire observably different applicants and do not get differential outcomes. The changes in the demographics and outcomes are driven by employers interested in hiring marginal applicants, who substitute to a slightly less marginal applicant who is covered by the money-back guarantee.

Turning to model (2), while treated employers interested in hiring low-expertise applicants leave higher feedback, it appears that treated employers interested in hiring high-expertise applicants are more likely to rehire an applicant. One possible reason for this interesting contradiction is that the types of jobs filled by employers looking to hire low-expertise applicants are hiring them for jobs where there is no possibility of rehire.

3.6 Experiment 2 (Guaranteed vs. Recommended) Empirical Results

We next turn to the second experiment (as described in section 3.4), where for a random sample of employers, oDesk recommends workers rather than guaranteeing them. This follow-up experiment allows us to further distinguish between the two possible explanations for the effects of the money-back guarantee whether: (1) the guarantee's value to employers is primarily due to its insurance value, or (2) its value is as a signal of latent worker quality (or possibly just salience of the applicant).

Probability of filling a job

Table 3.11 shows no significant difference between the probability of filling a job post for employers who saw MBG-eligible workers marked as "guaranteed" and employers who saw MBG-eligible applicants marked as "recommended." Cutting the data, as we did in the first experiment, by desired expertise level of the hired applicant, table 3.12 shows that there are no differential effects by type of applicant desired. However, here the point estimates are large and imprecisely estimated, especially for employers interested in hiring high-expertise applicants. This makes interpreting the heterogeneous effects of the MBG on the extensive margin difficult. However, the point estimates from table 3.11 are close to zero, indicating that the effect of the MBG on the extensive margin – the probability of filling a job posting – is most likely driven by signaling and not by reducing the risk involved in transacting on oDesk.com.

The user interface in the second experiment was slightly different from that in the first experiment. In the second experiment, employers in both the treatment group and the control group by default only saw applicants whose BestMatch score was above the MBG-eligible threshold. In the second experiment, only MBG-eligible applicants are now shown in the default-view. In the default view, applicants were all marked as "guaranteed" for treated employers and "recommended" for control employers. Non MBG-eligible applicants could now only be viewed by clicking to view all applicants. While the results are not comparable with those from the first experiment, the behavior of treatment employers can be compared to the behavior of control employers in the second experiment isolating the risk reducing effect of the money back guarantee. To some extent, we worry that the guarantee is not implicitly a recommendation, or that the signal sent by being marked as "guaranteed" is different that the signal sent by being marked as "recommended." There is no way we can truly be sure that the signaling component of marking an applicant as "guaranteed" is similar to that of marking an applicant as "recommended."

Table 3.13 only allows me to conclude that the employers on oDesk are not more likely to hire applicants from the default "guaranteed" list than to hire applicants from the default "recommended" list. Thus, table 3.11 and table 3.13 indicate that oDesk can replicate the positive effects of the MBG without having to take on the risk associated with the guarantee. Furthermore, if I believe that the recommendation is similar to a guarantee not backed up by any costly investment or risk-shifting, then this is evidence that the MBG increases the probability of transacting in the marketplace and the selection of higher quality applicants not by shifting risk but only through signaling to employers which applicants are of higher latent quality.

If I assume that the "Recommendation" indication is similar to the "Guaranteed" indication except that "Guaranteed" not only indicates which applicants are of higher quality, but also shifts some of the risk involved in hiring to the platform, then we can make use of the second experiment to separate the risk-shifting effect of the MBG from the signaling effect. The results seem to indicate that the risk-reducing effects of the MBG do not affect selection dramatically either on the extensive margin, by causing more employers to transact on oDesk, or on the intensive margin, by causing employers to dramatically substitute from non MBG-eligible applicants to much higher quality MBG-eligible applicants. Thus, it is not surprising that table 3.14 indicates that the risk-reducing effects of the MBG as opposed to the signaling effects do not have any effect on job outcomes.

3.7 Discussion

I find that when a money-back guarantee covering sellers of labor deemed to be above some quality threshold is introduced to a market for labor, buyers for whom quality (i.e. hiring high-expertise applicants) is most important are 3 percentage points more likely to locate and hire an applicant from a baseline of 40% fill rate. However, conditional on filling a job, the guarantee does not influence these employers to substitute to candidates who are any different on observable characteristics.

The introduction of the money-back guarantee causes employers who are looking at candidates just below the quality threshold to substitute to candidates just above that threshold and are covered by the guarantee. This leads to a 3% increase in feedback left by employers who substitute to candidates just above the threshold, but no observable change in feedback left by employers who specifically targeted higher expertise applicants. While the employers looking to hire high-expertise talent do not report any higher satisfaction, the increase along the extensive margin leads to a 3 percentage point increase in rehiring from a baseline of 14%.

This pattern of results is consistent with both a story of the MBG acting as a de facto certification mechanism and simply helping employers to identify top talent, and a story of the MBG acting as a risk-shifting mechanism, which lowers the risk involved in hiring for employers by shifting some risk to the platform. To disentangle these two effects, I ran a follow-up experiment comparing the effects of marking applicants not as "Guaranteed" but as "recommended." Guaranteeing applicants, as opposed to simply signaling their quality by marking them as "Recommended," does not have any effect on either the probability of filling a job posting or the likelihood that employers looking at candidates just below the quality threshold substitute to candidates just above that threshold.

The evidence suggests that the results found in the first experiment might be mostly driven by signaling effects of the MBG as opposed to the risk-shifting effects of the MBG. To help verify these results, I conducted a survey where clients were asked, "Please share with us any thoughts you have on the Money-Back Guarantee." While the reaction to guarantee was generally positive, with 52% of the three hundred fifty two clients polled stating they "Loved it," the comments by clients who supported the guarantee indicated that they mainly found value in the guarantee as a signal of quality. For example:

> *"Its a great concept. I can easily find out best freelancer. But i am confused about Hidden folder. I saw lot of qualified freelancer are there. I think i would be better to Leave Hidden folder system."*
> – client 3282708806

> *"It's nice to know that this is available. It doesn't substitute for doing a good screening and interview but it feels like it helps."*
> – client 3515664247

> *"Excellent idea. Why? Because the "rating" system that is in operation is (in my opinion) not very "reliable". Contractors really "beg" not to have "bad" comments/ reviews. Clients may not want to alienate contractors (because the client may need to ask the contractor for help on something they've done in the past - even if it was bad/ frustrating etc)."*
> – client 3664891102

3.8 Conclusion

I ran a series of two randomized experiments set in the context of an online labor platform to try and shed light on an inconsistency between the theoretical and empirical literature on warrenties and money-back guarantees. The theoretical literature on money-back guarantees suggests that they can be particularly relevant for experience goods (Moorthy and Srinivasan, 1995). However, the empirical literature has not fully detailed the effects of guarantees in markets for experience goods.

My findings highlight that while a money-back guarantee did have a substantial effect on participation in the labor market, the characteristics of the employees hired, and job outcomes, the effects are most likely driven by signaling and not the the risk-shifting effects of the money-back guarantee. Labor markets have long been filled with signals of quality such as education and third-party certifications, which can achieve the same results based on reputation as a money-back guarantee backed up by cash. Additionally, this paper highlights the heterogeneous effects of the money-back guarantee. On the extensive margin, a money-back guarantee affects a very different subset of the population than it does on the intensive margin.

It is important to note the limitations of the analysis in this paper. One very important limitation is that all firms in this market are hiring task-based labor. Thus, incentives to locate a top-notch employee are lower than in the traditional labor market, since both the expected tenure of an employee and the firing costs are much lower than in the traditional labor market. Additionally, the second experiment does not allow for a perfect seperation of the risk-reducing effects from the signaling effects, as there is a concern that the signaling value of a guarantee (which is backed by money) differs from that of a recommendation (which is only backed by reputation of the platform). As the reputational concerns of the platform are extremely high, I do not believe that this is of primary concern. More importantly, I would also argue that any additional strength in the signaling value of the money-back guarantee due to its being backed by money is not in actuality signaling value, but part of the reduced risk.

My results are important to understanding the full risks associated with hiring labor. One reason for the lack of a risk-shifting effect in the context of a labor market is that calculating the full risks of a new hire is extremely difficult, and insuring against these risks (assuming they can be calculated) would be extremely costly. If a firm mistakenly hires a bad applicant, the costs are not limited to lost wages. In fact, the cost of lost wages is very small compared to the effects of possibly missing deadlines, which can have enormous implications for both upstream and downstream firms in the supply chain.

The results are also important for understanding the effects of money-back guarantees on markets more generally. To my knowlage, this is the first paper to study the introduction of a money-back guarantee that does not unilaterally affect the entire market. The money-back guarantees that have perviously been studied in the context of platforms such as Amazon.com or eBay.com are universal in that they cover all sellers doing business on the platform. In general, all firms in a market do not offer a money-back guarantee. Thus, a money-back guarantee not only reduces risk, but also helps to signal high-quality providers. The guarantee studied in this context is unique in that it focuses on a situation where only some firms are covered by the guarantee, which is more like what happens in a non-managed market. This unique situation allows for the study of the effects of a money-back guarantee on a market in a way that allows us to unbundle the full effects of a money-back guarantee and understand the full market effects of some higher quality firms offering this guarantee.

Figure 3.1: BestMatch Scores of Hired Workers, By MBG-eligibility

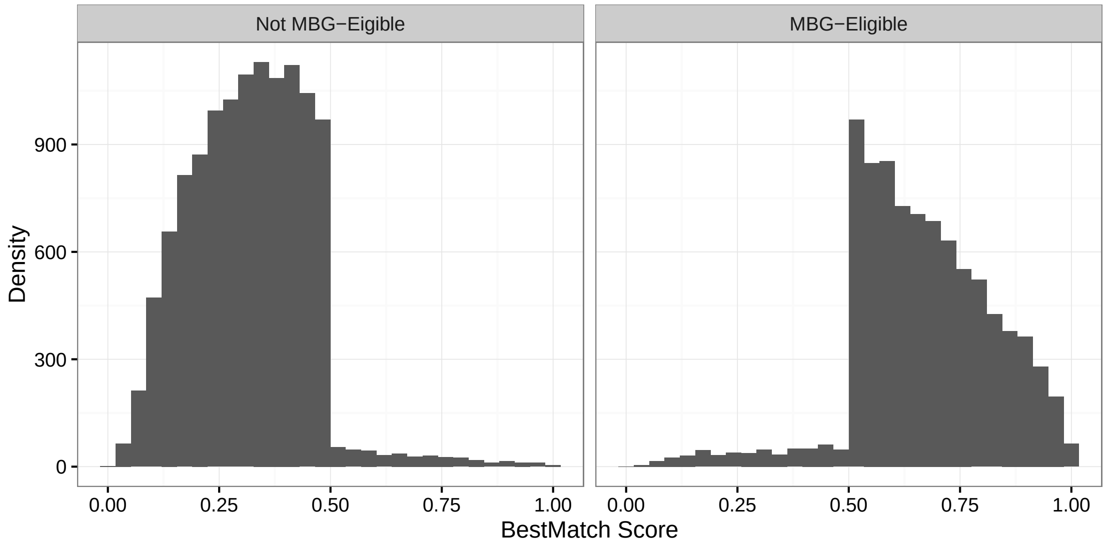

Note: Distribution of "BestMatch" scores by recommended and non-recommended hired workers.

Figure 3.2: Comparison of Employer Interface in Treatment (left) and Control (right)

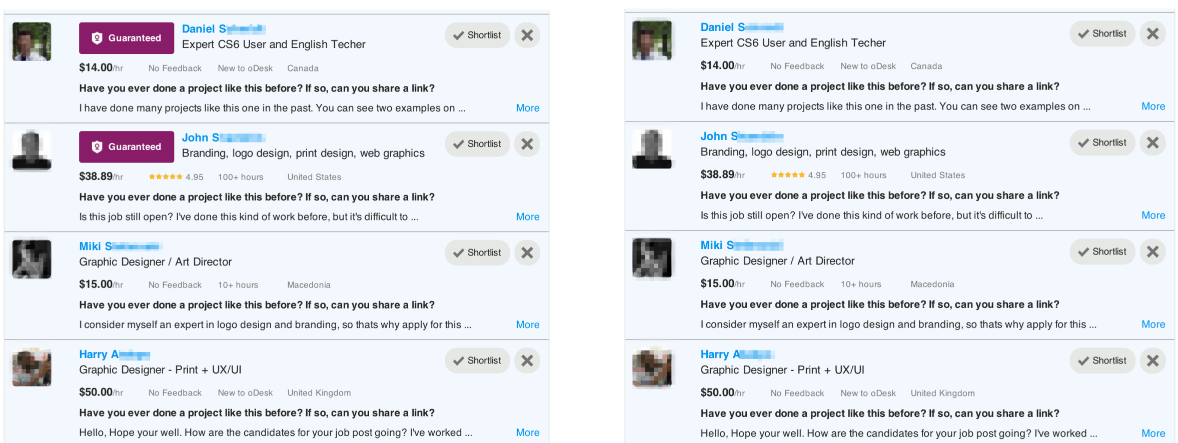

Notes: This figure shows the interfaces presented to employers. The left panel shows the interface presented to employers in the control group, whereas the right panel shows the interface for employers in the treatment group.

Figure 3.3: Mouse-over Description of the Money-back Guarantee

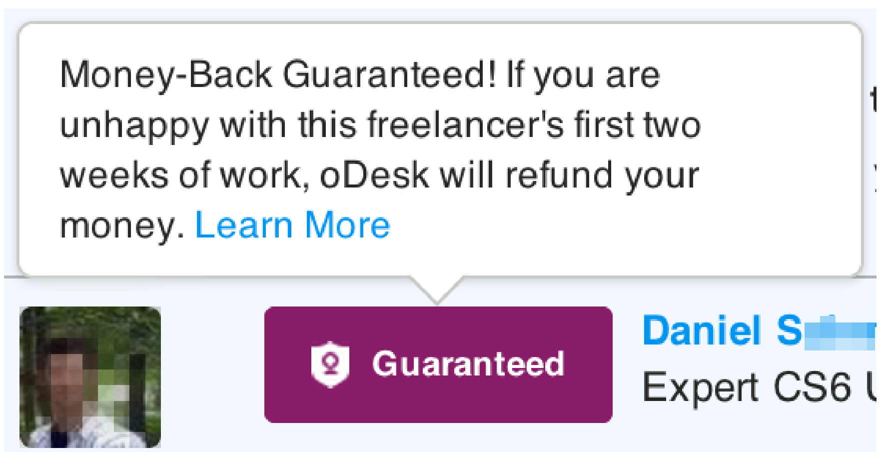

Notes: This figure shows the money-back guarantee label applied to MBG-eligible worker applicants in the treated group. It also shows the mouse-over description of the money-back guarantee program.

Figure 3.4: Comparison of Employer UI for Treatment (top) and Control (bottom) in the Second Experiment

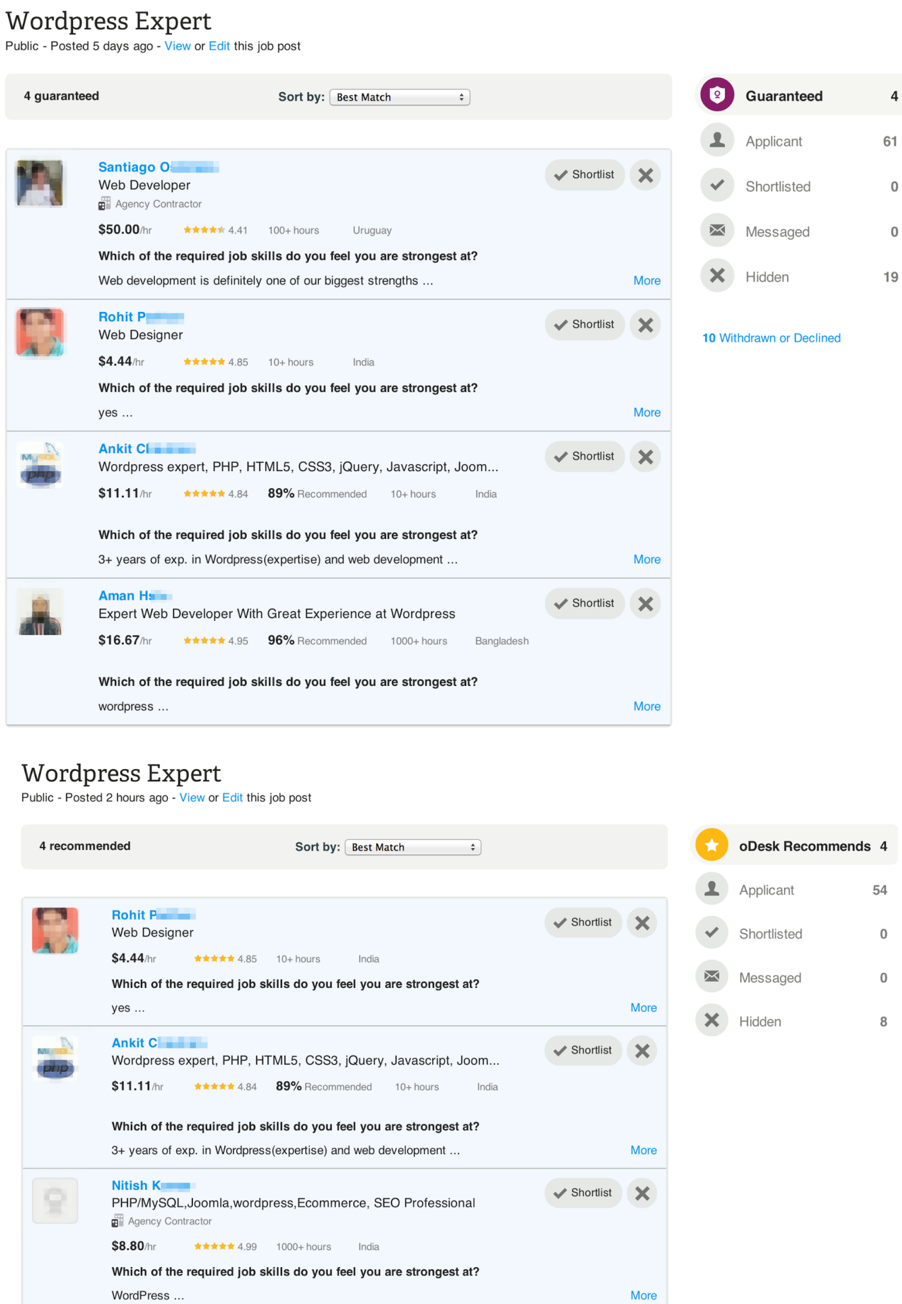

Notes: This figure shows the interfaces presented to employers in experiment 2. The left panel shows the interface presented to employers in the guaranteed cell, whereas the right panel shows the interface for employers in the recommended cell. Listed applicants are slightly different as treatment and control screenshots where taken at different times.

Figure 3.5: Distribution of BestMatch Scores of Hired Workers by Employer Treatment Group

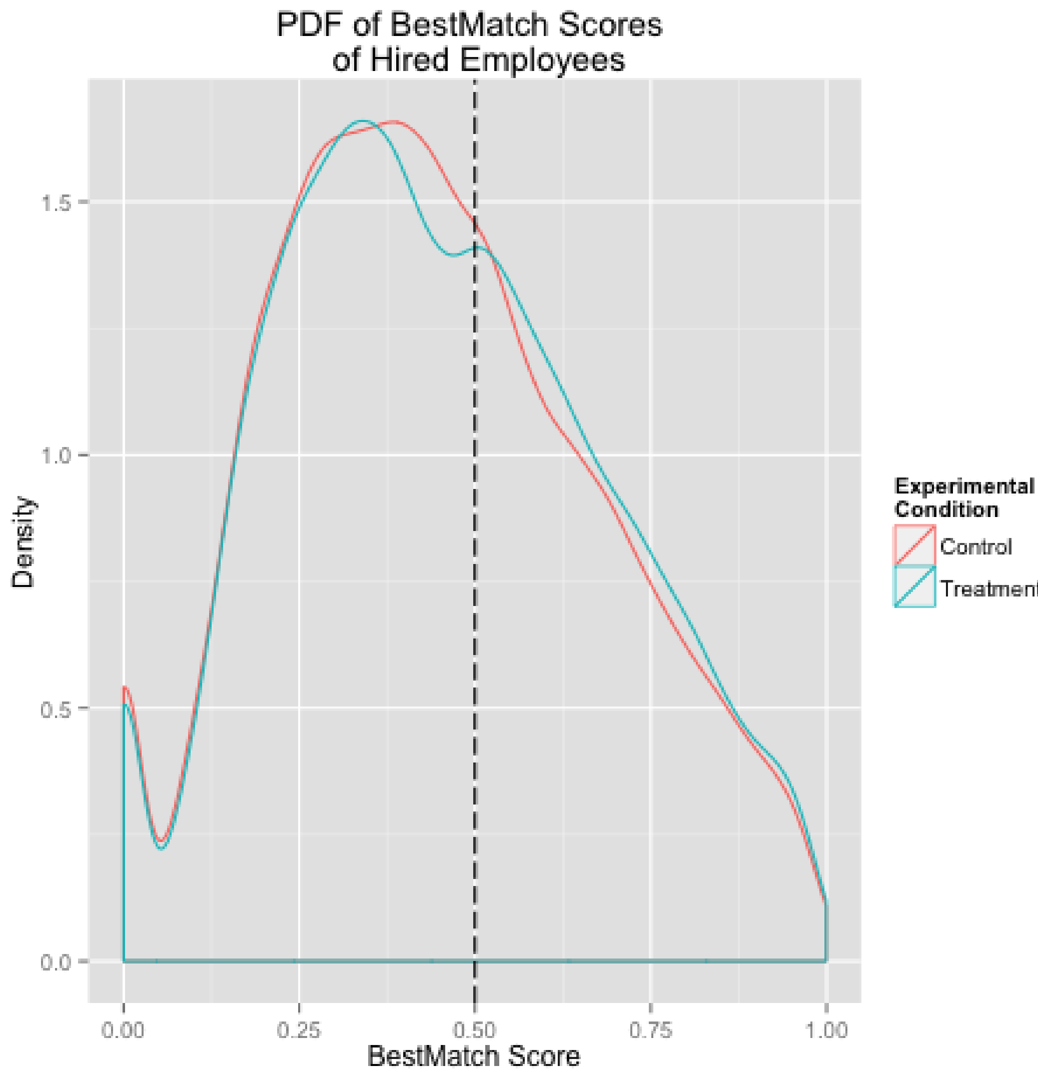

Note: Distribution of "BestMatch" scores for hired applicants by treatment and non-treatment employers.

Table 3.1: Balance Table of Employer, Job Posting, and Applicant Characteristics

	Control Mean: $\bar{X}_{CTL}$	Treatment Mean: $\bar{X}_{TRT}$	Difference In Means	p-value	
Employer Attributes					
Prior Job Postings	7.56	7.51	-0.05 (0.21)	0.82	
Prior Billed Jobs	3.25	3.20	-0.05 (0.12)	0.68	
Prior Spend by Employers	2867.08	2970.86	103.78 (247.97)	0.68	
Num Prior Contractors	3.28	3.26	-0.02 (0.12)	0.87	
Avg Feedback Score of Employer	4.80	4.79	-0.01 (0.01)	0.40	
Num of Reviews of Employer	2.34	2.34	-0.01 (0.10)	0.95	
Job Posting Attributes					
Number Non-invited Applicants	25.22	25.43	0.21 (0.33)	0.52	
Avg Best Match Score	0.36	0.36	0.00 (0.00)	0.08	*
Avg Bid	13.09	13.22	0.13 (0.11)	0.26	
Prefered Experience in Hours	31.53	30.03	-1.50 (1.16)	0.19	
Estimated Job Duration in Weeks	15.52	15.38	-0.14 (0.19)	0.48	
Applicant Attributes					
Hours Worked to Date	703.51	709.83	6.33 (8.19)	0.44	
Num Past Jobs Worked	15.07	14.94	-0.13 (0.15)	0.38	
Past Hourly Earnings	6026.28	6019.07	-7.22 (84.43)	0.93	
Num Prior Employers	12.05	11.97	-0.08 (0.12)	0.51	
Min Feedback Rating	0.35	0.35	-0.00 (0.00)	0.39	
Wage Bid	10.01	9.87	-0.14 (0.14)	0.31	
Profile Wage	9.96	9.83	-0.12 (0.11)	0.27	

Notes: This table reports means and standard errors across experimental groups of employer, job posting, and applicant characteristics. The top panel reports characteristics of employers allocated to treatment and control. The middle panel reports characteristics of job postings by treatment and control groups for the first job posted after allocation to the experiment for each employer. The bottom panel reports characteristics of employers at the time they were allocated to treatment or control groups. Reported p-values are the for two-sided t-tests of the null hypothesis of no difference in means across groups. In the bottom panel, standard errors are clustered at the employer level. Significance indicators: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : * * *$.

Table 3.2: Effect of MBG on Job Fill Rate

	Dependent variable:
	Fill Rate
MBG Treatment Group (TRT)	0.008
	(0.006)
Constant	0.344***
	(0.004)
Observations	28,975

Notes: The dependent variable is whether the employer filled their opening. In column (1) the sample is restricted to the first job post by the employer to have at least one MBG-eligible applicant. The estimation method is OLS.

Table 3.3: Effect of MBG on Job Fill Rate by Desired Expertise Level of Employee

	Dependent variable:
	Fill Rate
TRT	−0.007
	(0.010)
Intermediate Experience	−0.022**
	(0.009)
High Experience	−0.090***
	(0.010)
TRT × Intermediate Experience	0.008
	(0.012)
TRT × High Expereince	0.041***
	(0.015)
Constant	0.394***
	(0.007)
Observations	36,264

Notes: The omitted group contains employers who were in the experimental control group and indicated that they wanted to hire low expertise workers. The sample is restricted to the first job post by the employer to have at least one MBG-eligible applicant. The estimation method is OLS.

Table 3.4: Effect of MBG on Job Fill Rate by Employer Platform Experience

	Dependent variable:
	Fill Rate
TRT	0.013*
	(0.007)
Employer FJP	−0.133***
	(0.007)
TRT × FJP	−0.012
	(0.010)
Constant	0.431***
	(0.005)
Observations	36,446

Notes: The omitted group contains employers who have prior experience on the platform. The sample is restricted to the first job post by the employer to have at least one MBG-eligible applicant. The estimation method is OLS.

Table 3.5: Effect of MBG on Ex Ante Quality of Hired Employees

	Dependent variable:			
	I(MBG Eligible)		log(BestMatch)	
	(1)	(2)	(3)	(4)
TRT	0.049***	0.043***	0.013***	0.012***
	(0.007)	(0.013)	(0.002)	(0.004)
Intermediate Expertise		0.035***		0.011***
		(0.012)		(0.004)
High Expertise		0.078***		0.029***
		(0.015)		(0.005)
TRT × Intermediate Expertise		0.004		0.001
		(0.017)		(0.005)
TRT × High Expertise		0.018		0.001
		(0.021)		(0.007)
Constant	0.442***	0.412***	0.367***	0.357***
	(0.005)	(0.009)	(0.002)	(0.003)
Observations	19,128	19,000	19,128	19,000

Notes: The sample is restricted to hired applicants that applied to the first job posts by an employer to have at least one MBG-eligible applicant. The estimation method is OLS. The DV in model (1) and (2) is an indicator equal to one if the hired applicant is MBG eligible. The DV in model (3) and (4) is the log of the BestMatch score of the hired applicant. The excluded category in models (2) and (4) is employers looking to hire low expertise applicants.

Table 3.6: Effect of MBG on Characteristics of Hired Employees

	Dependent variable:			
	log(Prior Hours)	log(No. Prior Jobs)	log(Prior Earnings)	log(Wage Rate)
	(1)	(2)	(3)	(4)
TRT	0.067*	0.036*	0.069	−0.007
	(0.040)	(0.021)	(0.046)	(0.011)
Constant	4.627***	2.384***	6.652***	2.217***
	(0.028)	(0.014)	(0.033)	(0.008)
Observations	19,135	19,135	19,135	19,135

Notes: The sample is restricted to the hired applicants that applied to first job posts by an employer to have at least one MBG-eligible applicant. The estimation method is OLS.

Table 3.7: Effect of MBG on Characteristics of Hired MBG-Eligible Employees

	Dependent variable:				
	log(BestMatch)	log(Prior Hours)	log(No. Prior Jobs)	log(Prior Earnings)	log(Wage Rate)
	(1)	(2)	(3)	(4)	(5)
TRT	0.002	0.030	0.014	0.012	−0.022
	(0.003)	(0.052)	(0.028)	(0.056)	(0.016)
Constant	0.495***	5.262***	2.852***	7.634***	2.380***
	(0.002)	(0.038)	(0.020)	(0.040)	(0.012)
Observations	8,922	8,922	8,922	8,922	8,922

Notes: The sample is restricted to hired applicants with BestMatch scores above .5 that applied to the first job post by an employer to have at least one MBG-eligible applicant. The estimation method is OLS. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table 3.8: Effect of MBG on Job Outcomes

	Dependent variable:		
	log(Public Feedback)	log(Private Feedback)	I(Rehired)
	(1)	(2)	(3)
TRT	0.002	0.016*	0.007
	(0.004)	(0.009)	(0.005)
Constant	1.728***	2.204***	0.140***
	(0.003)	(0.006)	(0.003)
Observations	13,160	14,446	22,385

Notes: The sample is restricted to filled job posts that have at least one MBG-eligible applicant. The estimation method is OLS.

Table 3.9: Effect of MBG on Job Outcomes MBG-eligible Applicants Only

	Dependent variable:		
	log(Public Feedback)	log(Private Feedback)	log(Rehire)
	(1)	(2)	(3)
TRT	−0.0002	0.006	0.007
	(0.005)	(0.012)	(0.007)
Constant	1.736***	2.219***	0.144***
	(0.004)	(0.009)	(0.005)
Observations	6,123	6,616	9,969

Notes: The sample is restricted to filled job posts by an employer that hired a MBG-eligible applicant. The estimation method is OLS.

Table 3.10: Effect of MBG on Job Outcomes Rate by Desired Expertise Level of Employee

	Dependent variable:	
	log(Private Feedback)	I(Rehired)
	(1)	(2)
TRT	0.039**	−0.006
	(0.016)	(0.008)
Intermediate Expertise	0.050***	0.003
	(0.014)	(0.007)
High Expertise	0.046***	−0.015
	(0.018)	(0.010)
TRT × Intermediate Expertise	−0.029	0.019*
	(0.020)	(0.011)
TRT × High Expertise	−0.049*	0.026*
	(0.025)	(0.014)
Constant	2.171***	0.141***
	(0.011)	(0.006)
Observations	14,349	22,678

Notes: The omitted group in both model (1) and model (2) contains employers who were in the experimental control group and indicated that they wanted to hire low expertise workers. The sample is restricted to the first job post that have at least one MBG-eligible applicant. The estimation method is OLS.

Table 3.11: Effect of Guarantee vs. Recommendation on Job Fill Rate

	Dependent variable:
	Fill Rate
MBG Treatment Group (TRT)	-0.005
	(0.009)
Constant	0.395^{***}
	(0.006)
Observations	12,947

Notes: The dependent variable is whether the employer filled the opening. The sample is restricted to the first job post that has at least one MBG-eligible applicant. The estimation method is OLS.

Table 3.12: Effect of MBG vs. Recommendation on Job Fill Rate by Desired Expertise Level of Employee

	Dependent variable:
	Fill Rate
TRT	-0.006
	(0.016)
Intermediate Expertise	-0.021
	(0.014)
High Expertise	-0.044^{***}
	(0.017)
TRT × Intermediate Expertise	0.017
	(0.020)
TRT × High Expertise	-0.030
	(0.024)
Constant	0.415^{***}
	(0.011)
Observations	12,929

Notes: The omitted group contains employers who were in the experimental control group and indicated that they wanted to hire low expertise workers. The sample is restricted to the first job post that has at least one MBG-eligible applicant. The estimation method is OLS.

Table 3.13: Effect of MBG vs. Recommendation on Ex Ante Quality of Hired Employees

	Dependent variable:	
	I(MBG Eligible)	log(BestMatch)
	(1)	(2)
TRT	0.002	0.001
	(0.011)	(0.004)
Constant	0.631***	0.420***
	(0.008)	(0.003)
Observations	7,188	7,188

Notes: The sample is restricted to hired applicants that applied to the first job posts that have at least one MBG-eligible applicant. The estimation method is OLS. The DV in model (1) is an indicator equal to one if the hired applicant is MBG-eligible. The DV in model (2) is the log of the BestMatch score of the hired applicant.

Table 3.14: Effect of MBG vs. Recommendation on Job Outcomes

	Dependent variable:		
	log(Public Feedback)	log(Private Feedback)	log(Rehire)
	(1)	(2)	(3)
TRT	−0.001	0.0002	−0.008
	(0.006)	(0.016)	(0.008)
Constant	1.734***	2.183***	0.140***
	(0.004)	(0.011)	(0.006)
Observations	4,136	4,492	6,656

Notes: The sample is restricted to filled job posts that have at least one MBG-eligible applicant. The estimation method is OLS.

Chapter 4

Weather and Willingness to Pay

4.1 Introduction

Online auction sites such as eBay sell a multitude of goods and provide fertile ground for economic insight into the behavior of buyers and sellers. Economic understanding of strategic bidding behavior in online auctions has been driven by auction theory, which provides sharp predictions on equilibrium bidding behavior. Vickrey (1961) observed that in a second-price sealed-bid auction with private values, it is a weakly dominant strategy for a bidder to bid their reservation value. However, empirical studies of bidding behavior on eBay have been hard pressed to show that bidders on eBay bid according to a rational bidding strategy. Specifically, over multiple auctions for the same good, bidders on eBay may vary their bids up and down by drastic amounts. These changes in bid amount are not sufficiently explained by changes in supply and demand or auction-specific variation. This behavior is seemingly irrational if one's dominant strategy is to always bid one's valuation.

This paper seeks to uncover the behavioral underpinnings that could cause bidders to vary their valuation of an item from day to day. Behavioral anomalies, that is, systematic deviations from the von Neumann-Morgenstern axiom of rational behavior, could explain why a single bidder varies his bids on a homogeneous item dramatically from day to day and auction to auction. Specifically, if one believes in the constructed preferences hypothesis (Payne et al., 1999) under which people are assumed to construct preferences on the spot when asked to reach a particular judgment or decision, then changes in mood would cause a revaluation of willingness to pay (WTP) and more over, could explain variation in bidding. Freudian theory proposes that human motivation is a result of unconscious needs and drives. As such, consumer behavior could be at some level driven by unconscious needs to alter bad mood or promote good mood (Winslow, 1986; Billig, 1999).

The weather is commonly used as a proxy for mood as there is substantial psychological evidence that weather affects mood (Hirshleifer and Shumway, 2003). Seasonal affective disorder, more commonly known as the winter blues, is a well-documented disorder. To some extent nearly everyone experiences some differences in mood between warm, sunny, clear days and cold, cloudy, rainy days. A substantial amount of economic literature has used sunshine as a proxy for mood (Kliger, Raviv, et al., 2010 and others). Additionally, the neuroscience literature has developed a link between weather and mood. Neuroscientists believe that sunlight and weather affect mood though the body's level of dopamine and melatonin: production of melatonin by the pineal gland is inhibited by light to the retina and permitted by darkness. The body usually retains a balance between dopamine and melatonin. During the day dopamine levels are elevated and melatonin levels are down, and during the night dopamine levels decrease as melatonin levels increase. However, during cloudy weather, the lack of sunlight causes melatonin production to rise, which decreases dopamine levels and causes a more negative mood.[1]

While the weather is commonly used as a proxy for mood, it can also directly affect economic behavior. Thus, teasing out the effects of weather on bidding behavior through the channel of mood has been difficult. This paper uses a series of empirical tests to separate the various mechanisms through which the weather can affect a bidder's WTP.

This paper uses a series of empirical tests to study the effect of weather on bidding behavior on eBay. These tests will allow me to isolate the various mechanisms through which weather can affect bids submitted on eBay. In eBay auctions, bids are submitted by a proxy bidding system. When bidding on an item, users enter the maximum amount they would be willing to pay for the item. The seller and other bidders don't know the maximum bid. The eBay proxy bidding system will automatically place bids on the bidder's behalf using the automatic bid increment amount, which is based on the current highest bid. The system only bids as much as necessary to make sure that the bidder remains the high bidder, or to meet the reserve price, up to the entered maximum amount. Thus, by summing bidder's highest proxy bid in an auction with the posted shipping price, I am able to construct a variable which represents an individual's WTP in a particular auction for a particular good.[2]

This paper has three main findings. I first demonstrate that on rainy days bidders on eBay increase their bids by about $1.00 on average on consumer products. Secondly, I show that rain does not seem to have any statistical effect on bidders WTP for cash-like items such as gift cards. Finally, I show that rain has no effect on a bidder's bid submission behavior including the number of bids submitted, and how close the bid is submitted to the end of the auction. If the weather primarily influenced WTP through effecting cognitive behavior that influences bidding strategies, or by increasing travel costs by making roads dangerous, then theory would predict that the results should be similar for both consumer goods and cash-like substitutes. Additionally, if the weather was decreasing the opportunity cost of bidding by increasing probability that the bidder is indoors and in front of his or her computer, then theory would predict that weather would affect submission behavior including the number and timing of bidders bids. Together these results indicate that weather affects a bidder's

[1]See Nelson (2005) for a much more complete description of the body's circadian rhythm.

[2]Brown, Hossain, et al. (2007) show that bidders on eBay often fail to properly account for shipping. All of the results presented are robust to specifications that do not include shipping.

WTP primarily by influencing mood, which causes a revaluation of WTP for an item. The remainder of the chapter is organized as follows. Section 2 reviews the previous literature on mood and emotional effects in various decision-making settings. Section 3 describes the data. Section 4 discusses the mechanisms through which weather can influence a bidder's WTP on eBay. Section 5 describes the empirical specifications and results. And section 6 discusses and concludes the paper.

4.2 Literature Review

Standard economic theory assumes that individuals make decisions based on the desirability or utility of their available options. As a result, individual preferences can be inferred by observing choice. Behavioral economists have given substantial attention to the effect of specific emotions associated with decisions. Early economic references to emotion in choice include John Maynard Keynes' famous reference to "animal spirits" and Adam Smith's discussion of the "influence of passions." Specifically, behavioral economics has focused on the effects of emotion, and has investigated how specific emotions such as anger, reciprocity, and guilt affect decisions. These studies have advanced our understanding of preferences and how decisions are made.

The psychological literature makes a stark distinction between emotion and mood. According to psychologists, emotions tend to last for brief periods of time, but moods are expected to last longer (Watson and Vaidya, 2003). For example, an emotion of joy created by eating a chocolate (Macht and Dettmer, 2006) is prone to last for a moment, while positive mood is a relatively long-lasting emotional state that can affect decisions long after having been induced (Knowles et al., 1993).

The psychology literature, however, is mixed on the direction of the impact of mood on valuation. Two models predict the direction of this impact. The first is the affect infusion model (Forgas, 1995), which suggests that positive mood results in higher valuation. It hypothesizes that positive mood causes subjects to focus their attention on the positive attributes of an item, and thus predicts a higher valuation under a positive mood and a lower valuation under a negative mood. Alternately, the mood maintenance hypothesis (Isen and Geva, 1987) predicts that people in a negative mood bid higher to alter their negative mood by winning, while people in a positive mood refrain from risk taking to prevent a change in their situation.

Behavioral economics has long failed to differentiate between mood and emotion, and thus economic studies of the effects of people's global affective states or moods on individual behavior are scarce. However, economic forays into emotional response are substantial. Frank (1988) argues that emotions have important roles in decision making and personal interactions. When accounting for utility gain from emotions, seemingly irrational decisions can become rational. Using game theory, Rabin (1993) formalized the emotional response that people like to help people who are helping them, and people like to hurt those who are hurting them. Loewenstein and Lerner (2003) moved the behavioral economic interest in emotions from anticipated emotions, such as regret and disappointment, to immediate emotions. They studied how current emotions, such as anger, fear or pain – which can be modeled as a state dependent variables – affect economic behavior.

More recent laboratory experiments have further investigated the role of emotions in seemingly irrational decisions. In Capra (2004), the effect of induced mood was studied in strategic interaction through three classical games: the dictator, ultimatum and trust games. She found that choices under the induced bad mood are closer to the benchmark rational choices when compared to choices under a good mood. Individuals in a good mood tend to be more altruistic. Ben-Shakhar et al. (2007) used a one-shot, two-person, power-to-take game in which player 1 can claim any part of player 2's resources, and player 2 can react by destroying some or all of these resources and thus preventing player 1 from obtaining the claimed amount. They found that a physiological measure is related with both self-reported anger and the destructive decision. Sanfey et al. (2003) provides evidence which supporting Benshakhar's findings by scanning neural activities in subjects as they respond to fair and unfair proposals in the ultimatum game.

Moving beyond game theoretic applications, Kahneman et al. (1991) showed that individuals place a higher price on the object they own (selling price) than the same object they do not own (buying price). The induced disgust emotion was shown to reduce both selling and buying price, while sadness was shown to make the selling price higher than the buying price. A substantial amount of further research on the effect of mood can be found in the marketing literature. Spies et al. (1997) observed that consumers in a good mood spend more on impulsive purchases and are more satisfied with their shopping experience. However, this paper was non-experimental in that it correlated the non-random pleasantness of a store to mood as measured by a survey taken during the shopping experience to purchasing behavior.

While I have surveyed a substantial range of literature documenting the connection between mood and economic behavior, the nature of inducing mood means that many of these studies are conducted in laboratory settings, and there is still substantial room for the study of the effects of mood on real-world transactions.

Thus, I turn to studies which take advantage of the exogenous nature of weather and its well-established relationship to mood. Psychological studies using weather as a proxy for mood include Rind and Bordia (1996), who showed that sunnier days and beliefs about sunnier days are correlated with increased tipping at restaurants. Weather-related mood effects have also been explored in several economics and finance studies. Saunders (1993) shows that when it is cloudy in New York City the New York Stock Exchange index returns tend to be negative. This finding is further substantiated by Hirshleifer and Shumway (2003) who look at weather effects on stock returns from 1982 to 1997 in 26 different cities where stock exchanges are located. They find a negative relationship between cloud cover and stock returns in almost all the cities. Kliger and Levy (2003) show that investors' risk preferences are affected by seasonal mood, and Kliger, Gurevich, et al. (2009) document seasonal effects on initial public offerings' short-run and long-run performance.

The only study I am aware of that examines the link between weather and bidding behavior is a working paper by Kliger, Raviv, et al. (2010). Kliger, Raviv, et al. (2010)

studies art prices at auctions conducted from 1786 to 1909 in England and shows that length of the day had significant positive effect on the auction selling price. Presumably, the length of the day is directly linked to sunlight, which influences mood. In contrast to Kliger, Raviv, et al. (2010), this study focuses on items where the market value of the item is well known. The subjective nature of art makes it difficult to control for item-level variation in preferences; additionally, longer days could possibly be correlated with longer auctions and increased bidding especially in the era prior to electricity. Thus, their study is unable to fully explore the mechanisms though which bidding affects WTP.

4.3 Data

Bid-level eBay data was purchased from Advanced Ecommerce Research Systems (AERS). For each auction several characteristics are observable: the auction title, subtitle, and category listing. Additionally, I observe the auction starting time and starting price as well as ending time and ending price. A unique identifier for each seller is also observable. I also observe the seller's total feedback score, if the auction included a Buy It Now option and the Buy It Now price. Finally, I observe several auction-level indicator variables that denote if the auction was specially featured in any way. For example, the seller could pay extra for the auction to have a bolded title or high-resolution pictures. For each auction, I observe every bid amount submitted to eBay's proxy bidding system by each bidder. Bidders are identified by a unique bidder identifier. The bidder's total feedback score and the exact time of the bid is recorded. I also observe the current price of the auction when each bid is submitted.

I restrict my analysis to a group of carefully selected items covering a wide variety of types of items available on eBay. These include: TI-83 Plus calculators, second-generation iPod Nanos, Giorgio Armani Acqua Di Gio for men 3.4 oz sets of perfume, Target gift cards, and Windows Vista Home Premium and Ultimate operating systems. These goods were selected based on their popularity and homogeneous nature. The item space includes software, hardware, luxury goods and non-luxury goods. It also includes goods that depreciate in value over time as well as well as cash-substitutes that experience no depreciation. The data includes both used and new items. The AERS data covers all transactions for the above items for auctions ending from January 1, 2008, to December 31, 2010. Due to the multiple-year nature of the data, significant variation in supply and demand as well as they types of buyers and sellers may occur even within a product group. This variation will need to be addressed econometrically.

Historical weather data was obtained from Weather Underground (www.wundergound.com) using a web-scraper program. Data obtained from weather underground is preferable to that supplied by the National Climatic Data Center (NCDC) as it supplements the 22,000 weather stations supported by NCDC with 18,000 additional personal weather stations. This allows for weather to be observed at the zip code level as opposed to the weather station level. Observations include temperature, dew point, pressure, visibility, wind speed, precipitation cloud cover, and weather events such as hail, thunderstorms, or tornados. The historical weather data was matched to the bid-level eBay data by date and zip code. Zip code data is only observable for winners of auctions, but by matching based on unique bidder ID I am able to match 42% of the total bidders to zip codes.[3]

This matching method relies on the assumption that the bidder location remains fixed between the moments when he bids and does not win, and when he wins an auction. Another issue would be if purchases were shipped to a location different from where it is purchased. On average, a non-winning bid is matched, with zip code information, to a winning bid that is within 79.67 days. The median match is only 2.69 days apart, and the 75th percentile match is only 59.89 days. Since 75% of the matches are less than two months apart, I assume that the bidder did not move between when I observe his zip code and when he bid. When there are multiple zip-codes per bidder, the match that is closest in days is used. This matching methodology is necessary for the analysis, but it is important to realize that it is a potential source of error. Only 11.9% of Americans moved during 2007 and 2008 and 12.5% of Americans moved during 2009. Additionally, most of these moves were within the same county.[4] These statistics combined with the short number of days between most of the matches makes the assumption of fixed bidder location even more justifiable. A more conservative approach would be to run the regressions at an auction level using only the winning bid price. However, since bidders on eBay usually bid across many auctions until they win, this would reduce my sample to about one bid per individual. This would prohibit me from taking advantage of individual fixed effects regressions. Additionally, my question of interest is at the bid level not the auction level. I am not interested in if weather affects the price of a good sold on eBay, but if it predicts variation in a bidder's willingness to pay.

To help get a sense of the amount of variation in bidding amounts, I present some basic summary statistics below in table 4.1. These statistics will prove insightful when investigating if the calculated effect of weather on bid amounts is economically meaningful. The statistics are calculated at the individual-mean bid level by rainy and non-rainy days.

In the final data set, attention is restricted to successfully completed auctions with known shipping costs and normal durations. Furthermore, in an attempt to limit the product space auctions with uncommon condition values are removed. More precisely, item quality can be broken into three categories: new, used and unknown. The analysis focuses on the most common condition per product: new for perfume, gift cards, and Windows; used for calculators and iPods. Minority conditions and broken items are never considered. Additionally, the top and bottom 1% of bids by item group are trimmed. The top 1% of bids are removed to get rid of shill bids, which is the practice by sellers of bidding a high price on their own item when the auction might close at a low price, so they can repost the auction. The bottom 1% of bids are removed to eliminate false WTP generated by the practice of bidding a very low

[3]See Appendix C.1 for a comparison of bidding behavior between eventual winners and those bidders who never win an auction. Eventual winners on average bid much lower than those who never win, and also bid much further from the end of the auction. This indicates that eventual winners are much more selective about what auctions they participate in. Additionally, eventual winners have statistically lower variances in their bidding behavior.

[4]http://online.wsj.com/article/SB10001424052748704879704575236533316039428.html

amount on an auction to more easily follow the auction moving forward. eBay later added a Follow Auction which greatly reduced this behavior.

4.4 Theoretical Mechanisms

There are several mechanisms through which the weather can affect bidding behavior; these are detailed in figure 4.1. The first and most common mechanism is by directly changing the actual value of the item. Weather can directly influence the value of a consumer good. For example, when it rains, the value of an umbrella goes up. To remove this possible mechanism, I carefully chose a set of consumer products whose values do not directly increase or decrease with weather. The second channel through which the weather can affect bidding behavior is by altering the bidder's value of time or opportunity cost of bidding. On rainy days, for example, a bidder may be confined to the house, and thus more closely monitor active auctions on eBay. While closely monitoring auctions, the bidder may be more likely to get swept up in "bidding fervor," leading to a higher measured WTP that is not at all affected by mood. The third channel through which the weather can affect bidding is by altering mood. This alteration of mood can change WTP in two ways. Firstly, weather might affect mood directly altering the perceived value of the item. Secondly, weather might affect mood influencing cognitive processe, changing bidding strategies and indirectly changing the amount bid. For example, mood could cause a bidder to become more or less competitive. The value for the item is unchanged, but the value elicitation mechanism itself may be influenced by mood. It is possible that mood influences the way subjects decide how much to bid. Indeed, mood has been shown to alter cognition, which, in an auction environment, may result in an increase/reduction in deviations from the best strategy.

4.5 Empirical Analysis and Results

To explore the relationship between the weather and individual's WTP for a good in a particular auction, I begin with a regression of the following form:

$$\textsc{Bid} + \textsc{Shipping}_{iptz} = \beta_0 + \beta_1 \textsc{Rain}_{tz} + \textsc{Individual Cov}_{it} + \textsc{Auction Cov}_{it} + \delta_{iptz} + \epsilon_{iptz} \tag{4.1}$$

where the dependent variable is the bid amount plus shipping cost placed by individual i on product p at year-month t in zip code z.

The explanatory variable of interest is an indicator of rain at time t in zip code z.[5] Despite not having a clear theoretical prediction about optimal bidding strategies, I conduct a

[5]In table 4.2 in C.2 I use the bid without shipping as the dependent variable, and the results are nearly identical.

reduced form analysis of how mood (affected by weather) affects bidding. I control for within-bidder, time dependent experience as well as auction level heterogeneity. I use individual, item, and time fixed effects. This means that the effect is measured from variation in an individual's bid across different auctions for the same item within the same month. I do not analyze the effect of temperature, as temperature is not a monotonic measure of mood. For example, mood may be higher on a 70-degree day than on a 50-degree day, but mood is probably not higher on a 95-degree day than on a 70-degree day. To test the robustness of the results, I ran various specifications using weather as a covariate. This did not change any of the results.

The regressor of interest (weather) varies only at the zip code-day level, but the dependent variable (bid + shipping) varies at the individual-zip code-day level. Thus, I expect that there is substantial serial correlation in the residual within zip code. To solve this I replace the standard OLS variance estimator with a robust cluster-variance estimator with clustering by zip code.[6] More conservatively, all regressions were also run with clusters at the county level, with no change in what variables are significant.

The analysis necessitates the use of a multi-month, multi-year dataset. This introduces complications as significant variation may occur across products, conditions, and categories during the sample. Any specification that fails to include month-year fixed effects is inherently flawed as there are unobserved time trends in bidding that are correlated with season and by extension, weather. For example, the period between Thanksgiving and Christmas is regarded as the biggest shopping time of the year. In fact, aiming to boost its standing as a holiday shopping destination, eBay launched a holiday season marketing campaign in 2009, which it renewed in 2010.[7] This time of year also tends to be characterized by more cloudy, snowy, and rainy weather for much of the country. This unobserved correlation between weather, time, and unobserved demand and supply shocks must therefore be controlled for using month fixed effects. Additionally, I use month-year fixed effects as there are also unobserved shocks which vary by month and year. For example, Windows 7, the successor to Windows Vista, was released on October 22, 2009.

I am still concerned about unobserved individual treatment-level heterogeneity. I am concerned that people select themselves into the environment in which they live. Hypothetically, it's possible that people who love cloudy weather and rain and shun material items – reducing their valuation of items on eBay – choose to live in the Pacific Northwest, while those whose moods are most profoundly affected by sunlight and who place higher value on material goods select into living in southern California. This would cause an unobserved correlation between weather and bidding across geographic region. To correct for this possibility, I add individual fixed effects to the model. While the above story could be corrected for by using zip code-level fixed effects, I decide to follow the even more conservative methodology and use individual fixed effects which adjust for the additional possibility that there is some

[6] I am actually concerned by both correlation within zip code as well as over time, I aggregate only at the group level as recommended by Arellano (1987) and Angrist and Pischke (2009).

[7] http://online.wsj.com/article/SB10001424052748704746304574505543900212118.html

additional unobserved variable correlated with individual characteristics and willingness to pay. The results using zip code-level fixed effects are nearly identical to those with individual fixed effects. I am now only using variation within an individual, month, year, and item to drive the results.

To empirically isolate the different channels though which the weather, namely rain, affects willingness to pay, I designed a series of empirical tests that take advantage of the different items in the dataset. I first note that Target gift cards are fundamentally different from the other items in the dataset in that their value is common to all bidders. In fact, having an associated dollar amount effectively induces a valuation. Since a \$100 Target gift card is worth exactly \$100 (minus some amount due to lack of spending flexibility), it is unlikely that the perceived value of the gift card will be altered by mood. If weather affects the willingness to pay for a Target gift card, it will likely do so by altering the way subjects decide what to bid. The results reported in table 4.2 omit target gift cards from the analysis.

I begin by showing that rain has an effect on bids submitted on eBay on non cash-like items. This effect could be due to mood or opportunity costs – which I will explore further. Model (1) of table 4.2 includes only covariates that control for the seller's specified choice of auction, the seller's history and reputation, and the bidder's experience and reputation. Rain on the day of the bid corresponds to a \$2.14 increase in the average bid. This implies that rain has a significantly positive effect on a bidder's valuation of a good on eBay. In model (2) I add in month-year fixed effects. Including time fixed effects does not have any meaningful effects on the outcome of the regression. In model (3) I run the preferred specification which uses individual bidder fixed effects. Here the coefficient is slightly smaller, but still positive and significant.

I next run the same regression, but limit the sample to only Target gift cards. This will help to establish the mechanism through which weather and by proxy mood influence willingness to pay. If weather has a similar effect on bidding for Target gift cards as it did for non-Target gift cards, it is more likely that weather does not have any effect on the perceived value of the item, but instead is affecting cognition which alters bidding strategy. Figure 4.3 shows that there is no significant difference between willingness to pay on days when it rains and days when it does not rain. Thus, I conclude that if weather is affecting willingness to pay through mood then its most likely doing so by directly altering the perceived value of the item to the bidder.

Since I demonstrate that worse weather seems to cause an increase in bidding on eBay, it is possible for one to think that I am not actually picking up any effect of mood on bidding. It's possible that on days with worse weather (e.g. rain) a bidder would rather be indoors instead of outside. Being stuck indoors allows the bidder to follow an auction more closely which causes him to get caught up in bidding fervor and increase his bid. To check if this is a viable alternative hypothesis, I devise two robustness checks. First, I calculate the time from when the bid took place to the end of the auction, as I hypothesize that if a bidder was following an auction more closely, he would wait longer to submit his bid. If this alternate story were true, I also expect bidders to submit more bids per auction on days when they are stuck indoors.

Model (1) of table 4.4 regresses the time to the end of the auction in seconds on a rainy day indicator and auction-level covariates. I include individual, and item-month-year fixed effects just as in the main specification. Rain does not seem to have any significant effect on how close to the end of the auction a bid was submitted. Model (2) regresses the number of bids submitted per day on a rainy day indicator. I see that once again the rainy day indicator is not significantly different from zero. These robustness checks seem to indicate that the effect of weather on bidding is probably not only due to bidders paying more attention to an auction.

4.6 Discussion and Conclusion

I find that a rainy day is associated with a bidder on eBay submitting bids on consumer items that are about 1% higher on average than those submitted on non-rainy days. This relationship only holds for general consumer items. There does not seem to be a statistically significant relationship between rain on the day a bid is submitted and the bid submitted for Target gift cards, which can be considered a type of cash substitute. Finally, rain does not seem to have any effect on either the number of bids submitted on that day, or how close to the end of the auction the last bid on a given auction was submitted.

These empirical tests are meant to help tease out the mechanism through which weather and possibly mood affect bidders' willingness to pay for items on eBay. The products I analyze can be separated into two general categories: consumer goods, which can generally be considered private value goods, and cash-like substitutes, which can be considered public value goods. This differentiation is important as it allows me to unpack the mechanism through which mood effects willingness to pay. The literature in psychology has highlighed that mood can directly affect an individual's perceived value of an item. The literature has also detailed that mood influences cognitive processes, changing bidding strategies and indirectly changing the amount bid. If mood were to affect cognitive processes more generally, it seems likely that the effect would also be seen on cash-like substitutes. Finally, to help alleviate concerns that the effects of the weather on bidding are not due to mood at all, but instead to changes in opportunity costs of bidding, I show that weather has no effect of the timing and number of bids submitted in an auction.

Taken all together, the evidence presented seems to suggest that the weather does have a significant yet small effect on bidding behavior on eBay. It further suggests that the mechanism through which weather affects bidding behavior on eBay is by altering mood, which alters a bidder's perceived valuation of an item.

While weather has long been used as a proxy for mood, and has been demonstrated to influence economic behavior, the previous literature has failed to fully separate and identify the various mechanisms through which weather could be affecting behavior. The evidence presented in the paper, while only suggestive of the true mechanism through which weather affects valuation, is to my knowlage the first attempt to empirically separate the various mechanisms through which weather alters bidders' willingness to pay.

It is important to note the limitations of the analysis in this paper. One very important limitation is that due to my method of matching, my sample is limited to bidders who have won at least one of my analyzed items on eBay from January 2008 through December 2010. This means that I have a selected sample of winners compared to the average bidder on eBay.

My results are potentially relevant for understanding how preferences are constructed and how individuals assign economic value to goods. Specifically, these results lend support for the constructed preferences hypothesis (Payne et al., 1999) under which people are assumed to construct preferences on the spot when asked to reach a particular judgment or decision. They further demonstrate that the effects of mood can help elucidate consistent revaluation of consumer products by individuals.

Figure 4.1: Mechanisms Through Which Rain Effects Submitted Bids

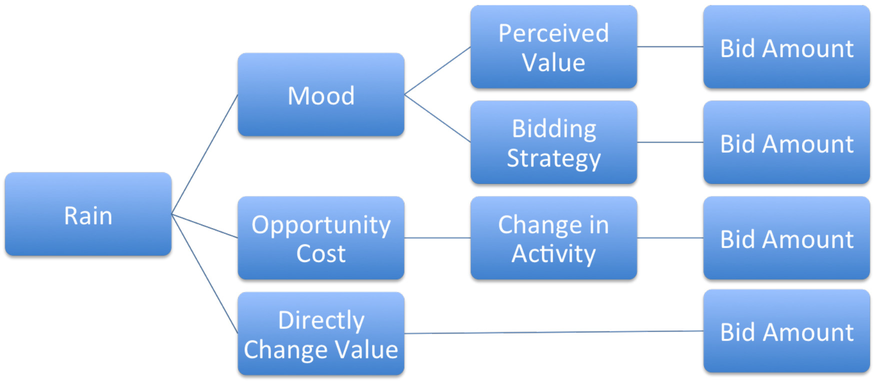

Table 4.1: Mean Bidding Behavior By Product and Rain

	No Rain Mean: $\bar{X}_{CTL}$	Rain Mean: $\bar{X}_{TRT}$	Difference In Means	p-value	
Target Gift Card					
Number of Observations	22903	7785			
Avg Bid	123.17	125.57	2.41 (1.76)	0.17	
Avg Bid + Shipping	123.73	126.19	2.45 (1.76)	0.16	
Bid/Amazon Price	0.88	0.88	-0.00 (0.00)	0.20	
Min to Auction End	2594.93	2515.61	-79.32 (37.36)	0.03	**
Windows Software					
Number of Observations	19745	8233			
Avg Bid	103.88	109.38	5.50 (0.71)	<0.001	***
Avg Bid + Shipping	109.30	114.86	5.56 (0.71)	<0.001	***
Bid/Amazon Price	0.34	0.35	0.01 (0.00)	<0.001	***
Min to Auction End	1537.31	1654.07	116.76 (33.46)	<0.001	***
Apple iPod					
Number of Observations	7282	2320			
Avg Bid	54.56	55.57	1.02 (0.61)	0.10	*
Avg Bid + Shipping	59.70	60.83	1.13 (0.62)	0.07	*
Bid/Amazon Price	0.81	0.82	0.01 (0.01)	0.19	
Min to Auction End	510.51	549.70	39.19 (33.01)	0.24	
TI 83 Calculator					
Number of Observations	2261	794			
Avg Bid	36.53	40.65	4.13 (0.56)	<0.001	***
Avg Bid + Shipping	42.30	46.02	3.72 (0.54)	<0.001	***
Bid/Amazon Price	0.89	0.97	0.08 (0.01)	<0.001	***
Min to Auction End	227.26	181.23	-46.04 (30.88)	0.14	
Armani Perfume					
Number of Observations	4609	1920			
Avg Bid	31.44	31.40	-0.04 (0.23)	0.87	
Avg Bid + Shipping	37.59	37.65	0.06 (0.23)	0.80	
Bid/Amazon Price	0.79	0.79	0.00 (0.00)	0.71	
Min to Auction End	839.01	922.02	83.01 (46.36)	0.07	*

Notes: This table reports means and standard errors for bidding behavior for each item in the dataset on rainy and non-rainy days. Observations are at the individual-auction level. Only the highest bid for an individual in each auction is kept. Reported p-values are for two-sided t-tests of the null hypothesis of no difference in means across groups. In the bottom panel, standard errors are clustered at the employer level. Significance indicators: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

Table 4.2: The Effect of Rain on Willingness to Pay, Consumer Goods

	Dependent variable:		
	Bid + Shipping		
	(1)	(2)	(3)
Rain Indicator	2.145^{***}	1.824^{**}	1.167^{*}
	(0.392)	(0.898)	(0.665)
Opening Level Covariates	Yes	Yes	Yes
Month-Year FE	No	Yes	Yes
Item FE	No	Yes	Yes
Bidder FE	No	No	Yes
Observations	47,110	47,110	47,110

Notes: Covariates include: the minimum bid, the bidder's previous bid count, weekend indicator, number of simultaneous auctions, seller feedback, indicator for reserve, indicator for bold auction title, indicator for picture in auction. All models are OLS with robust standard errors clustered at the zip code level. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table 4.3: The Effect of Rain on Willingness to Pay, Gift Cards

	Dependent variable:		
	Bid		
	(1)	(2)	(3)
Rain Indicator	−0.484	0.023	0.176
	(1.448)	(2.752)	(1.534)
Opening Level Covariates	Yes	Yes	Yes
Month-Year FE	No	Yes	Yes
Bidder FE	No	No	Yes
Observations	30,579	30,579	30,579

Notes: Covariates include: the minimum bid, the bidder's previous bid count, weekend indicator, number of simultaneous auctions, seller feedback, indicator for reserve, indicator for bold auction title, indicator for picture in auction. All models are OLS with robust standard errors clustered at the zip code level. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table 4.4: The Effect of Rain on Timing/Quantity of Bids, Non Gift Cards

	Dependent variable:	
	Min To End	Bid/Day
	(1)	(2)
Rain Indicator	35.993	0.029
	(34.487)	(0.030)
Opening Level Covariates	Yes	Yes
Month-Year FE	Yes	Yes
Item FE	Yes	Yes
Bidder FE	Yes	Yes
Observations	47,110	34,966

Notes: Covariates include: the minimum bid, the bidder's previous bid count, weekend indicator, number of simultaneous auctions, seller feedback, indicator for reserve, indicator for bold auction title, indicator for picture in auction. All models are OLS with robust standard errors clustered at the zip code level. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Bibliography

Abell, Peter Malcolm, Teppo Felin, and Nicolai J Foss (2007). "Building micro-foundations for the routines, capabilities, and performance links". *Capabilities, and Performance Links.*

Agrawal, Ajay K, Nicola Lacetera, and Elizabeth Lyons (2013). "Does Information Help or Hinder Job Applicants from Less Developed Countries in Online Markets?"

Agrawal, Ajay, John Horton, et al. (2013). "Digitization and the Contract Labor Market". *Economic Analysis of the Digital Economy*, p. 219.

Arrow, Kenneth J (1996). "The economics of information: An exposition". *Empirica* 23.2, pp. 119–128.

Autor, David H (2001a). "Why do temporary help firms provide free general skills training?" *Quarterly Journal of Economics*, pp. 1409–1448.

— (2001b). "Wiring the labor market". *Journal of Economic Perspectives*, pp. 25–40.

Autor, David H and David Scarborough (2008). "Does job testing harm minority workers? Evidence from retail establishments". *The Quarterly Journal of Economics*, pp. 219–277.

Bagwell, Kyle (1992). "Pricing to signal product line quality". *Journal of Economics & Management Strategy* 1.1, pp. 151–174.

Bagwell, Kyle and Michael H Riordan (1991). "High and declining prices signal product quality". *The American Economic Review*, pp. 224–239.

Balachander, Subramanian (2001). "Warranty signalling and reputation". *Management Science* 47.9, pp. 1282–1289.

Barach, Moshe A. (2015). "Search, Screening, and Information Provision: Personnel Decisions in an Online Labor Market".

Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat (2010). "Information gathering and marketing". *Journal of Economics & Management Strategy* 19.2, pp. 375–401.

Barney, Jay (1991). "Firm resources and sustained competitive advantage". *Journal of Management* 17.1, pp. 99–120.

Barney, Jay B (1986). "Strategic factor markets: Expectations, luck, and business strategy". *Management science* 32.10, pp. 1231–1241.

Barron, John M, Mark C Berger, and Dan A Black (1997). "Employer search, training, and vacancy duration". *Economic Inquiry* 35.1, pp. 167–192.

Barron, John M, John Bishop, and William C Dunkelberg (1985). "Employer search: The interviewing and hiring of new employees". *The Review of Economics and Statistics*, pp. 43–52.

Barron, John M, Dan A Black, and Mark A Loewenstein (1989). "Job matching and on-the-job training". *Journal of Labor Economics*, pp. 1–19.

Beales, Howard, Richard Craswell, and Steven C Salop (1981). "The efficient regulation of consumer information". *The Journal of Law & Economics* 24.3, pp. 491–539.

Ben-Shakhar, Gershon et al. (2007). "Reciprocity and emotions in bargaining using physiological and self-report measures". *Journal of economic psychology* 28.3, pp. 314–323.

Billig, Michael (1999). "Commodity fetishism and repression: reflections on Marx, Freud and the psychology of consumer capitalism". *Theory & Psychology* 9.3, pp. 313–329.

Blasco, Sylvie and Barbara Pertold-Gebicka (2013). "Employment policies, hiring practices and firm performance". *Labour Economics* 25, pp. 12–24.

Brown, Jennifer, Tanjim Hossain, and John Morgan (2007). "Shrouded attributes and information suppression: Evidence from the field". *The Quarterly Journal of Economics*, pp. 859–876.

Brown, Meta, Elizabeth Setren, and Giorgio Topa (2012). "Do informal referrals lead to better matches? Evidence from a firm's employee referral system". *Evidence from a Firm's Employee Referral System (August 1, 2012). FRB of New York Staff Report* 568.

Bryant, W Keith and Jennifer L Gerner (1978). "The price of a warranty: The case for refrigerators". *Journal of Consumer Affairs* 12.1, pp. 30–47.

Budish, Eric, Benjamin N. Roin, and Heidi Williams (2015). "Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials". *American Economic Review* 105.7, pp. 2044–85.

Burdett, Kenneth and Elizabeth J Cunningham (1998). "Toward a theory of vacancies". *Journal of Labor Economics* 16.3, pp. 445–478.

Burks, Stephen V et al. (2013). "The value of hiring through referrals".

Capra, C Monica (2004). "Mood-driven behavior in strategic interactions". *The American Economic Review* 94.2, pp. 367–372.

Coff, Russell and David Kryscynski (2011). "Drilling for micro-foundations of human capital–based competitive advantages". *Journal of Management*, p. 0149206310397772.

Daughety, Andrew F and Jennifer F Reinganum (1995). "Product Safety: Liability, R&D, and signaling". *The American Economic Review*, pp. 1187–1206.

Devine, Theresa J, Nicolas M Kiefer, et al. (1991). "Empirical Labor Economics: The Search Approach". *OUP Catalogue*.

Dube, Arindrajit, Eric Freeman, and Michael Reich (2010). "Employee replacement costs". *Institute for Research on Labor and Employment*.

Fallick, Bruce Chelimsky (1992). "Job security and job search in more than one labor market". *Economic Inquiry* 30.4, pp. 742–745.

Forgas, Joseph P (1995). "Mood and judgment: the affect infusion model (AIM)." *Psychological bulletin* 117.1, p. 39.

Frank, Robert H (1988). *Passions within reason: the strategic role of the emotions.* WW Norton & Co.

Garvin, David A (1983). "Quality on the line". *Harvard Business Review* 61.5, pp. 64–75.

Gerner, Jennifer L and W Keith Bryant (1981). "Appliance warranties as a market signal?" *Journal of Consumer Affairs* 15.1, pp. 75–86.

Grossman, Sanford J (1981). "The informational role of warranties and private disclosure about product quality". *The Journal of Law & Economics* 24.3, pp. 461–483.

Grossman, Sanford J and Oliver D Hart (1980). "Disclosure laws and takeover bids". *The Journal of Finance* 35.2, pp. 323–334.

Grossman, Sanford J and Joseph E Stiglitz (1980). "On the impossibility of informationally efficient markets". *The American Economic Review*, pp. 393–408.

Hall, Richard (1993). "A framework linking intangible resources and capabilities to sustainable competitive advantage". *Strategic Management Journal* 14.8, pp. 607–618.

Hall, Robert E and Alan B Krueger (2010). "Evidence on the determinants of the choice between wage posting and wage bargaining".

Heal, Geoffrey (1977). "Guarantees and risk-sharing". *The Review of Economic Studies* 44.3, pp. 549–560.

Hirshleifer, David and Tyler Shumway (2003). "Good day sunshine: Stock returns and the weather". *The Journal of Finance* 58.3, pp. 1009–1032.

Hoffman, Mitch, Lisa B Kahn, and Danielle Li (2015). "Discretion in Hiring".

Hoffman, Mitchell (2014). "How is Information Valued? Evidence from Framed Field Experiments".

Holzer, Harry J (1987). "Hiring procedures in the firm: their economic determinants and outcomes".

— (1996). *What employers want: Job prospects for less-educated workers*. Russell Sage Foundation.

Horton, John J (2010). "Online Labor Markets". *Internet and Network Economics*, p. 515.

— (2013). "The effects of subsidizing employer search".

Horton, John J and Joseph M Golden (2015). "Reputation Inflation: Evidence from an Online Labor Market".

Isen, Alice M and Nehemia Geva (1987). "The influence of positive affect on acceptable level of risk: The person with a large canoe has a large worry". *Organizational Behavior and Human Decision Processes* 39.2, pp. 145–154.

Jovanovic, Boyan (1982). "Truthful disclosure of information". *The Bell Journal of Economics*, pp. 36–44.

Kahneman, Daniel, Jack L Knetsch, and Richard H Thaler (1991). "Anomalies: The endowment effect, loss aversion, and status quo bias". *The Journal of Economic Perspectives* 5.1, pp. 193–206.

Katz, Michael L and Carl Shapiro (1994). "Systems competition and network effects". *The Journal of Economic Perspectives* 8.2, pp. 93–115.

Kihlstrom, Richard E and Michael H Riordan (1984). "Advertising as a Signal". *The Journal of Political Economy*, pp. 427–450.

Kliger, Doron, Gregory Gurevich, and Abraham Haim (2009). "When Chronobiology Met Economics-Seasonal Affective Impact on the Demand for IPOs".

Kliger, Doron and Ori Levy (2003). "Mood-induced variation in risk preferences". *Journal of Economic Behavior & Organization* 52.4, pp. 573–584.

Kliger, Doron, Yaron Raviv, et al. (2010). "Auction Prices and the Weather: New Evidence from Old Masters".

Knowles, Patricia A, Stephen J Grove, and W Jeffrey Burroughs (1993). "An experimental examination of mood effects on retrieval and evaluation of advertisement and brand information". *Journal of the Academy of Marketing Science* 21.2, pp. 135–142.

Kőszegi, Botond (2003). "Health anxiety and patient behavior". *Journal of Health Economics* 22.6, pp. 1073–1084.

Kraemer, Carlo, Markus Nöth, and Martin Weber (2006). "Information aggregation with costly information and random ordering: Experimental evidence". *Journal of Economic Behavior & Organization* 59.3, pp. 423–432.

Kübler, Dorothea and Georg Weizsäcker (2004). "Limited depth of reasoning and failure of cascade formation in the laboratory". *The Review of Economic Studies* 71.2, pp. 425–441.

Loewenstein, George and Jennifer S Lerner (2003). "The role of affect in decision making". *Handbook of Affective Science* 619.642, p. 3.

Lutz, Nancy A (1989). "Warranties as signals under consumer moral hazard". *The Rand Journal of Economics*, pp. 239–255.

Macht, Michael and Dorothee Dettmer (2006). "Everyday mood and emotions after eating a chocolate bar or an apple". *Appetite* 46.3, pp. 332–336.

Martinelli, César (2006). "Would rational voters acquire costly information?" *Journal of Economic Theory* 129.1, pp. 225–251.

McCall, John Joseph (1970). "Economics of information and job search". *The Quarterly Journal of Economics*, pp. 113–126.

McWilliams, Bruce (2012). "Money-back guarantees: Helping the low-quality retailer". *Management Science* 58.8, pp. 1521–1524.

Milgrom, Paul R and Robert J Weber (1982). "A theory of auctions and competitive bidding". *Econometrica: Journal of the Econometric Society*, pp. 1089–1122.

Milgrom, Paul and John Roberts (1986). "Price and advertising signals of product quality". *The Journal of Political Economy*, pp. 796–821.

Moorthy, Sridhar and Kannan Srinivasan (1995). "Signaling quality with a money-back guarantee: The role of transaction costs". *Marketing Science* 14.4, pp. 442–466.

Mortensen, Dale T (1970). "Job search, the duration of unemployment, and the Phillips curve". *The American Economic Review*, pp. 847–862.

Mortensen, Dale T and Christopher A Pissarides (1999). "New developments in models of search in the labor market". *Handbook of Labor Economics* 3, pp. 2567–2627.

Nelson, Randy J (2005). *An introduction to behavioral endocrinology* . Sinauer Associates.

Nosko, Chris and Steven Tadelis (2015). "The limits of reputation in platform markets: An empirical analysis and field experiment".

O'Connell, Matthew and Mei-Chuan Kung (2007). "The Cost of Employee Turnover." *Industrial Management* 49.1.

Oyer, Paul and Scott Schaefer (2011). "Personnel economics: hiring and incentives". *Handbook of Labor Economics* 4, pp. 1769–1823.

Pallais, Amanda (2014). "Inefficient Hiring in Entry-Level Labor Markets". *American Economic Review* 104.11, pp. 3565–99.

Pallais, Amanda and Emily Glassberg Sands (2013). "Why the Referential Treatment? Evidence from Field Experiments on Referrals".

Payne, John W et al. (1999). "Measuring constructed preferences: Towards a building code". *Elicitation of Preferences*. Springer, pp. 243–275.

Persico, Nicola (2004). "Committee design with endogenous information". *The Review of Economic Studies* 71.1, pp. 165–191.

Peteraf, Margaret A and Jay B Barney (2003). "Unraveling the resource-based tangle". *Managerial and Decision Economics* 24.4, pp. 309–323.

Pissarides, Christopher A (1994). "Search unemployment with on-the-job search". *The Review of Economic Studies* 61.3, pp. 457–475.

Priest, George L (1981). "A theory of the consumer product warranty". *The Yale Law Journal* 90.6, pp. 1297–1352.

Rabin, Matthew (1993). "Incorporating fairness into game theory and economics". *The American Economic Review*, pp. 1281–1302.

Reed, Richard and Robert J DeFillippi (1990). "Causal ambiguity, barriers to imitation, and sustainable competitive advantage". *Academy of Management Review* 15.1, pp. 88–102.

Rees, Albert (1966). "Information networks in labor markets". *The American Economic Review*, pp. 559–566.

Rind, Bruce and Prashant Bordia (1996). "Effect on restaurant tipping of male and female servers drawing a happy, smiling face on the backs of customers' checks". *Journal of Applied Social Psychology* 26.3, pp. 218–225.

Rosen, Sherwin (1982). "Authority, control, and the distribution of earnings". *The Bell Journal of Economics*, pp. 311–323.

Russo, Giovanni, Cees Gorter, and Ronald Schettkat (2001). "Searching, hiring and labour market conditions". *Labour Economics* 8.5, pp. 553–571.

Russo, Giovanni, Piet Rietveld, et al. (2000). "Recruitment channel use and applicant arrival: An empirical analysis". *Empirical Economics* 25.4, pp. 673–697.

Sanfey, Alan G et al. (2003). "The neural basis of economic decision-making in the ultimatum game". *Science* 300.5626, pp. 1755–1758.

Sattinger, Michael (1993). "Assignment models of the distribution of earnings". *Journal of Economic Literature*, pp. 831–880.

Saunders, Edward M (1993). "Stock prices and Wall Street weather". *The American Economic Review* 83.5, pp. 1337–1345.

Schettkat, Ronald (1996). "Labor Market Flows Over the Business Cycle: An Asymmetric Hiring Cost Explanation". *Journal of Institutional and Theoretical Economics* 152.4, pp. 641–653.

Shieh, Shiou (1996). "Price and Money-Back Guarantees as Signals of Product Quality". *Journal of Economics & Management Strategy* 5.3, pp. 361–377.

Spence, Michael (1977). "Consumer misperceptions, product failure and producer liability". *The Review of Economic Studies* 44.3, pp. 561–572.

Spies, Kordelia, Friedrich Hesse, and Kerstin Loesch (1997). "Store atmosphere, mood and purchasing behavior". *International Journal of Research in Marketing* 14.1, pp. 1–17.

Stanton, Christopher and Catherine Thomas (2012a). "Landing the first job: The value of intermediaries in online hiring". *Available at SSRN 1862109*.

— (2012b). "Landing the first job: the value of intermediaries in online hiring". *Available at SSRN 1862109*.

Stigler, George J (1961). "The economics of information". *The Journal of Political Economy*, pp. 213–225.

Teece, David J (2007). "Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance". *Strategic Management Journal* 28.13, pp. 1319–1350.

Van Ours, Jan C and Geert Ridder (1993). "Vacancy durations: search or selection?" *Oxford Bulletin of Economics and Statistics* 55.2, pp. 187–198.

Van Ours, Jan and Geert Ridder (1992). "Vacancies and the recruitment of new employees". *Journal of Labor Economics*, pp. 138–155.

Vickrey, William (1961). "Counterspeculation, auctions, and competitive sealed tenders". *The Journal of Finance* 16.1, pp. 8–37.

Watson, David and Jatin Vaidya (2003). "Mood measurement: Current status and future directions". *Handbook of Psychology*.

Wiener, Joshua Lyle (1985). "Are warranties accurate signals of product reliability?" *Journal of Consumer Research*, pp. 245–250.

Winslow, EG (1986). "Keynes and Freud: Psychoanalysis and Keynes's Account of the "Animal Spirits" of Capitalism". *Social Research*, pp. 549–578.

Appendix A

A.1 Additional Details on Hiring

Figure A.1 shows the timing of application, interviews, and hiring decisions relative to the time the job was posted. Clearly most applications are submitted soon after the job is posted, with the probability of an application being posted dropping steadily over time. Interviewing and hiring follow a bimodal distribution. Employers begin interviewing candidates shortly after the job is posted and then make job offers. Those who cannot come to an agreement with their first-choice applicant, then interview some more and make a 2nd round of offers.

When an employer "views" and application by clicking on it, the application is expanded and the first thing shown below the bid information already observable on the unexpanded list of applicants is the applicant's cover letter and the text responses to any questions posed in the job posting. Figure A.2 shows an example of this information. Directly below this is the "work history" of the applicant, which was shown above (in text) in figure 2.2.

A.2 Search and Screening

In this section I explore a series of robustness checks and additional specifications related to the effects of hiding past wage rate information on search and screening behavior.

Table A.1 uses a log-linear model to rerun the results provided in table 2.4. The quasi-poisson model is the preferred specification as only 37% of jobs Skype any applicants, and this model better adjusts for the number of zeros, which skew the distribution of the data.

Table A.2 uses the preferred quasi-Poisson specification to show that the results on information acquisition are similar regardless of the dependent variable used as a rough proxy for the amount of intensive information acquisition the employer engaged in.

As a further robustness check, table A.3 shows that running table 2.6 using a log-linear specification instead of a quasi-Poisson regression does not materially affect the direction of the results.

Table A.4 presents application level regressions, which indicate that the odds of asking a question conditional on sending a message is higher for treated employers than control employers, although not significantly so in all specifications. Interestingly, the results on

Figure A.1: PDF of Timing of Screening Behavior

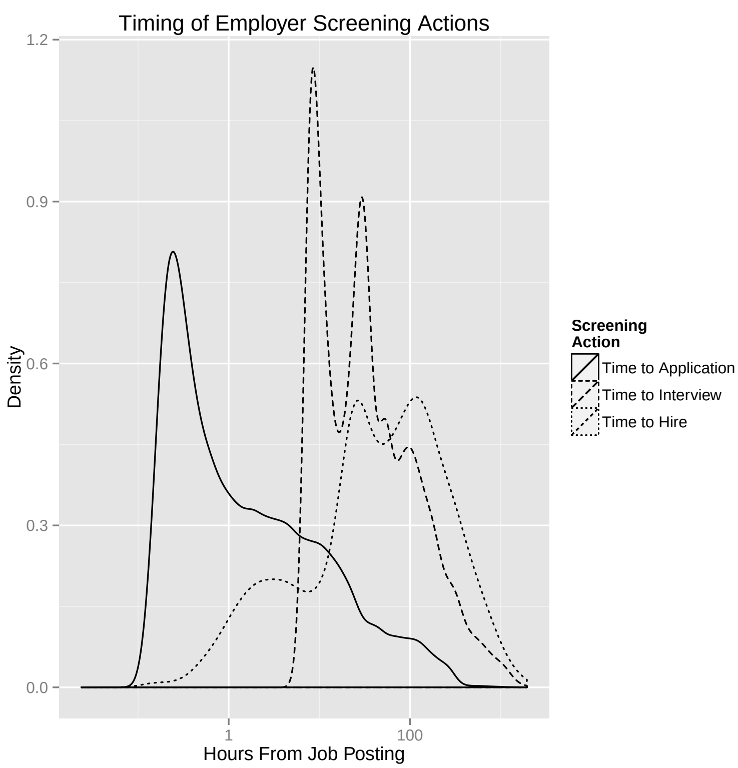

Note: This figure is limited to only non-invited applicants to public job postings.

Figure A.2: Expanded View of ATS Application

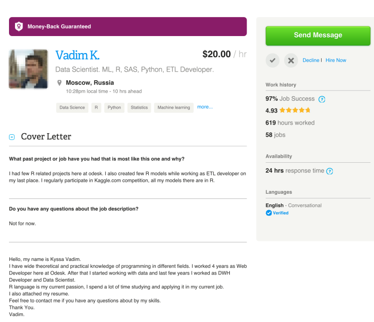

Note: Upon expanding an application, the employer also sees an expanded profile.

Table A.1: Search Intensity, Log Linear Models

	Dependent variable:
	Num Viewed OLS
Hidden Wages Treatment Group (TRT)	0.049*
	(0.026)
Constant	1.430***
	(0.120)
Opening Level Covariates	Yes
Observations	5,855

Notes: The sample is restricted to hourly first job posts by an employer. All models include covariates including: category indicators, prior jobs billed by the employer, employer's prior spendings on the platform, the number of applications to the job opening, the number of recommended applications to the job opening, and the average applicant's bid. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

odds of asking for a Skype ID are a tightly measured zero, indicating that employers have a tendency to ask either all or none of the applicants they message for Skype IDs.

A.3 Job Outcomes

Figure A.3 plots the distribution of wage-to-bid ratios that are not equal to one for the first contract signed for each treatment and control employer after the start of the experiment. Some employees manage to negotiate wages higher than their initial bid amount. Interviews with workers on oDesk reveal that for top level talent, it is not unheard of to use other offers to attempt to negotiate up the contract wage relative to the bid. Additionally, figure A.3 allows us to conclude that the treatment effect does not appear to be driven by outliers such as bidding only one cent on a job opening.

Figure A.3: Distribution of the Non-equal Wage-to-Bid Ratio by Treatment Group

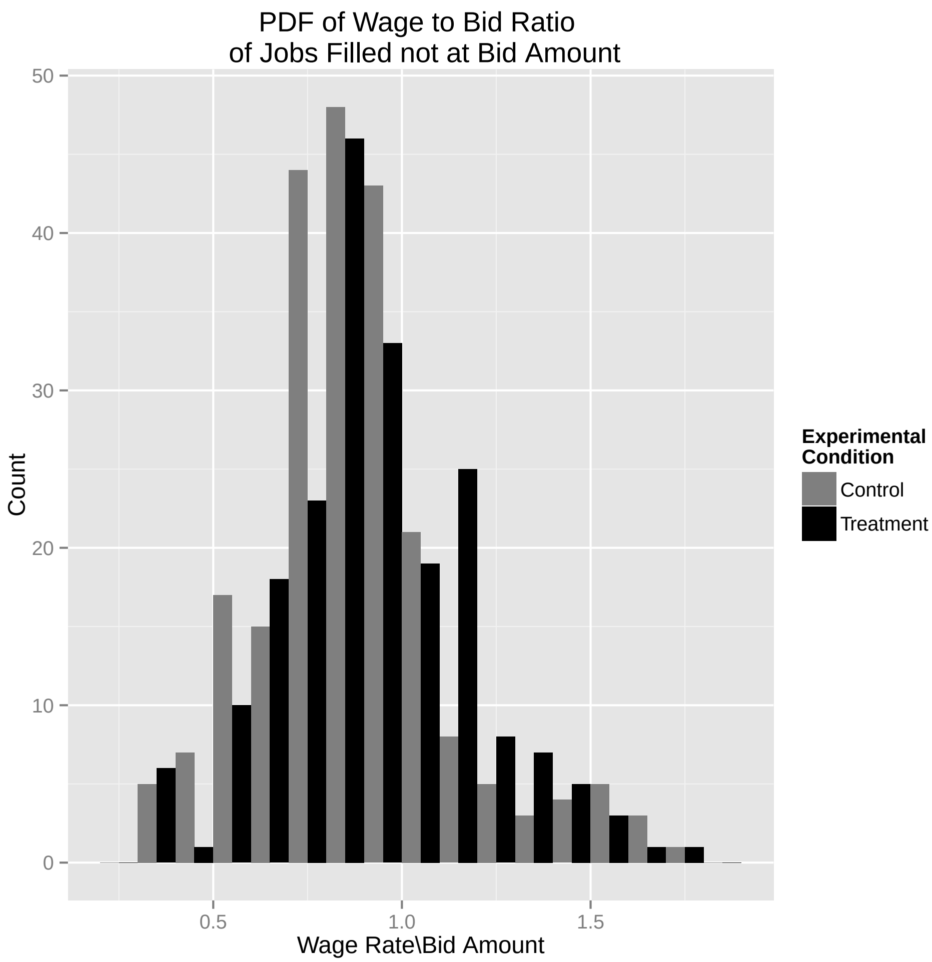

Note: This table plots the distributions on wage-to-bid ratio for treatment and control jobs. Wage-to-bid ratio is calculated as the ratio of the wages paid to the winning applicant to the hourly bid submitted by the winning applicant in that job. The top and bottom .5% of wage-to bid ratios were dropped. All wage-to-bid ratios equal to one were dropped.

Table A.2: Information Acquisition, Alternate DVs

	Dependent variable:		
	Total Number of Words	Total Number of Question Marks	Total Number of Question Words
	Poisson	Poisson	Poisson
	(1)	(2)	(3)
Hidden Wages			
Treatment Group (TRT)	0.064	0.080	0.082
	(0.066)	(0.060)	(0.059)
Constant	5.847***	0.923***	0.966***
	(0.285)	(0.255)	(0.260)
Opening Level Covariates	Yes	Yes	Yes
Observations	5,855	5,855	5,855

Notes: Estimates are from quasi-maximum likelihood Poisson models. Heteroskedasticity-robust standard errors are reported. The sample is restricted to hourly first job posts by an employer. All models include covariates: category indicators, prior jobs billed by the employer, employer's prior spending on the platform, the number of applications to the job openings, the number of recommended applications to the job opening. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : * * *$.

Table A.3: Information Acquisition, Log Linear Models

	Dependent variable:			
	log1p(Apps Messaged) OLS	log1p(Apps Skyped) OLS	log1p(Apps with "?") OLS	log1p(Question Words) OLS
	(1)	(2)	(3)	(4)
Hidden Wages				
Treatment Group (TRT)	0.016	0.012	0.024	0.040***
	(0.020)	(0.025)	(0.015)	(0.015)
Constant	1.255***	0.181	0.926***	0.824***
	(0.096)	(0.144)	(0.073)	(0.075)
Opening Level Covariates	Yes	Yes	Yes	Yes
Observations	3,285	2,195	2,799	2,678

Notes: This table shows the relationship between measures of information acquisition and the treatment using log-linear regressions. The level of observation is the job posting. Estimates are from quasi-maximum likelihood Poisson models. Heteroskedasticity-robust standard errors are reported. The sample is restricted to hourly first job posts by an employer. All models include covariates: category indicators, prior jobs billed by the employer, employer's prior spending on the platform, the number of applications to the job opening, the number of recommended applications to the job opening, and the average applicant's bid. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Table A.4: Effect of Hiding Past Wages on Information Acquisition, Application Level

	Dependent variable:			
	I(Applicant Messaged)	I(Applicant Skyped)	I(Applicant Qword)	I(Applicant Qmark)
	(1)	(2)	(3)	(4)
Hidden Wages				
Treatment Group (TRT)	0.088	−0.004	0.182**	0.096
	(0.055)	(0.080)	(0.086)	(0.095)
Constant	−2.020***	−0.783**	0.772**	0.717*
	(0.214)	(0.360)	(0.378)	(0.415)
Opening Level Covariates	Yes	Yes	Yes	Yes
Observations	186,635	10,541	10,541	10,541

Notes: This table shows the relationship between measures of information acquisition and the treatment. The level of observation is the application to a job posting. Estimates are from logit models. Heteroskedasticity-robust standard errors clustered at the employer level are reported. The sample is restricted to hourly first job posts by an employer. All models include opening level covariates including a category indicator, employer past earnings, employer past jobs, number of applications, number of recommended applications, and requested skills. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.

Appendix B

B.1 Characteristics of MBG-eligible Applicants

In table B.1 I compare the characteristics of applicants who are eligible for the money-back guarantee (BestMatch $\geq .5$) and applicants who are not eligible for the money-back guarantee (BestMatch $< .5$). Clearly applicants selected by the algorithm have substantially more experience, feedback, and charge more for their services.

In table B.2, I show that the applicants who were deemed eligible for the money-back guarantee looked no different between treatment and control jobs posted by treatment and control employers.

Table B.1: Characteristics of MBG-eligible Applicants vs. Non-eligible Applicants

	Non-Eligible mean: $\bar{X}_{CTL}$	Eligible mean: $\bar{X}_{TRT}$	Difference In Means	p-value	
Applicant Attributes					
Hours Worked to Date	583.23	1114.88	531.65 (40.09)	<0.001	***
Num Past Jobs Worked	11.97	27.60	15.63 (0.96)	<0.001	***
Past Hourly Earnings	7238.29	20002.97	12764.69 (762.99)	<0.001	***
Num Prior Employers	9.75	20.81	11.05 (0.70)	<0.001	***
Min Feedback Rating	0.26	0.62	0.36 (0.01)	<0.001	***
Wage Bid	15.08	22.72	7.64 (0.58)	<0.001	***
Profile Wage	13.76	20.93	7.17 (0.46)	<0.001	***

Notes: This table reports means and standard errors across eligibility for the money-back guarantee. Standard errors are clustered at the employer level. Significance indicators: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

Table B.2: Balance Table of MBG-Eligible Applicant Characteristics

	Control Mean: $\bar{X}_{CTL}$	Treatment Mean: $\bar{X}_{TRT}$	Difference In Means	p-value
Applicant Attributes				
Hours Worked to Date	1080.93	1150.05	69.12 (82.69)	0.40
Num Past Jobs Worked	27.45	27.76	0.31 (2.11)	0.88
Past Hourly Earnings	20645.51	19337.12	-1308.40 (1664.50)	0.43
Num Prior Employers	20.42	21.20	0.78 (1.53)	0.61
Min Feedback Rating	0.62	0.61	-0.01 (0.01)	0.23
Wage Bid	22.41	23.05	0.64 (1.26)	0.61
Profile Wage	21.01	20.85	-0.17 (1.06)	0.87

Notes: This table reports means and standard errors across experimental groups for characteristics of applicants who were eligible for the money-back guarantee. Standard errors are clustered at the employer level. Significance indicators: $p \leq 0.10 : \dagger$, $p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

Appendix C

C.1 Eventual eBay Winners vs. Non-Winners

Table C.1 shows that in my sample those bidders who eventually win an auction bid differently from those bidders who fail to ever win an auction in my sample. This is worrying, as my results are based on a selected sample of those bidders who eventually win an auction as I only obtain zip code information and thus weather information for those bidders who win.

C.2 Alt. DV - No Shipping

In table C.2, I show that the regression results are robust to specifications where shipping is not included in the calculation of willingness to pay.

Table C.1: Winners vs. Non Winners

	Never Winner Mean: $\bar{X}_{CTL}$	Eventual Winner Mean: $\bar{X}_{TRT}$	Difference In Means	p-value	
Target Gift Card					
Number of Observations	25125	30688			
Avg Bid	123.78	91.87	-31.91 (1.08)	<0.001	***
Avg Bid + Shipping	124.36	92.53	-31.83 (1.08)	<0.001	***
Bid/Amazon Price	0.88	0.60	-0.28 (0.00)	<0.001	***
Min to Auction End	2574.81	3175.04	600.23 (24.65)	<0.001	***
Windows Software					
Number of Observations	55914	27978			
Avg Bid	105.50	69.89	-35.61 (0.37)	<0.001	***
Avg Bid + Shipping	110.94	75.51	-35.43 (0.37)	<0.001	***
Bid/Amazon Price	0.34	0.24	-0.10 (0.00)	<0.001	***
Min to Auction End	1571.67	2731.15	1159.48 (20.97)	<0.001	***
Apple iPod					
Number of Observations	26826	9602			
Avg Bid	54.80	40.30	-14.50 (0.29)	<0.001	***
Avg Bid + Shipping	59.97	45.63	-14.34 (0.30)	<0.001	***
Bid/Amazon Price	0.81	0.64	-0.18 (0.00)	<0.001	***
Min to Auction End	519.98	1189.92	669.94 (23.06)	<0.001	***
TI 83 Calculator					
Number of Observations	7376	3055			
Avg Bid	37.60	31.01	-6.59 (0.31)	<0.001	***
Avg Bid + Shipping	43.27	36.18	-7.09 (0.31)	<0.001	***
Bid/Amazon Price	0.91	0.74	-0.17 (0.01)	<0.001	***
Min to Auction End	215.30	561.14	345.84 (25.59)	<0.001	***
Armani Perfume					
Number of Observations	15839	6529			
Avg Bid	31.43	23.55	-7.88 (0.15)	<0.001	***
Avg Bid + Shipping	37.61	29.63	-7.98 (0.15)	<0.001	***
Bid/Amazon Price	0.79	0.62	-0.17 (0.00)	<0.001	***
Min to Auction End	863.42	2147.45	1284.03 (34.60)	<0.001	***

Notes: This table reports means and standard errors for bidding behavior for each item in the dataset for never-winners and eventual winners. Observations are at the individual-auction level. Only the highest bid for an individual in each auction is kept. Reported p-values are the for two-sided t-tests of the null hypothesis of no difference in means across groups. In the bottom panel, standard errors are clustered at the employer level. Significance indicators: $p \leq 0.10$:† ,$p \leq 0.05 : *$, $p \leq 0.01 : **$ and $p \leq .001 : ***$.

Table C.2: Willingness to Pay - Consumer Goods, No Shipping

	Dependent variable:		
	Bid		
	(1)	(2)	(3)
Rain Indicator	2.056^{***}	1.716^{*}	1.091^{*}
	(0.386)	(0.889)	(0.644)
Opening Level Covariates	Yes	Yes	Yes
Month-Year FE	No	Yes	Yes
Item FE	No	Yes	Yes
Bidder FE	No	No	Yes
Observations	47,110	47,110	47,110

Notes: Covariates include: the minimum bid, the bidder's previous bid count, weekend indicator, number of simultaneous auctions, seller feedback, indicator for reserve, indicator for bold auction title, indicator for picture in auction. All models are OLS with robust standard errors clustered at the zip code level. Significance indicators: $p \leq 0.10 : *$, $.p \leq 0.05 : **$ and $p \leq .01 : ***$.